Model-Based Calibration Toolbox™ 3
Getting Started Guide

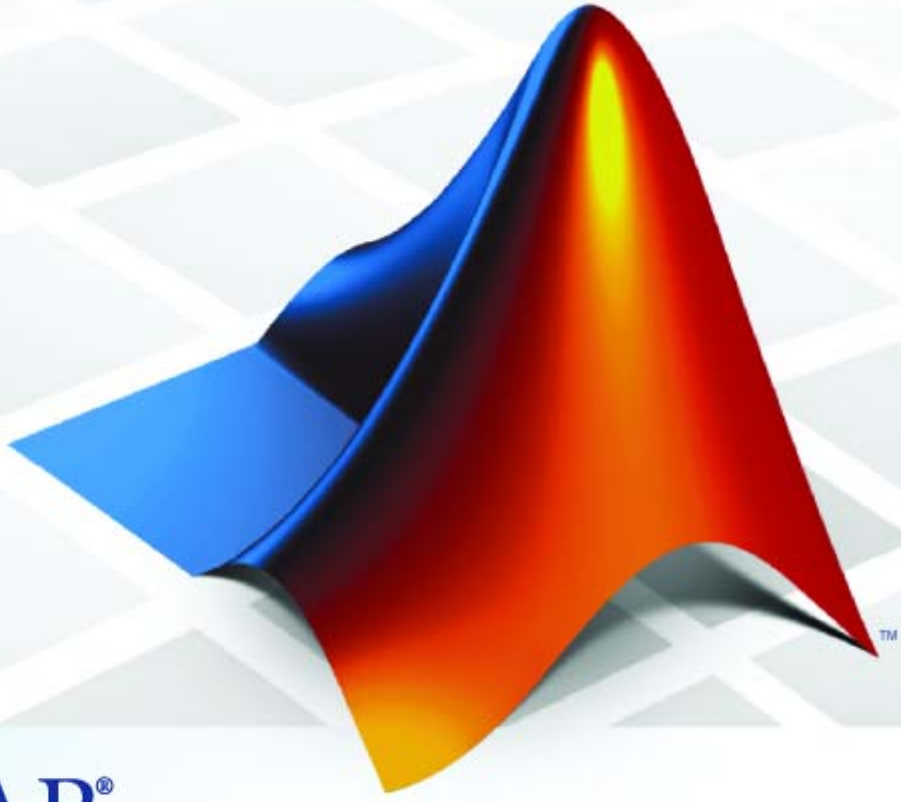

MATLAB®
&SIMULINK®

The MathWorks™
Accelerating the pace of engineering and science

How to Contact The MathWorks

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

`suggest@mathworks.com`	Product enhancement suggestions
`bugs@mathworks.com`	Bug reports
`doc@mathworks.com`	Documentation error reports
`service@mathworks.com`	Order status, license renewals, passcodes
`info@mathworks.com`	Sales, pricing, and general information

508-647-7000 (Phone)

508-647-7001 (Fax)

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

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

Model-Based Calibration Toolbox™ Getting Started Guide

© COPYRIGHT 2005–2009 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

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

Revision History

November 2005	Online only	New for Version 3.0 (Release 14SP3+)
September 2006	Online only	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)

Contents

1 Introduction

Product Overview 1-2
About the Model Browser 1-2
About CAGE 1-3

How to Use This Manual 1-5
Learning the Model-Based Calibration Toolbox Product . . 1-5
Learning the Model Browser 1-5
Learning CAGE 1-7
Training Material 1-7

System Requirements 1-8
Operating System Requirements 1-8
Required MathWorks Products 1-8
Optional MathWorks Products 1-8

Limitations 1-10
Waitbar May Have Transparent Background 1-10

2 Gasoline Engine Calibration Case Study

Gasoline Case Study Overview 2-3
Case Study Introduction 2-3
Why Use Design of Experiment and Engine Modeling? . . . 2-4
Problem Definition 2-5
Introduction to Two-Stage Modeling 2-6

Designing the Experiment 2-7
Overview of Design Process 2-7
Creating a Test Plan 2-7
Specifying the Model Inputs 2-8

Creating Designs 2-12
Data Source 2-17

Importing and Filtering Data 2-18
Import and View Data 2-18
Filter Data 2-24

How Is a Two-Stage Model Constructed? 2-27

Building the Models 2-31
Completing Model Setup 2-31
Specifying the Local Model Type 2-32
Selecting Data and Responses to Model 2-33

Selecting Local Models 2-40

Creating Boundary Models 2-44

Selecting Global and Two-Stage Models 2-49
Inspect the Global Models 2-49
Create Multiple Models to Compare 2-51
Create a Two-Stage Model 2-55
Adding New Response Models 2-57

Using Validation Data 2-60

Exporting the Models 2-63

Optimized Calibration 2-64
Problem Definition 2-64
Benefits of Automated Calibration 2-65

Importing Models into CAGE 2-67

Setting Up Calibration Tables to Fill 2-70
Altering Variable Ranges 2-70
Setting Up Tables 2-70

Setting Up the Optimization **2-73**

Defining Variable Values **2-78**

Running the Optimization **2-82**

Setting Up the Sum Optimization **2-85**
Setting Up the Optimization Initial Values and Objective **2-85**
Creating Table Gradient Constraints **2-88**
Running the Sum Optimization **2-91**

Filling Tables with Optimization Results **2-92**

MBT Spark Estimator Problem **2-94**
What Is an Estimator? **2-94**
View the Feature **2-94**
View Variables **2-95**
Edit and Import Boundary Model **2-96**
Use the Feature Fill Wizard **2-97**
Inspect Results **2-101**
CAGE Import Tool **2-104**

3 Diesel Engine Calibration Case Study

Diesel Case Study Overview **3-2**
Introduction **3-2**
Problem Definition **3-2**

Design of Experiment **3-4**
Introducing Design of Experiment **3-4**
Constraining the Design **3-7**
Creating Candidate Designs **3-14**
Data Collection **3-15**

Modeling **3-17**

Overview of Modeling Process 3-17
Building Models 3-18
Building and Evaluating Alternative Models 3-20
Creating Boundary Models 3-26
Export to CAGE 3-29

Optimized Calibration 3-31
Problem Definition 3-31
Benefits of Automated Calibration 3-32

Importing Models of Engine Responses into CAGE . . . 3-34

Defining Additional Variables and Models 3-35

Setting Up Calibration Tables to Fill 3-37

Setting Up the Optimization 3-38

Setting Up Constraints 3-42

Defining Variable Values 3-47

Running the Optimization 3-50

Setting Up the Sum Optimization 3-52
Setting Up Initial Values and Sum Objective 3-52
Creating the Brake Specific NOx Constraint 3-54
Setting Weights for the Sum Objective and Constraint . . . 3-55
Set Parameters and Run Optimization 3-57

Filling Tables with Optimization Results 3-59

4 Command-Line Interface to the Model-Based Calibration Toolbox Product

Introduction to the Command-Line Interface 4-2

Processes You Can Automate 4-3

Understanding Model Structure 4-6
Overview of Model Structure 4-6
Projects 4-6
Test Plans 4-6
Responses 4-6

How the Model Tree Relates to Command-Line Objects 4-9

Command-Line Demos 4-14

5 Tutorial: Model Quickstart

Two-Stage Models 5-2

Starting the Toolbox and Loading Data 5-4

Setting Up the Model 5-8
Creating the Test Plan 5-8
Setting Up the Local Model 5-10
Setting Up the Global Model 5-14
Selecting Data 5-18
Specifying the Response Model 5-22

Verifying the Model 5-24
Verifying the Local Model 5-24
Verifying the Global Model 5-26
Selecting the Two-Stage Model 5-28

Comparing the Local Model and the Two-Stage Model ... 5-35
Maximum Likelihood Estimation ... 5-36
Response Node ... 5-39

Exporting the Model ... 5-41

Creating Multiple Models to Compare ... 5-43
Methods For Creating More Models ... 5-43
Creating New Local Models ... 5-43
Adding New Response Features ... 5-46
Comparing Models ... 5-47
Creating New Global Models ... 5-49
Creating Multiple Models Using Build Models ... 5-52

6 Tutorial: Design of Experiment

What Is Design of Experiment? ... 6-2
Why Use Design of Experiment? ... 6-2
Design Styles ... 6-3
Structure of This Design Editor Tutorial ... 6-4

Getting Started with the Design Editor ... 6-6
Setting Up a Model ... 6-6
Starting the Design Editor ... 6-7
Creating a New Design ... 6-8

Creating Optimal Designs ... 6-10
Introducing Optimal Designs ... 6-10
Start Point Tab ... 6-11
Candidate Set Tab ... 6-12
Algorithm Tab ... 6-15

Viewing Design Displays ... 6-17

Using the Prediction Error Variance Viewer ... 6-21
Introducing the Prediction Error Variance Viewer ... 6-21
Improving the Design ... 6-24

Creating Classical Designs **6-27**
Adding a Classical Design **6-27**
Classical Design Browser **6-29**
Setting Up and Viewing a Classical Design **6-30**

Using the Design Evaluation Tool **6-34**

Creating Space-Filling Designs **6-37**

Applying Constraints **6-39**

Saving Designs **6-45**

7 Tutorial: Data Editor

Introduction to the Data Editor **7-2**

Loading the Data **7-3**
Entering the Data Editor **7-3**
Loading a Data File **7-4**

Viewing and Editing the Data **7-6**
Viewing Data **7-6**
Using Notes to Sort Data for Plotting **7-8**
Removing Outliers and Problem Tests **7-9**
Reordering and Editing Data **7-10**

User-Defined Variables and Filtering **7-13**
Adding New Variables **7-13**
Applying a Filter **7-14**
Sequence of Variables **7-16**
Deleting and Editing Variables and Filters **7-17**

Storage **7-18**

Test Groupings **7-20**

Matching Data to Designs **7-24**
Introducing Matching Data to Designs **7-24**
Tolerances and Cluster Information **7-27**
Understanding Clusters **7-30**

8 Tutorial: Feature Calibration

What Are Feature Calibrations? **8-2**

Setting Up Calibrations **8-4**
Starting CAGE **8-4**
Setting Up Variables **8-4**
Setting Up Models **8-7**

Setting Up a Feature Calibration **8-10**
Setting Up a New Feature **8-10**
Setting Up the Strategy **8-12**
Setting Up the Tables **8-13**

Calibrating a Feature **8-16**
Process For Feature Calibration **8-16**
Calibrating the Normalizers **8-17**
Calibrating the Tables **8-21**
Calibrating the Feature **8-28**

Exporting Calibrations **8-33**

9 Tutorial: Tradeoff Calibration

What Is a Tradeoff Calibration? **9-2**

Setting Up a Tradeoff Calibration **9-3**
Creating a Tradeoff **9-3**

Adding Tables to a Tradeoff Calibration 9-5
Displaying the Models 9-6

Performing the Tradeoff Calibration 9-8
Process Overview 9-8
Checking the Normalizers 9-10
Setting Values for Other Variables 9-11
Filling Key Operating Points 9-13
Filling the Table by Extrapolation 9-17
Exporting Calibrations 9-19

10 Tutorial: Data Sets

Setting Up the Data Set 10-2
Tutorial Overview 10-2
Opening an Existing Calibration 10-3
Importing Experimental Data into a Data Set 10-3
Adding an Item to a Data Set 10-6

Comparing the Items in a Data Set 10-8
Viewing the Data Set as a Table 10-8
Viewing the Data Set as a Plot 10-10
Displaying the Error 10-11
Coloring the Display 10-13

Reassigning Variables 10-16

11 Tutorial: Filling Tables from Data

Setting Up a Table and Experimental Data 11-2
Tutorial Overview 11-2
Adding Variables 11-3
Adding a New Table 11-4
Importing Experimental Data 11-6

Filling the Table from the Experimental Data **11-9**

Selecting Regions of the Data **11-13**

Exporting the Calibration **11-15**

12 Tutorial: Optimization and Automated Tradeoff

Tutorial Overview **12-2**
Tutorial Problem Definitions **12-2**
The Optimization View **12-2**
Setting Up Your Optimization Session **12-4**

Single-Objective Optimization **12-5**
Process Overview **12-5**
Using the Optimization Wizard **12-5**
Setting Objectives, Constraints, and Operating Points . . . **12-8**
Running the Optimization **12-13**
Using Optimization Results to Fill Tables **12-14**
Using a Custom Fill Routine to Fill Tables **12-17**

Multiobjective Optimization **12-19**
Setting Up and Running the Multiobjective Optimization **12-19**
Optimization Output View **12-23**
Selecting Best Solutions **12-29**

Sum Optimization **12-33**

Automated Tradeoff **12-38**

Worked Example Optimization **12-41**
Example Overview **12-41**
Using the Worked Example Optimization **12-42**

Creating an Optimization from Your Own Algorithm . . **12-48**

Process Overview 12-48
Step 1: Verify the Algorithm 12-50
Step 2: Create a CAGE Optimization Function 12-52
Step 3: Define the Optimization Options 12-55
Step 4: Add the Algorithm to the Optimization Function . . 12-58
Step 5: Register Your Optimization Function with CAGE 12-62
Step 6: Verify Your New Optimization 12-63

Examples

A

Case Studies **A-2**

Commandline MBC **A-2**

Model Browser **A-2**

CAGE **A-2**

Index

1

Introduction

The following sections introduce Model-Based Calibration Toolbox™ software.

- “Product Overview” on page 1-2
- “How to Use This Manual” on page 1-5
- “System Requirements” on page 1-8
- “Limitations” on page 1-10

Product Overview

This section contains the following topics:

In this section...
“About the Model Browser” on page 1-2
“About CAGE” on page 1-3

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

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

The Model Browser part of the toolbox is a powerful tool for experimental design and statistical modeling. The models you build with the Model Browser can be imported into the CAGE Browser part of the toolbox to produce optimized calibration tables.

High accuracy engine models are a key component for reducing calibration effort and engine development time.

The time spent calibrating an engine control unit has been increasing, due to new control actuators. The new actuators give the potential for increased performance, reduced emissions, and improved fuel consumption. It is necessary to apply advanced modeling and optimization techniques to achieve the full benefits available from the addition of new actuators. Advanced modeling techniques offer increased understanding of the complex, nonlinear engine responses. High accuracy models can be used throughout the design process, including the calibration of base maps with the optimal settings for the control parameters, determined by constrained optimizations.

About the Model Browser

The Model Browser is a flexible, powerful, intuitive graphical interface for building and evaluating experimental designs and statistical models:

- Design of experiment tools can drastically reduce expensive data collection time.
- You can create and evaluate optimal, space-filling, and classical designs, and constraints can be designed or imported.
- Hierarchical statistical models can capture the nature of variability inherent in engine data, accounting for variation both within and between tests.
- The Model Browser has powerful, flexible tools for building, comparing, and evaluating statistical models and experimental designs.
- There is an extensive library of prebuilt model types and the capability to build user-defined models.
- You can export models to CAGE or to MATLAB® or Simulink® software.

Starting the Model Browser

To start the application, type

```
mbcmodel
```

at the MATLAB command prompt.

About CAGE

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

As engines get more complicated, and models of engine behavior more intricate, it is increasingly difficult to rely on intuition alone to calibrate lookup tables. CAGE provides analytical methods for calibrating lookup tables.

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 Model Browser models. From these models, CAGE builds steady-state ECU calibrations.

CAGE also compares lookup tables directly to experimental data for validation.

Starting the CAGE Browser

To start the application, type

```
cage
```

at the MATLAB command prompt.

How to Use This Manual

This section contains the following topics:

In this section...
"Learning the Model-Based Calibration Toolbox Product" on page 1-5
"Learning the Model Browser" on page 1-5
"Learning CAGE" on page 1-7
"Training Material" on page 1-7

Learning the Model-Based Calibration Toolbox Product

The case studies chapters contain step-by-step examples that guide you through using the whole toolbox to solve engine modeling and calibration problems:

- Chapter 2, "Gasoline Engine Calibration Case Study"
- Chapter 3, "Diesel Engine Calibration Case Study"

The following tutorial chapter explains how to get started using the command-line interface to the toolbox:

- Chapter 4, "Command-Line Interface to the Model-Based Calibration Toolbox Product"

Learning the Model Browser

These tutorials guide you through using specific Model Browser tools:

- Chapter 5, "Tutorial: Model Quickstart " provides a quick introduction to modeling with the toolbox. The tutorial describes how to set up and view a two-stage model using engine data.
- Chapter 6, "Tutorial: Design of Experiment" covers the tools for Design of Experiments with a step-by-step guide to setting up, viewing, and comparing one of each of the design types: classical, space-filling, and

optimal. The tutorial also describes how to define and apply constraints and export designs.

- Chapter 7, “Tutorial: Data Editor” is a guide to using the Data Editor to load and manipulate data for modeling. You can load data from files or the workspace or custom Microsoft® Excel® sheets. You can view plots of the data and define new variables and filters. You can store and import user-defined variables and filters, and define test groupings.

Learning CAGE

These tutorials guide you through using specific tools in CAGE:

- Chapter 8, "Tutorial: Feature Calibration" describes how to set up and calibrate lookup tables by reference to a model.
- Chapter 9, "Tutorial: Tradeoff Calibration" describes how to calibrate lookup tables using tradeoff calibrations.
- Chapter 10, "Tutorial: Data Sets" describes how to validate calibrations using experimental data.
- Chapter 11, "Tutorial: Filling Tables from Data" describes how to fill lookup tables using experimental data.
- Chapter 12, "Tutorial: Optimization and Automated Tradeoff" describes how to set up and run optimizations, including single objective, multiobjective, sum and user-defined optimization, and automated tradeoff.

Training Material

The files for the tutorial chapters are in the *matlabroot*`/toolbox/mbc/mbctraining` directory.

System Requirements

This section lists the following:

- Operating system requirements
- Required MathWorks products
- Optional MathWorks products

Operating System Requirements

The Model-Based Calibration Toolbox product is a Windows® 32-bit PC only product.

You can see the MATLAB system requirements online at `http://www.mathworks.com/products/system.shtml/Windows`

Required MathWorks Products

The Model-Based Calibration Toolbox product requires the following other MathWorks products:

- Simulink
- Optimization Toolbox™
- Statistics Toolbox™
- Symbolic Math Toolbox™

Optional MathWorks Products

The Model-Based Calibration Toolbox product can use the following MathWorks products:

- Parallel Computing Toolbox™
- Neural Network Toolbox™
- Genetic Algorithm and Direct Search Toolbox™

Note If you want to import Microsoft Excel files or use the custom Excel® file facility of the toolbox, you must also have the Microsoft Excel application.

Limitations

Waitbar May Have Transparent Background

The waitbar dialogs that are displayed to inform you of the progress of lengthy operations in the toolbox often display with transparent backgrounds. Instead of the normal window grey, they pick up their background from whatever is on the screen when they are displayed.

This bug is purely cosmetic and will not cause any data loss.

2

Gasoline Engine Calibration Case Study

This case study provides a step-by-step guide to using the Model-Based Calibration Toolbox product to solve a gasoline engine calibration problem. This section discusses the following topics:

- "Gasoline Case Study Overview" on page 2-3
- "Designing the Experiment" on page 2-7
- "Importing and Filtering Data" on page 2-18
- "How Is a Two-Stage Model Constructed?" on page 2-27
- "Building the Models" on page 2-31
- "Selecting Local Models" on page 2-40
- "Creating Boundary Models" on page 2-44
- "Selecting Global and Two-Stage Models" on page 2-49
- "Using Validation Data" on page 2-60
- "Exporting the Models" on page 2-63
- "Optimized Calibration" on page 2-64
- "Importing Models into CAGE" on page 2-67
- "Setting Up Calibration Tables to Fill" on page 2-70
- "Setting Up the Optimization" on page 2-73
- "Defining Variable Values" on page 2-78
- "Running the Optimization" on page 2-82

- "Setting Up the Sum Optimization" on page 2-85
- "Filling Tables with Optimization Results" on page 2-92
- "MBT Spark Estimator Problem" on page 2-94

Gasoline Case Study Overview

In this section...
"Case Study Introduction" on page 2-3
"Why Use Design of Experiment and Engine Modeling?" on page 2-4
"Problem Definition" on page 2-5
"Introduction to Two-Stage Modeling" on page 2-6

Case Study Introduction

This case study demonstrates how to systematically develop a set of optimal steady-state engine calibration tables using the Model-Based Calibration Toolbox product. This case-study uses a 2.2L inline 4 cylinder, naturally aspirated DOHC (Dual Overhead Cams) 4 valve per cylinder spark ignition (SI) engine equipped with dual-independent variable cam-phasing (DIVCP) hardware and electronic throttle.

Optimal steady-state engine calibration tables for intake cam phase, exhaust cam phase, and spark advance are developed as part of the case study process.

This example takes you through the following steps:

1. Create a design for your experiment — see "Designing the Experiment" on page 2-7.
2. Import the resulting data (taken using the design) to examine and filter it in preparation for modeling — see "Importing and Filtering Data" on page 2-18.
3. Make statistical models based on the data.
   a. "How Is a Two-Stage Model Constructed?" on page 2-27
   b. "Building the Models" on page 2-31
   c. "Selecting Local Models" on page 2-40
   d. "Creating Boundary Models" on page 2-44
   e. "Selecting Global and Two-Stage Models" on page 2-49

   f "Using Validation Data" on page 2-60

4. Export these models to the CAGE part of the toolbox to generate optimal calibration tables — see "Exporting the Models" on page 2-63.

5. The Model Browser section of the case study involves design of experiment, data handling, and model construction and export. In the CAGE browser section of the case study you use the models to complete the optimization of the calibration tables, see "Optimized Calibration" on page 2-64.

The following sections introduce the benefits of applying model-based calibration methods to solve this case study problem:

- "Why Use Design of Experiment and Engine Modeling?" on page 2-4
- "Problem Definition" on page 2-5
- "Introduction to Two-Stage Modeling" on page 2-6

Why Use Design of Experiment and Engine Modeling?

These approaches can be used to ensure that you develop optimal engine calibrations for complex engines with many controllable variables (such as variable valve timing, variable valve lift, and cylinder deactivation) at minimum cost and time.

Test bed time is expensive, and Design of Experiments methodology can help you choose the most effective points to run to get the maximum information in the shortest time. You can break the exponential dependency between the complexity of the engine (number of inputs) and the cost of testing (number of tests). You can collect the most statistically useful data, and just enough of it to fit the models.

Experimental design test points can be constrained based on previous experience to avoid damaging expensive engine hardware prototypes at unrealistic operating points.

The act of statistically modeling engine data can help identify the effect of interactions between calibration settings and engine performance, which can be vital to understanding how to optimally meet emissions constraints.

Accurate statistical models of engine data can also be used to develop calibration tables that have smooth transitions between the operating range of the engine and the edge regions of calibration tables where the engine will not be operated.

Optimal calibrations can be generated from statistical engine models in a methodical, repeatable process to ensure that maximum performance is achieved subject to emissions, driveability, and material limit constraints.

Problem Definition

The aim of this case study is to produce optimized tables for

- Intake Cam Phase
- Exhaust Cam Phase
- Spark Timing Schedules

  as a function of Load and rpm, subject to the following constraints

  - Constrain solutions to lie within the boundary constraint model (to keep the engine within its operating region)
  - Constrain cam phase solutions so they do not change by more than 10^o between table cells (that is, no more than 10^o per 500 RPM change and per 0.1 load change).
  - Constrain residual fraction <= 25% at each drive cycle point (to ensure stable combustion). Residual fraction is the percentage of burned gas mass in the cylinder at intake valve close, relative to the total mass in the cylinder at intake valve close. Constraining maximum residual fraction is a simple and reasonable way of ensuring stable combustion. Residual fraction = 100 * Burned Gas Mass from Last Cycle / (Burned Gas Mass From Last Cycle + Fresh Air Mass)

To produce these tables, you need to make accurate models of the behavior of torque, exhaust temperature, and residual fraction at different values of speed, throttle area, spark, and cam timings. You need engine data to build these models, so the first step is constructing an experimental design to collect the most useful set of points.

Before you can design an experiment you need to set up a two-stage test plan and define your model inputs and model type.

Introduction to Two-Stage Modeling

What is a two-stage test plan? You use a test plan to set up models in the Model Browser. The two stages refer to the way that engine data is often collected. For example, in each test, spark (the local variable) is swept while the other variables (such as speed and load) are held constant — these are referred to as global variables. Each test is taken at a different point in the global variables. Building the statistical models to take into account these individual sweeps makes it possible to incorporate engineering knowledge in the process. You can see plots of torque/spark sweeps, and use variables such as MBT (maximum brake torque) in modeling, rather than solely abstract mathematical properties of curves. You can then apply previous knowledge about the expected behavior of these variables to help you select good models.

You can easily identify outliers when you can see the sweep in which they were taken. The Model Browser allows you to visualize the data in a way that can help you identify and investigate suspect sweeps, and decide what kind of models will produce the best fit to the shapes of the data. The more controllable variables there are in an engine the more useful it is to have these visual aids to investigate complex data. Constructing models to take into account the way the data is collected helps build good models that you can have more confidence in. Statistically, it is the correct thing to do as it allows you to partition the errors within sweeps and the errors between sweeps separately.

You use a two-stage test plan to build your models because this data is suited to it. Spark is varied as the other variables are held constant, so the data is collected in a hierarchical structure; your models attempt to capture this information. You come to more detail on how this two-stage model is constructed after creating a design and obtaining data.

Designing the Experiment

In this section...
"Overview of Design Process" on page 2-7
"Creating a Test Plan" on page 2-7
"Specifying the Model Inputs" on page 2-8
"Creating Designs" on page 2-12
"Data Source" on page 2-17

Overview of Design Process

Creating a design in the Model Browser comprises several steps. You need to open the tool and create a new two-stage test plan. Then you need to enter the ranges and names of the input variables being used and choose a default model. Then you can create an initial design and set up the constraints on the input space. These constraints will be the same for all designs. From this constrained design, a series of child designs can be made with varying numbers of points added and slightly different models used. The final design can be chosen by comparing statistics of the various child designs and considering how many points you can afford to run. These steps are described next.

Creating a Test Plan

1. Start the Model Browser part of the toolbox by typing `mbcmodel` at the MATLAB command line.

2. From the startup project view, to create a new test plan, click **New** in the **Test Plans** list pane at the bottom.

   The **New Test Plan** dialog box appears.

3. Click the `Two-Stage` test plan icon

and click **OK**.

The default name of the new test plan, `Two-Stage`, appears in the Model Browser tree, in the **All Models** pane.

The view switches to the new test plan node in the tree, `Two-Stage`. The **Model Browser** window displays a diagram representing the two-stage model.

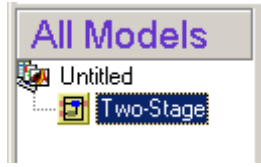

Specifying the Model Inputs

The models you are building are intended to predict the torque, fuel flow, and manifold pressure of the engine as a function of spark angle at specified operating points defined by the engine's speed, load, and cam timings. The input to the local model is the spark angle.

1 To specify spark angle as the local input, double-click the Local Inputs icon on the model diagram.

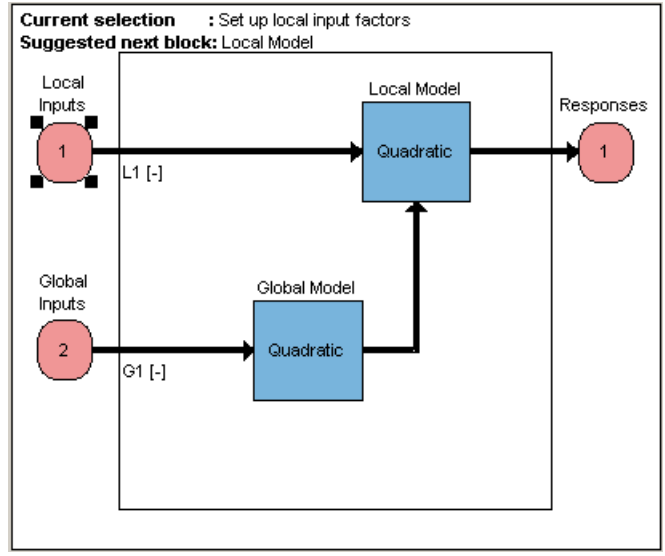

The Local Input Factor Setup dialog box appears.

a Set **Symbol** to `S`.

b Set **Signal** to `SPARK`. This is optional and matches the raw data.

c Set the range you want to model by changing **Max** to `50` (and leave **Min** at `0`).

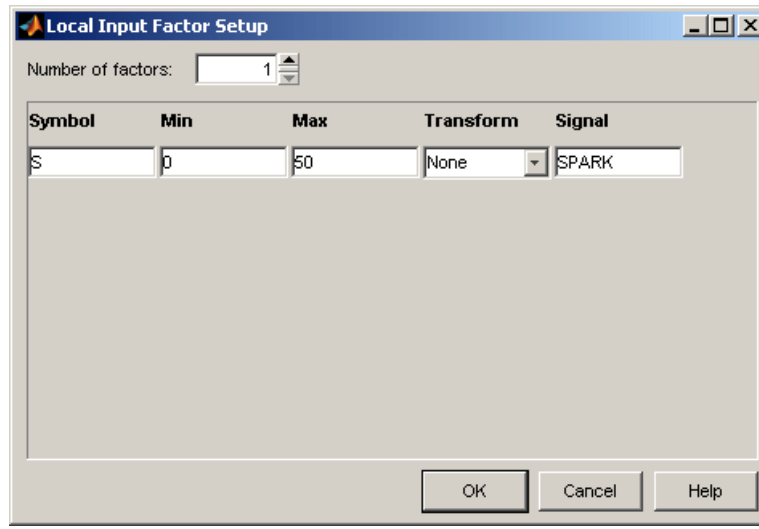

2 Click **OK** to dismiss the dialog box.

Notice that the new name of the local model input now appears on the two-stage model diagram.

The global inputs are the variables that are held constant at each operating point while spark is swept. In this case, these global variables are engine speed, scaled throttle area, intake cam angle, and exhaust cam angle.

3 To specify the global inputs, double-click the Global Inputs icon on the model diagram.

The Global Input Factor Setup dialog box appears.

By default, there is one input to the global model. Because this engine model has four input factors, you need to edit the input factors as follows:

a Click the up arrow button to increase the number of factors to four.

b Edit the four factors to create the engine model input. In each case, change the symbols, signal names, and ranges to the following:

Symbol	Signal	Min	Max
N	SPEED	500	6000
L	LOAD	0.05	0.95
ICP	INT_ADV	-5	50
ECP	EXH_RET	-5	50

Load = aircharge/maximum aircharge.

Cam angles are in units of degrees crankshaft, with intake values indicating advance from base timing, and exhaust values indicating retard from base timing.

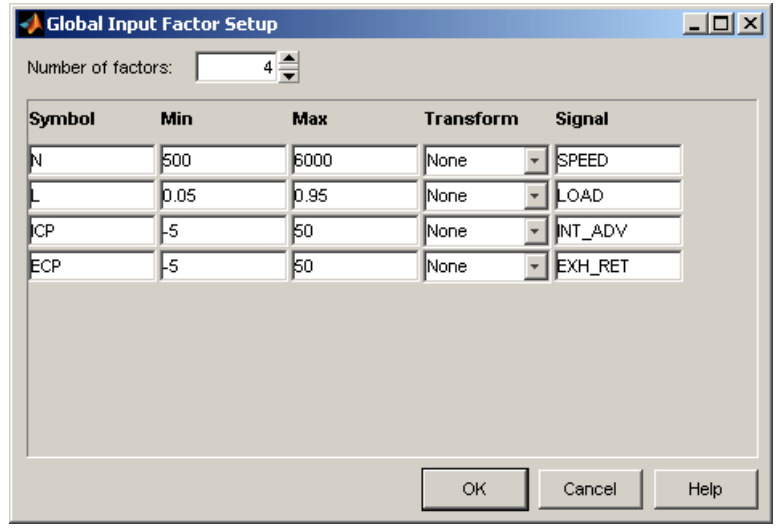

c Click **OK** to dismiss the dialog box.

4 To change the global model type, double-click the Global Model block in the two-stage model diagram. The Global Model Setup dialog box appears.

Deciding on the model to design for is vital for optimal designs only, when you already have some knowledge of how you expect the system to behave. In these cases, optimal designs can help you find the most efficient points for fitting the most robust models. In this case, you will create a space-filling design, which is best for exploring a new system where prior knowledge is low and you want to spread the available points to capture as much information as possible. These do not depend on model type; however, for this example you set a new model type now.

Remember that the statistical properties of different designs depend on the model type. For example, if you think you need cubic instead of quadratic in a factor, the number of points required rises dramatically and this has a highly adverse effect on the statistical quality of the designs. However, you do need to bear in mind that the final model will not be either of the possibilities listed here, because some terms will have been removed, or it might even be an RBF. You choose the most suitable model you can to construct a design, then when you have collected the data, you might find that a different model type produces the best fit.

5 `Polynomial` should already be selected from the **Linear model subclass** pop-up menu. Under **Model options**, verify the order for each of the four variables is two, to fit quadratic curves in each case.

6 Set **Stepwise** to `Minimize PRESS` (PREdicted Sum Square error).

This option will be important when you are fitting models to the data. You use the **Stepwise** feature to avoid overfitting the data; that is, you do not want to use unnecessarily complex models that "chase points" in an attempt to model random effects. Predicted error sum of squares (PRESS) is a measure of the predictive quality of a model. `Minimize PRESS` throws away terms in the model to improve its predictive quality, removing those terms that reduce the PRESS of the model. See "PRESS statistic" in the Model Browser documentation. You can also open the **Stepwise** window after model fitting to try to improve the fit with the **Stepwise** tools.

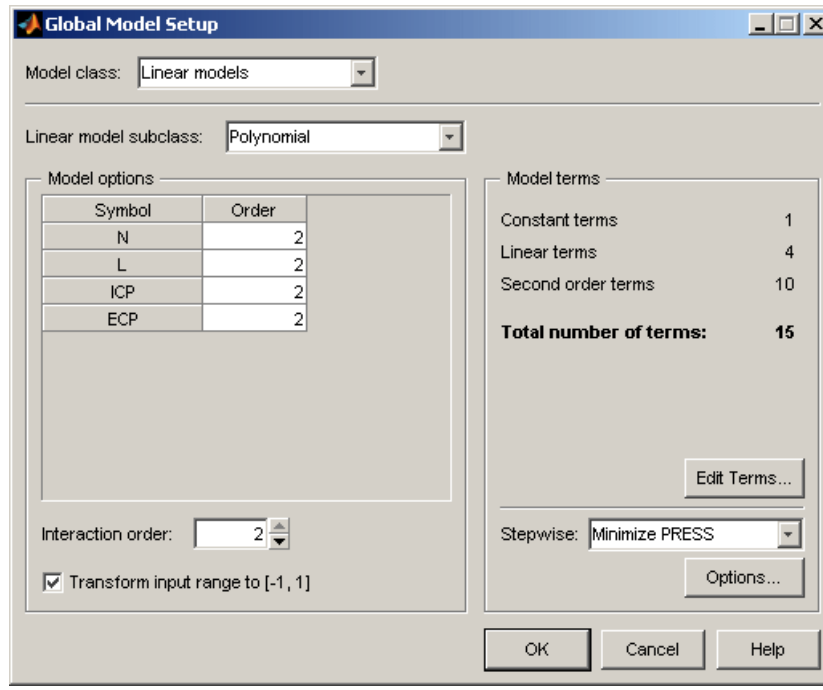

7 Click **OK** to dismiss the dialog box.

Creating Designs

Now you have set up the modeling test plan you can create an initial design and set up the constraints on the input space — these will be the same for all designs. From this constrained design, a series of child designs can be made with varying numbers of points added and slightly different models used. The final design can be chosen by comparing statistics of the various child designs and considering how many points you can afford to run.

1 Right-click the global model in the diagram and choose **Design Experiment**.

The Design Editor appears.

2 Click the button in the toolbar or select **File > New Design**. A new node called `Linear Model Design` appears.

The new `Linear Model Design` node is automatically selected. An empty Design Table (or any view you last used in the Design Editor) appears because you have not yet chosen a design.

3 Constrain the design space. Select **File > Import Constraints**. The Import Constraints dialog box appears.

4 In the **Import from** list, select `Boundary Constraints (.mat file)`. You will import a boundary model from an example file. In this way, you can use a boundary constraint from a previous investigation on a similar engine to constrain new designs.

5 Browse to the file `Gasoline_project.mat` in the `mbctraining` directory.

6 Click to select the `Boolean` type constraint as shown. This is the combination of both boundary constraint models.

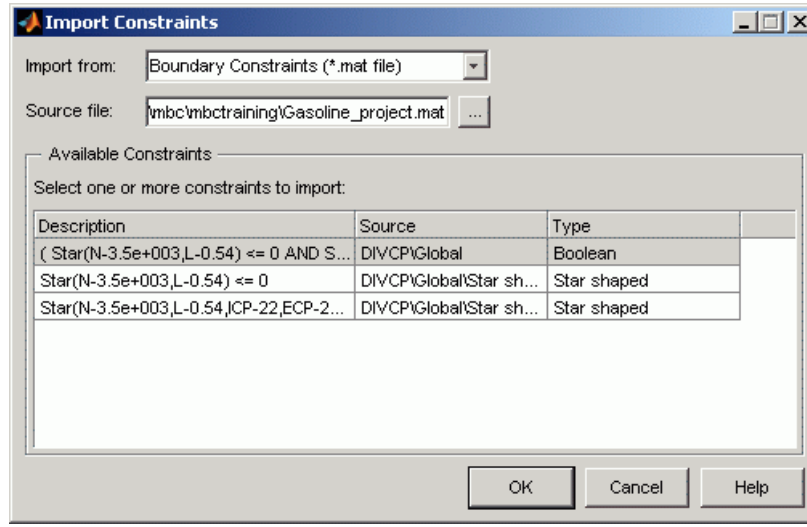

7 Click **OK** to import the boundary constraint.

Click **OK** in the following data matching dialog as all the signal names are automatically selected in this case.

8 Examine the constrained design space by right-clicking the title bar of a Design Table view and selecting **Current View > 3D Constraints**.

9 Select **Design > Space Filling > Design Browser**, or click the **Space Filling Design** button on the toolbar.

The Space Filling Design Browser appears.

Space-filling designs are best when there is little or no information about the underlying effects of factors on responses. For example, they are most useful when you are faced with a new type of engine, with little knowledge of the operating envelope. These designs do not assume a particular model form. The aim is to spread the points as evenly as possible around the operating space. Space-filling designs are also best for radial basis function models. You can use a mix-and-match approach: start with a space-filling design to survey the space, then continue testing with an optimal design once you have more understanding of the response and constraints. Once you have an idea of what model type will fit the response best, you can optimally add points in the most efficient places for the most robust model fit.

The most important thing to decide is how many design points you want. Testing is expensive and time-consuming, so you need to bear in mind how many points you have time for. When you consider the number of points, you also need to remember that a sweep will be done at each point and this will take some time. Do you need to allow time to fix problems or redo experimental points that can't be achieved?

10 Enter `800` for the **Number of points** and press **Enter**. A space-filling design is constructed, using the latin hypercube sampling method. Click the 3-D and 2-D tabs to examine the plots of new design distribution.

11 Click **Generate** to create a different design, and repeat until you achieve approximately 200 points in the **Size of constrained design** reported above the preview. This iteration is necessary because the space-filling design uses the whole variable space and some of the points will be removed if they fall outside the constraint. The preview is identical to the final design.

12 Click **OK** when you are satisfied with the design.

13 Right-click a view and select **Split View > 3D Design Projection** to view the design points.

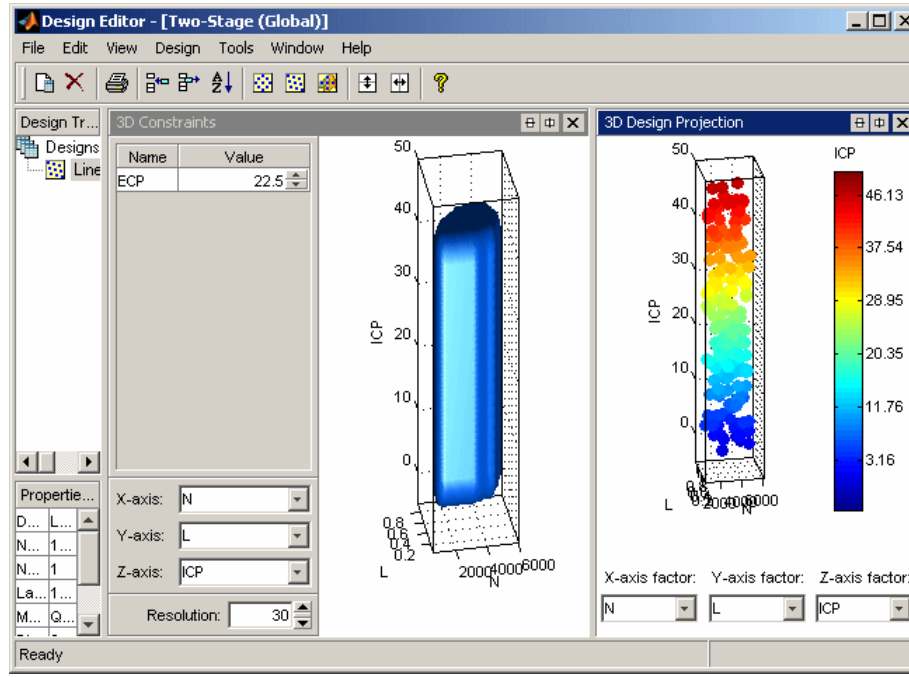

14 Add another space-filling design for some points with parked cam phasers. These points are important because we need an accurate model when cams are parked.

a Click the button in the toolbar to add a new design. The view switches to the new design in the tree. You could rename it `CAM_Parked`. The design inherits the same constraints as the parent design.

b Select **Design > Space Filling > Design Browser**, or click the **Space Filling Design** button on the toolbar.

In the dialog that appears, choose to replace the current points with a new design and click **OK**.

c Enter a number of points (try 40) and click **Generate** until you achieve a constrained design of about 10 points, and click **OK**.

d Display the design as a table, and edit the values of `ICP` and `ECP` in the new design to be zero.

15 Select **File > Merge Designs**. Select both your designs in the list, leave the **Merge to new design** option button selected, and click **OK**. Examine the merged design in the table view to confirm the parked cam points (`ICP` and `ECP` values of `0`) are at the end of the list of design points.

16 Add another space-filling design for collecting validation data.

   a Click the button in the toolbar to add a new design. The view switches to the new design in the tree. You could rename it `Validation`. The design inherits the same constraints as the parent design.

   b Select **Design > Space Filling > Design Browser**, or click the **Space Filling Design** button on the toolbar.

      In the dialog that appears, choose to replace the current points with a new design and click **OK**.

   c Enter a number of points (try 100) and click **Generate** until you achieve a constrained design of about 25 points, and click **OK**.

   d To add a point where the cams are parked, select **Edit > Add Point**.

      In the dialog that appears, select `User-specified` from the **Augment method** list, edit `ICP` and `ECP` to zero, and click **OK**.

You can use the Design Editor to make a selection of child designs to compare. When you have chosen the best design you can export it to file. In this case, you will import the example design for this case study. This design was used to collect the data for this case study, and you will later match these design points to data. Import the design as follows:

1 Select **File > Import Design**.

2 Leave the default `Design Editor file (.mvd)` in the **Import from** drop-down menu.

3 Browse to the design file `DIVCP.mvd` in the `mbctraining` directory, select it and click **Open**, then click **OK** to import the design. A new design node appears.

4 Rename the imported design (click, and press **F2**) `DIVCP`. You can right-click to change the Current View to examine the design points in 2D and 3D.

5 Choose `DIVCP` as the preferred design for future reference by selecting **Edit > Select As Best**.

6 Close the Design Editor.

Data Source

The data was collected using a constrained space-filling design on speed, load, intake cam phase, and exhaust cam phase. The points specified in the design were measured using the GT-Power engine simulation tool from Gamma Technologies (see `http://www.gtisoft.com/broch_gtpower.html`).

Simulink and StateFlow® simulation tools controlled the GT-Power model, running on a cluster of 14 desktop PC machines. The simulation time for the testing was 20 hours. The GT-Power model used predictive combustion, which gives good realistic results but is computationally expensive. Ten cycles were run at each spark advance setting after attaining steady-state speed/load conditions.

Importing and Filtering Data

In this section...
“Import and View Data” on page 2-18
“Filter Data” on page 2-24

Import and View Data

We provide an example data file, resulting from the example experimental design.

1. In the Model Browser, click the top root node, `Untitled`, in the **All Models** tree to go to the project view.
2. Click New Data in the toolbar or select **Data > New Data**.

   The Data Editor appears.
3. Click the Open File icon in the toolbar to load data from a file.

   The Data Import Wizard appears to select a file.
4. Use the Browse button to find and select the `DIVCP_Main_DoE_Data.xls` data file in the `mbctraining` folder. Double-click to load the file, and click **Next**.
5. The Data Import Wizard displays a summary screen showing the total number of records and variables imported, and you can view each variable's range, mean, standard deviation, and units in the list box. Click **Finish** to accept the data.

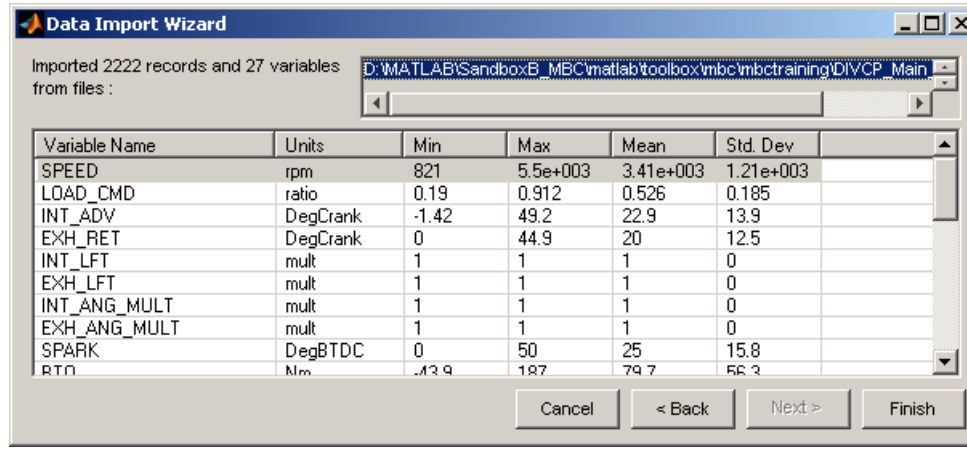

6 First you must define test groupings. Select **Tools > Change Test Groupings** (or use the toolbar button).

7 In the Test Groupings dialog box, clear the check box **One test/record**.

8 Locate and double-click `GDOECT` in the **Variables** list box. `GDOECT` appears in the left list.

9 Edit the **Tolerance** for `GDOECT` to `0.5`. You should see 202 tests defined, as shown.

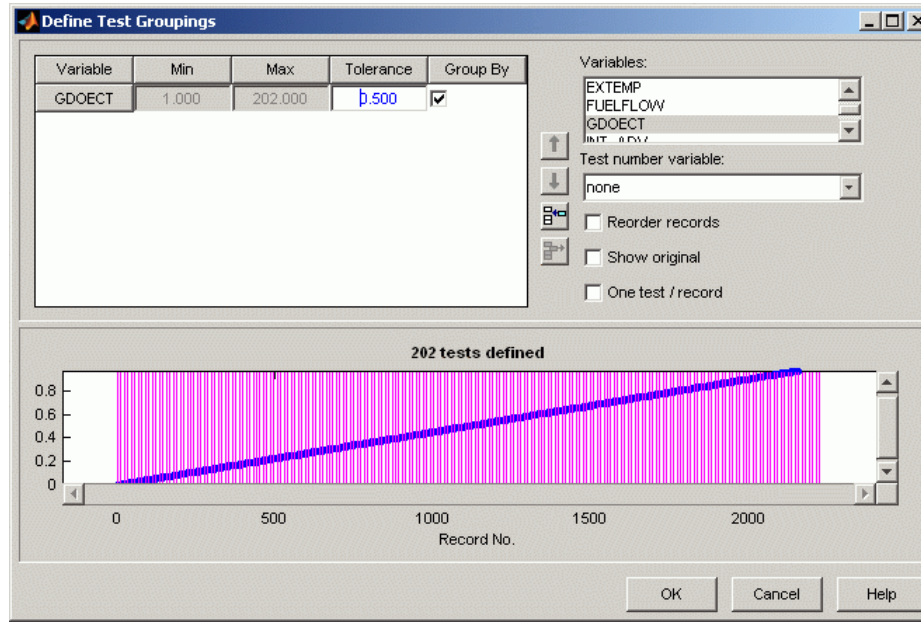

10 Click **OK** to close the Define Test Groupings dialog box.

11 Examine the data. The Data Editor is a powerful tool for displaying and sorting your data. You can use the right-click context menu to split the views, or use the toolbar buttons or the **View** menu. You can choose 2-D plots, 3-D plots, multiple data plots, data tables, and list views of filters, variables, test filters, and test notes.

For example, if you do not already have a 2-D plot, right-click the title bar of any plot and select **Current View > 2-D Data Plot**.

12 Click in the left lists to plot torque (`BTQ`) against spark (`SPARK`), then select one or more tests to display.

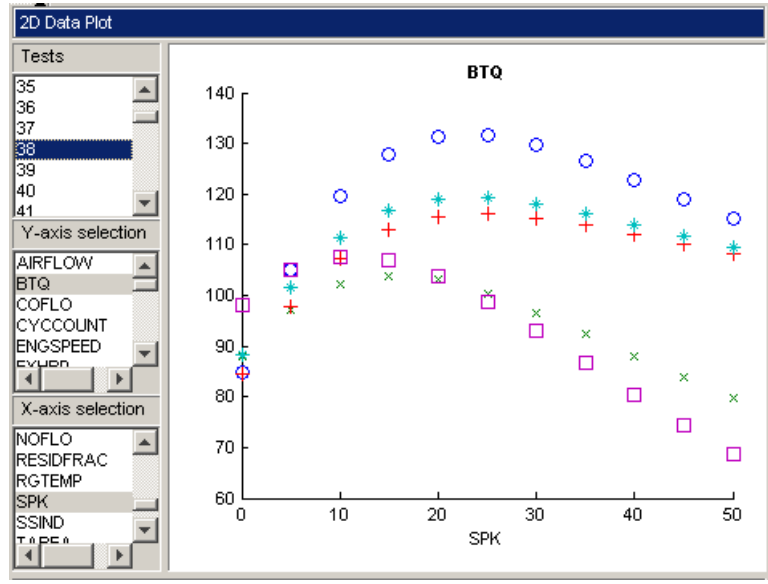

13 Right-click and select **Split View > 3-D Data Plot** to split the currently selected view and add a 3-D plot. Select one or more tests to display in the list at the left of the Data Editor, then choose three variables for the axes.

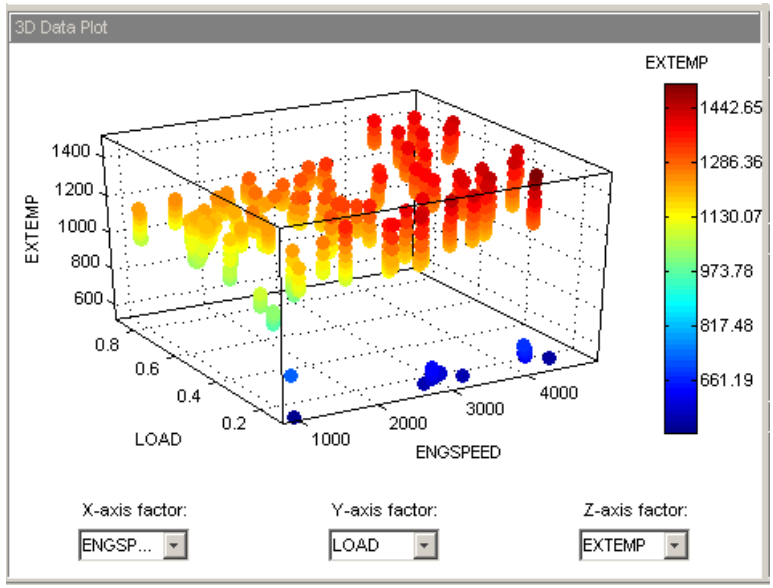

14 Right-click and select **Split View > Notes View** to split the currently selected view and add a test notes list view. This is empty until you add any test notes.

15 Select **Tools > Test Notes > Add**.

The Test Note editor appears.

a Enter an expression that defines the tests you want to note; for example `mean(BTQ) < 0` will evaluate the mean torque for each test and note those tests where the value is less than zero.

b Enter the text for this note in the edit box, e.g., `Negative Torque`.

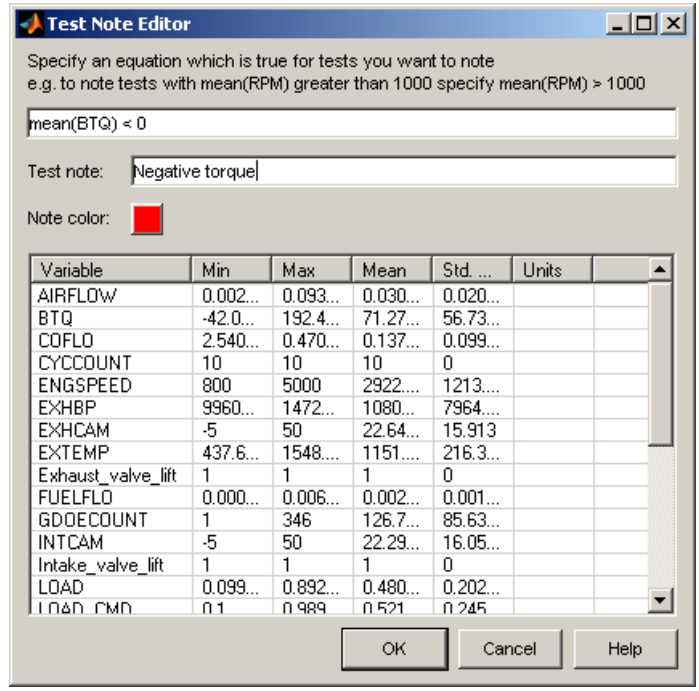

c Click **OK** to apply the test note.

16 Note that the new test note appears in the Notes list view.

Notes

Tests	
1	Negative torque
12	Negative torque
15	Negative torque
16	Negative torque
30	Negative torque
78	Negative torque
83	Negative torque
94	Negative torque
98	Negative torque
108	Negative torque
110	Negative torque
149	Negative torque
184	Negative torque
192	Negative torque
2	

You can sort the column for the `Negative Torque` note by clicking the column header (once to sort ascending and once more to sort descending). This allows you to quickly identify which tests satisfy the note definition. Investigate these test points in the other views. Select a test in the notes view and that test is displayed in the table view, 3-D plot, and multiple data plots views (but not the 2-D plots, which have their own test selection controls). If you select multiple tests, they are all shown in the data plots, but only the first test in the list is highlighted in the table view.

Filter Data

1. If you decide certain operating points are unsuitable for modeling, for instance, unstable points on the edge of the engine's operating envelope, you can use the Data Editor tools to help you identify and remove them. You can remove suspect data in the following ways:

   a. You can remove individual points by selecting them as outliers in 2D and multiple data plots by clicking, then selecting **Tools > Filters > Remove Outliers**. You can always replace them again with the **Tools** menu.

   b. You can define filters to remove all data points that do not fulfil a certain expression. Select **Tools > Filters > Add**. The Filter Editor appears.

c You can enter an expression that defines the records you want to keep. You can use any MATLAB function for filtering. Enter the following expression, which keeps records with `AFR` value greater than 14.25:

```
AFR>14.25
```

Click **OK**. Observe red records that have been filtered out in the top information bars. If you have turned on the option to **Allow Editing** in the data table view you can also see removed records in red.

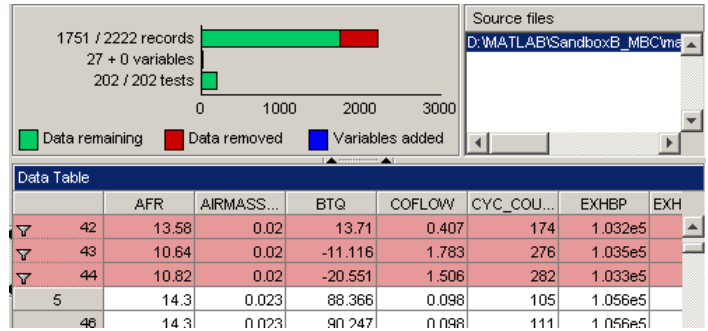

d Define another filter. Select **Tools > Filters > Add**. The Filter Editor appears.

e Enter the following expression, which keeps records with `RESIDFRAC` value less than 35:

```
RESIDFRAC<35
```

f You can exlude whole tests by defining test filters to remove them. Select **Tools > Test Filter > Add**. The Test Filter Editor appears.

You want to keep only those tests with sufficient points to fit the model (at least 5 points).

Enter `length(BTQ)>4` and click **OK**.

2 You can add a Filter Definitions and Test Filter Definitions list view using the right-click menu to see whether the filters have been successfully applied and how many records or tests are removed by each filter.

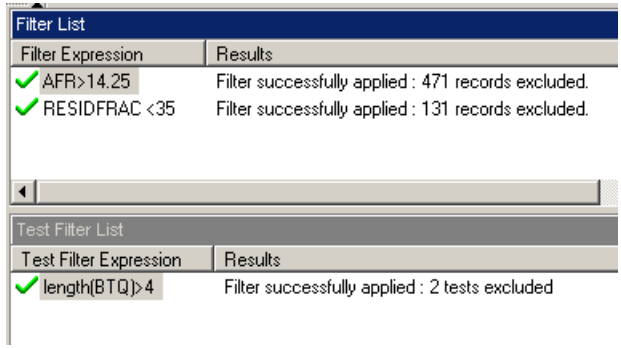

The bars at the top of the Data Editor always display the total numbers and proportion of removed data.

3 Close the Data Editor and return to the Model Browser.

How Is a Two-Stage Model Constructed?

Local models find the best fit of a curve to the data in each test. Each test in this case is a sweep of torque against spark angle, with speed, load, and cams held at a constant value for each sweep. The following illustrates a single sweep with a local model fitted.

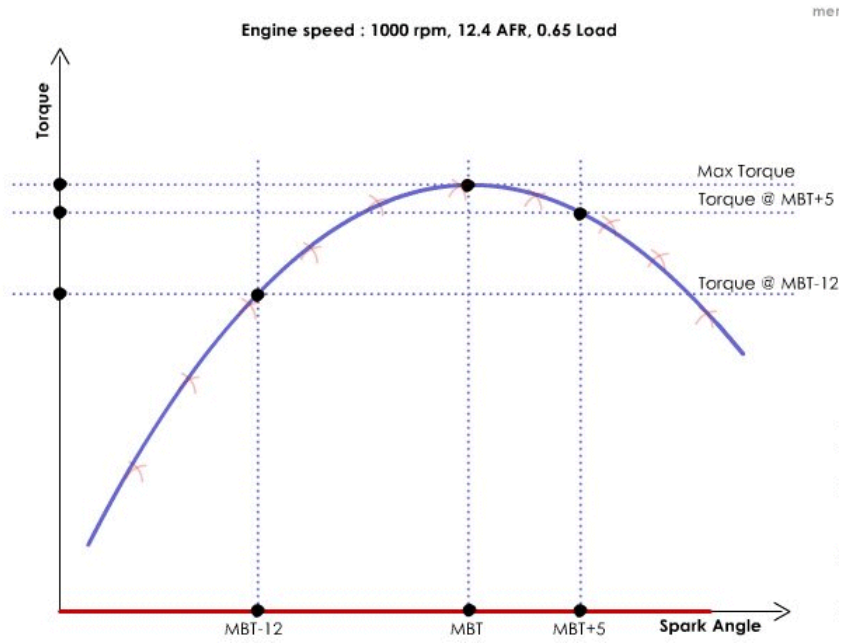

The local models provide the coefficients to generate global models. The equations describing those local model curves have certain coefficients such as max and knot, which for this data are peak torque and MBT spark (the spark angle that generates maximum brake torque).

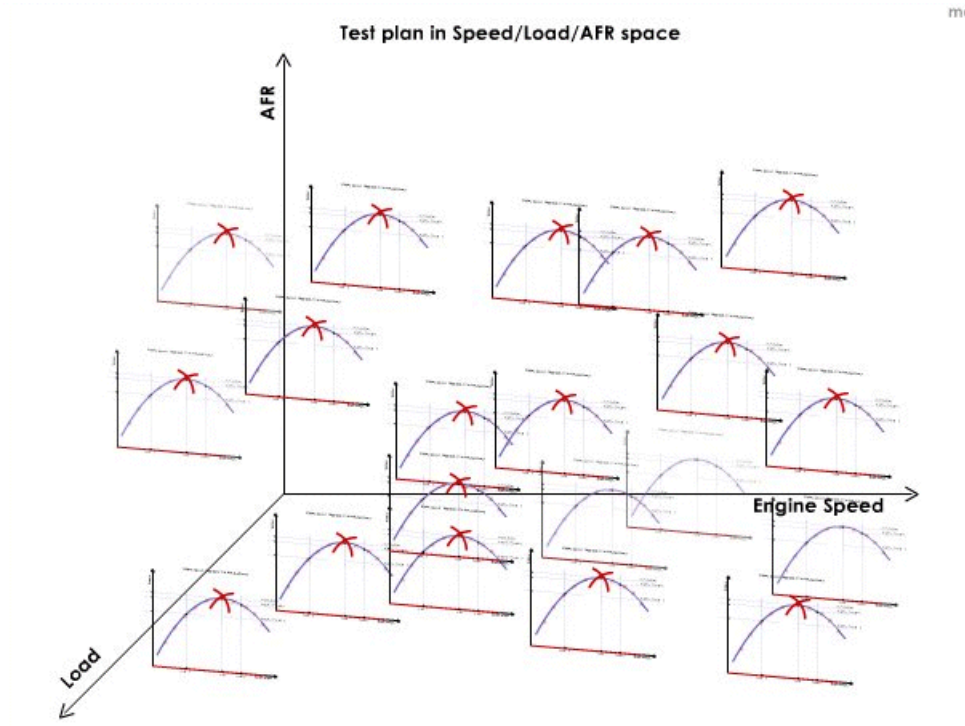

Local models are fitted to each test, in different places across the global space, as illustrated above. Each local model has coefficients for MBT and peak torque, etc. These coefficients become the data to which the global models are fitted. Coefficients such as peak torque and MBT are used to make the second stage of modeling more intuitive; an engineer will have a much better understanding of how a feature such as MBT spark varies through the global factor space than some esoteric curve fit parameter. Familiar variables like these are helpful to engineers trying to decide how well a model describes engine behavior. Better intuitive understanding allows much greater confidence in your models.

Global models are the best fit of a curve to the values of, for example, MBT for each test. This is repeated for each coefficient, producing several global models fitted to different coefficients of the local models. These coefficients are

referred to as response features of the local models. The following example shows a global model for maximum torque across the speed/load global space.

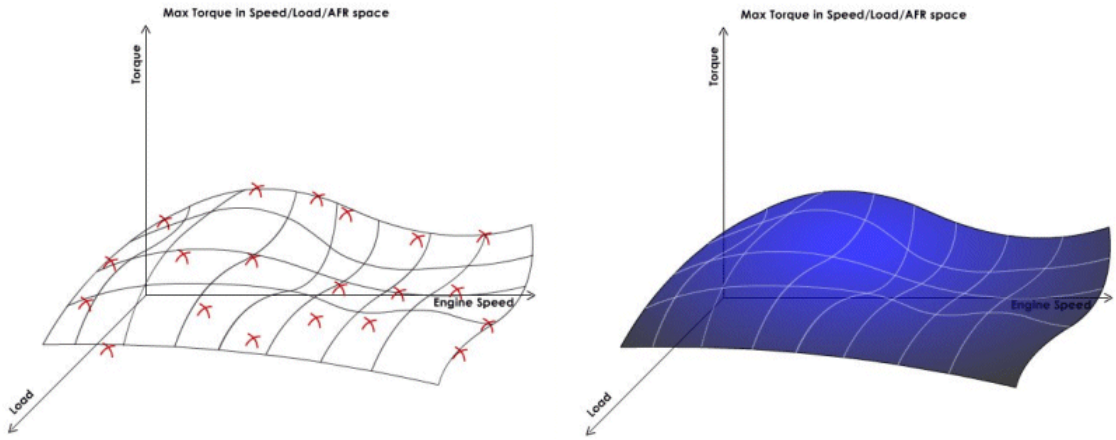

The two-stage model is a surface fitted across all the global models, to describe the behavior across all global variables.

It can be useful to think of local and global models as a series of 2-D slices, while the two-stage model fits a 3-D surface across the curves of the global model slices. It is difficult to visualize more dimensions! The following example shows a variety of 3-D plots of global models for properties of the local torque/spark curves (such as MBT, peak torque, and torque a number of degrees before and after MBT), showing how these properties vary across the speed/load global space. The 2-D plot of the global MBT model (on the right) demonstrates how MBT varies with engine speed.

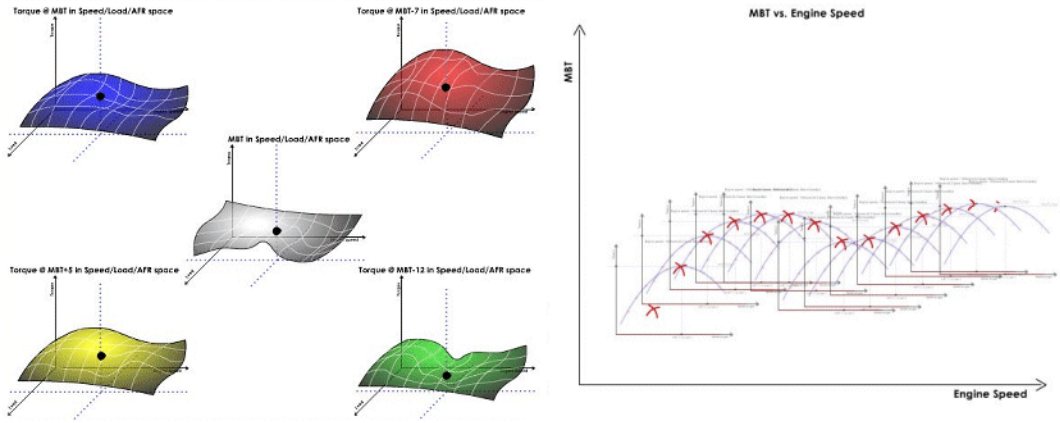

The two-stage model can take values of each coefficient at a certain value of, say, speed, to generate a new curve of torque against spark. This is a slice through the two-stage model surface.

In other words, you can test your two-stage model by comparing it with the local fit and with the data. For example, you can reconstruct a local torque/spark curve at an operating point by taking the values of MBT and peak torque and the curvature from the two-stage model, and then validate this reconstructed curve against the original fit and the data. The two-stage model can also predict responses between tests, for new sweeps at intermediate values for which there is no data. If the two-stage model shows an accurate fit when compared to the local sweeps, this is a good sign that the engine behavior is well described by the model across the global variables.

For more details on two-stage modeling, see “Two-Stage Models” on page 5-2, and see “Two-Stage Models for Engines” in the Model Browser documentation for more statistical depth. (In the Help Browser you can right-click and select **Back** to return to previous pages).

Building the Models

In this section...
"Completing Model Setup" on page 2-31
"Specifying the Local Model Type" on page 2-32
"Selecting Data and Responses to Model" on page 2-33

Completing Model Setup

You have already set up the local and global inputs and the global model type before constructing the design. Now you have imported, examined, and filtered the data in preparation for modeling. To complete your model setup, you need to specify the local model type, select the data to model, and choose the responses (model outputs) you want to model.

Remember that the aim of this case study is to produce optimized tables for

- Intake cam phase
- Exhaust cam phase
- Spark timing schedules

These tables are all functions of load and rpm, subject to constraints of operating region and residual fraction.

To produce these tables, you need to make accurate models of the behavior of torque, exhaust temperature, and residual fraction at different values of speed, load, spark, and cam timings. You have set the local model input as spark, and the global model inputs as engine speed, load, intake cam phase, and exhaust cam phase. Therefore, the responses you want to model are

- Torque
- Exhaust Temperature
- Residual fraction

Specifying the Local Model Type

The first response you want to model is torque against spark. The shape of torque/spark curves is well understood and you have examined some in the Data Editor. Polynomial spline curves are useful for fitting these shapes, where different curvature is required above and below the maximum. Therefore, you should set the local model type to polynomial spline. A spline is a curve made up of pieces of polynomial, joined smoothly together. The points of the joins are called knots. In this case, there is only one knot, at the maximum. The location of the knot in this case marks MBT.

To specify polyspline as the local model type,

1. First select the `Two-Stage` test plan node in the model tree, so you can see the test plan diagram.
2. Double-click the local model icon in the test plan diagram.

   The Local Model Setup dialog box appears.

   a. Select `Polynomial Spline` from the **Local Model Class**.

   b. Set **Spline Order** to `2` below and `2` above knot.

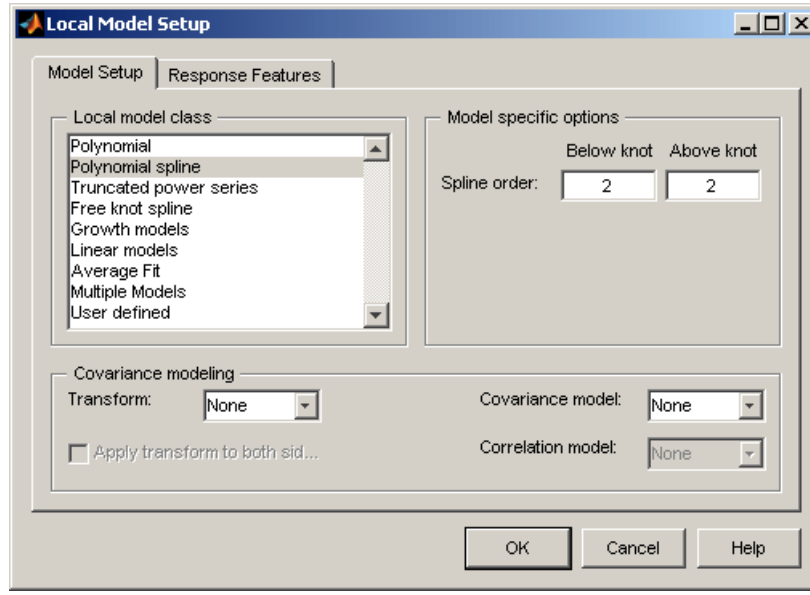

3 Click **OK** to dismiss the dialog box.

Notice that the new name of the local model class, `PS` (for polyspline) `2,2` (for spline order above and below the knot) now appears on the two-stage model diagram.

Selecting Data and Responses to Model

You have set up model types and model inputs. Now you can select the data for modeling and the responses (model outputs) you want to model.

1 Double-click the Response block in the test plan diagram. The Data Selection Wizard appears.

2 Because the test plan contains a design, the radio button **Match selected data to design** is selected. The imported `DIVCP` design appears in the left list (and other designs you created), and the data object in the right list.

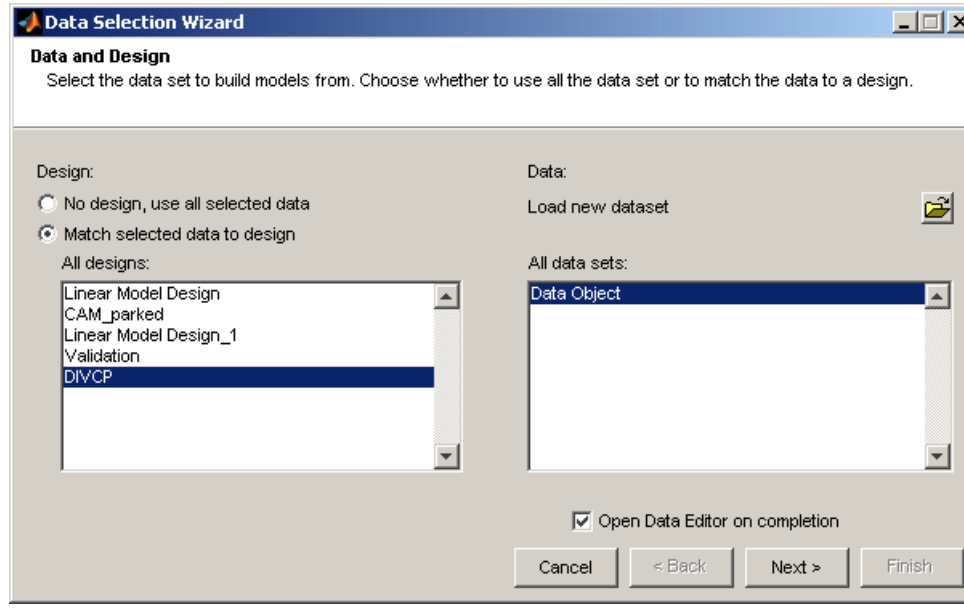

Select the `DIVCP` design and click **Next**.

3 The wizard tries to match each symbol to a data signal (shown as **Signal Name**). Make sure `SPARK`, `SPEED`, `LOAD`, `INT_ADV` and `EXH_RET` are selected, as shown. Edit incorrect matches by selecting the corresponding entries in the input and data lists and click the green arrow to select. Do not select the copy range check box.

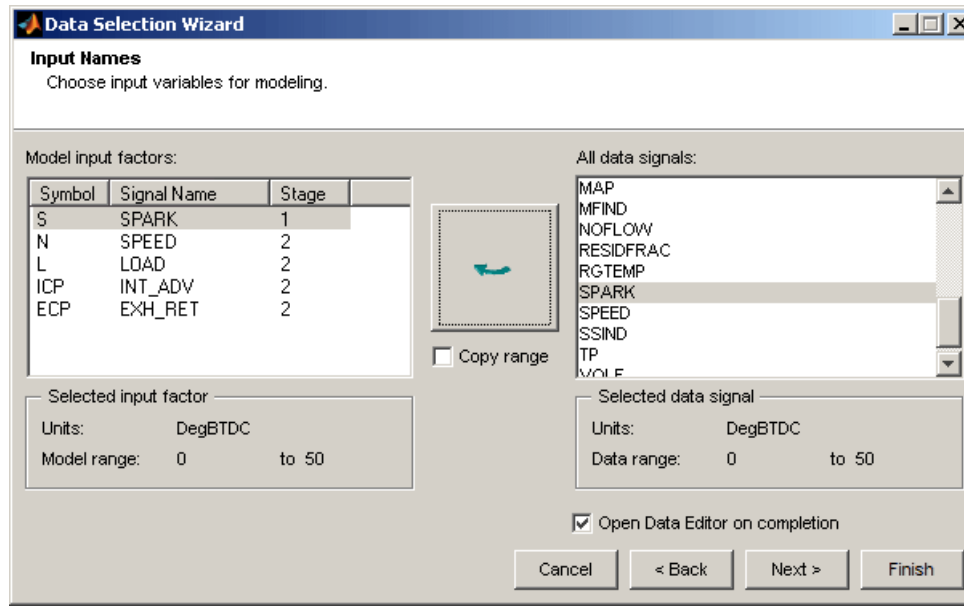

Click **Next**.

4 On this screen you select the response to model.

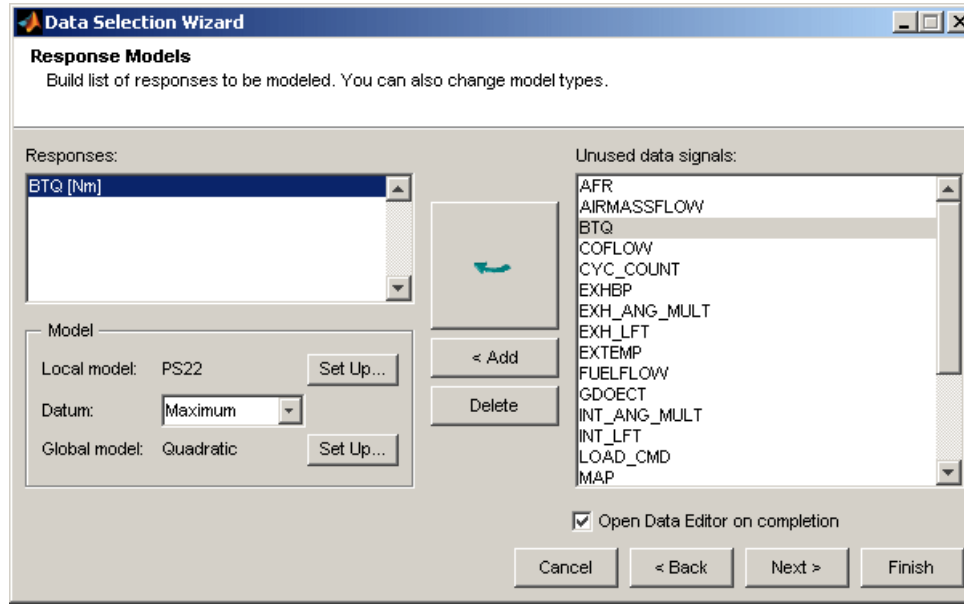

a Select `BTQ` as the response you want to model and click **Add**. Notice the local model settings you set earlier (PS22).

b Select `Maximum` from the **Datum** drop-down menu. In this case, the maximum of the torque/spark curves is MBT (spark angle at maximum brake torque), so this can be a useful feature to model.

c Click **Next**.

5 On this screen, you see the settings for matching data to designs.

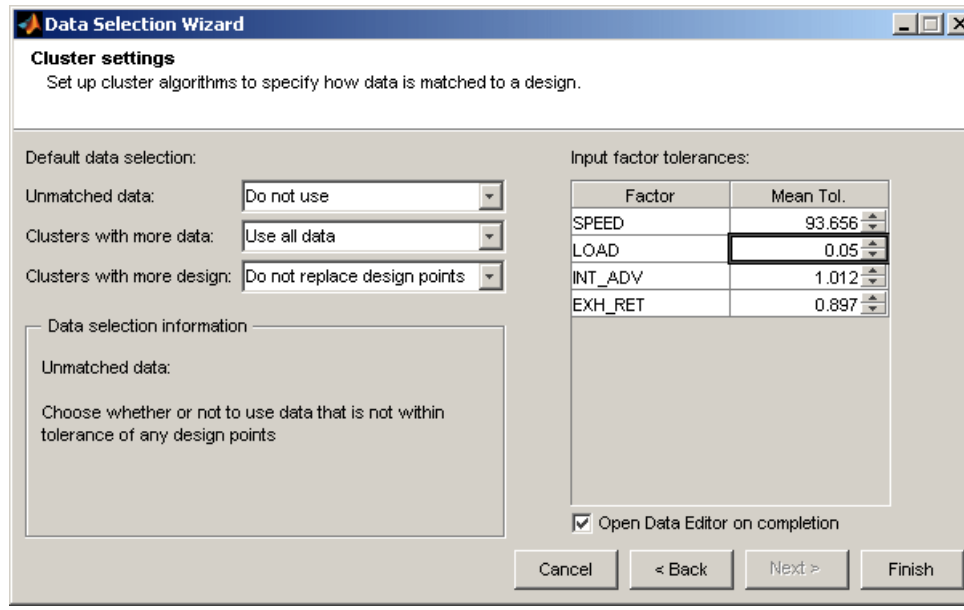

a Enter `0.05` in the LOAD tolerance edit box.

b Select `Do Not Use` from the drop-down menu for **Unmatched Data**. Data points that do not fall within tolerance of the design points will not be used for modeling.

c Leave the other settings at the defaults and click **Finish**.

The Data Editor appears so you can select data for modeling. The response model is built when you close the Data Editor.

6 Right-click a view and select **Current View > Cluster View**. Use the drop-down menus to select `SPEED` and `LOAD` for plotting.

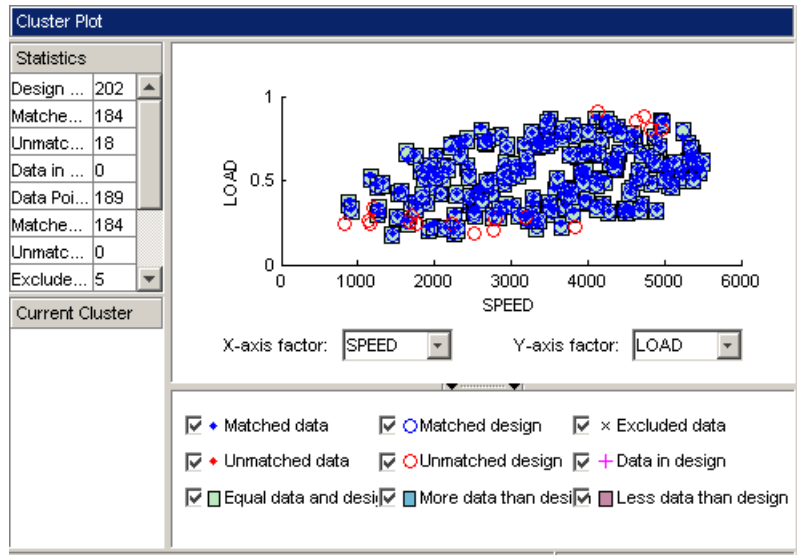

7 Inspect the data and design points by using the check boxes in the cluster plot. Clear the check box for `Equal data and design` to remove matched (green) clusters from the display. These data points fall within tolerance of the design point, so these points are selected for modeling. When you remove them from the display, you can see other points more clearly.

8 Clear the box for **Unmatched design** points, and look at the remaining excluded data points (plotted as crosses).

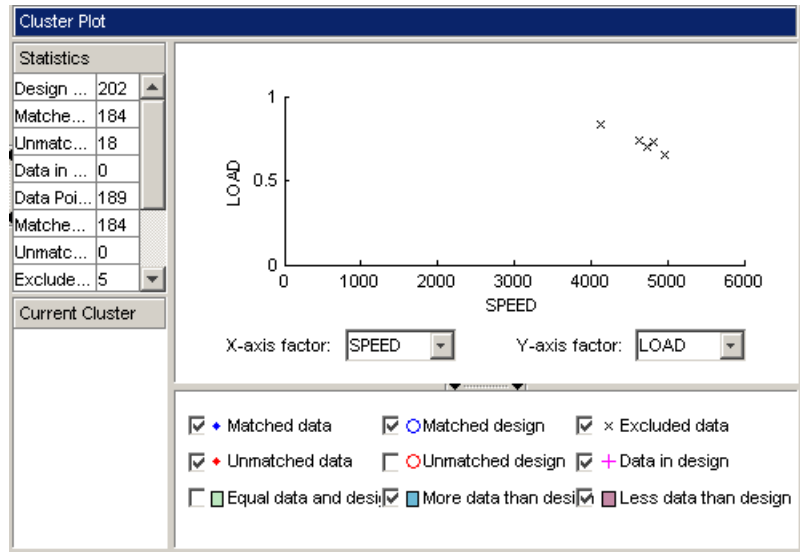

These points have not been matched to any design points because they do not lie within tolerance. The value of load achieved at these points was not close enough to the desired value of load, indicating a problem with these operating points. Notice that these points lie near the edge of the constrained area. In the Data Wizard, you selected `Do not use` unmatched points, so these points have not been selected for modeling. You can always change the tolerances and decide to include unmatched data points later; the choice in the Data Wizard is not irrevocable.

9 To accept the matched data for modeling, close the Data Editor. A dialog appears to check that you want to build the response model for torque and update the `Actual Design` to include all data selected for modeling. Click **Yes**, and the models are created.

Selecting Local Models

You should inspect the local and global models in turn, removing outliers if appropriate and trying different model types, before creating a two stage model. If the fit is good at the local and global levels you have the best chance of creating a two-stage model that accurately predicts the engine behavior. First, inspect the local models.

1 When you create the models the view switches automatically to the new local model node (PS22) in the model tree, in the **All Models** pane, and the local model view appears on the right. Make sure the local model node (PS22) is selected before following these steps.

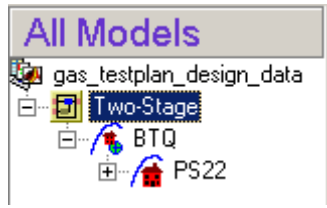

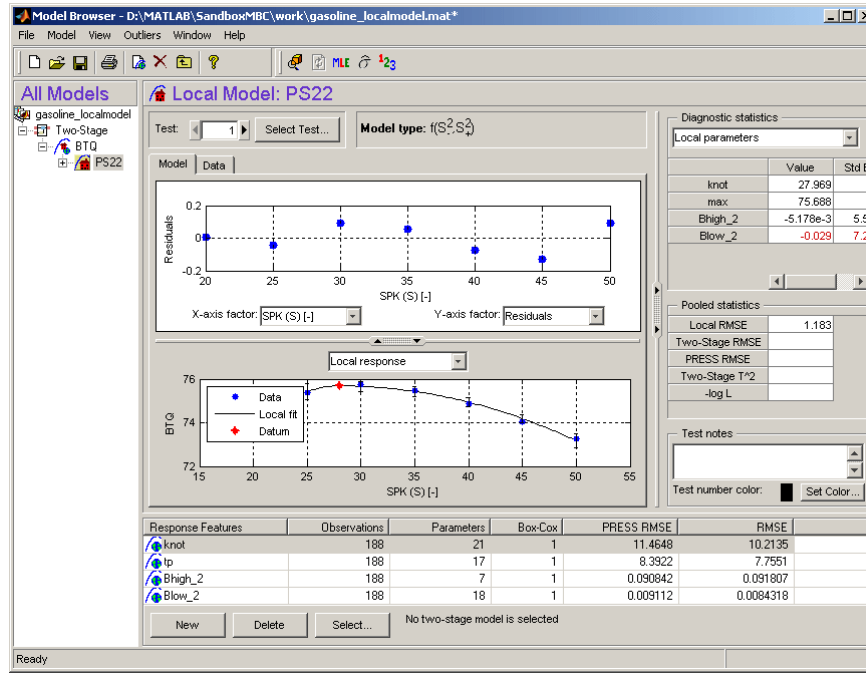

2 Look through the tests to inspect the fits. Use the **Test** controls.

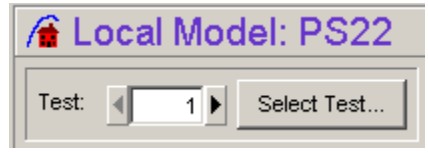

3 To quickly identify problem tests, click RMSE Plots in the toolbar (or **View** menu).

Inspect tests with high error values in the local model view. You can navigate to a test of interest from the RMSE Explorer by double-clicking a point in the plot to select the test in the Model Browser local model view. Plot all response features in the RMSE Explorer and investigate tests with extreme values.

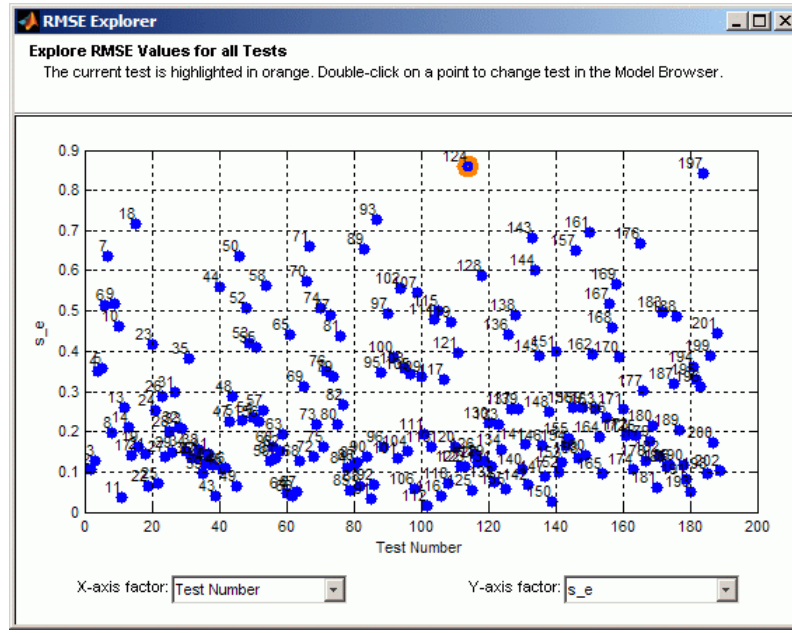

4. Consider removing outliers to improve fits if some points are badly distorting the torque spark curve (use right-click or **Outliers** menu). Pay particular attention to end points. For example, for tests where the majority of points are at higher spark angles than the maximum (at MBT), it can improve the fit to remove some of these long "tails". It can be useful to remove outliers in this region, because there is likely to be knock at spark values much higher than MBT where the engine is less stable. Similarly, as there is no knock in simulation data, points can be collected far in advance of MBT, and it can improve the fit to remove these.

5. If removing some outliers does not bring MBT within the range of the data points, consider removing the whole test (use the **Outliers** menu).

6. After making changes, check the RMSE Explorer plots again for problem tests.

7 Check all tests before inspecting the global models. Try looking at the `Local diagnostics` (select from the list in the **Diagnostic Statistics** pane). Check for high values of **Cond(J)** (e.g., > 10^8). High values of this condition indicator can be a sign of numerical instability.

In the next section, "Creating Boundary Models" on page 2-44, you will construct a boundary model before inspecting the global models.

Creating Boundary Models

You can create a boundary model at the test plan node. A model describing the limits of the operating envelope can be useful when you are evaluating optimization results and global models. It is useful to create the boundary model before viewing global models, so you can see the model areas inside the boundary on plots. You will build two boundary models and combine them, to map the envelope in speed and load, and also in all four inputs.

1. Select the test plan node in the model tree.
2. Select **TestPlan > Boundary Models**.

   The Boundary Editor opens.
3. Click New Boundary Model in the toolbar. A dialog opens where you can choose to build a boundary model of the response, local, or global inputs. In this case, you are interested only in making a model of the global boundary. You are not interested in the `Local` boundary (you know the range of spark already). Select **Global** and click **OK**.

   A dialog opens where you can select constraint inputs.

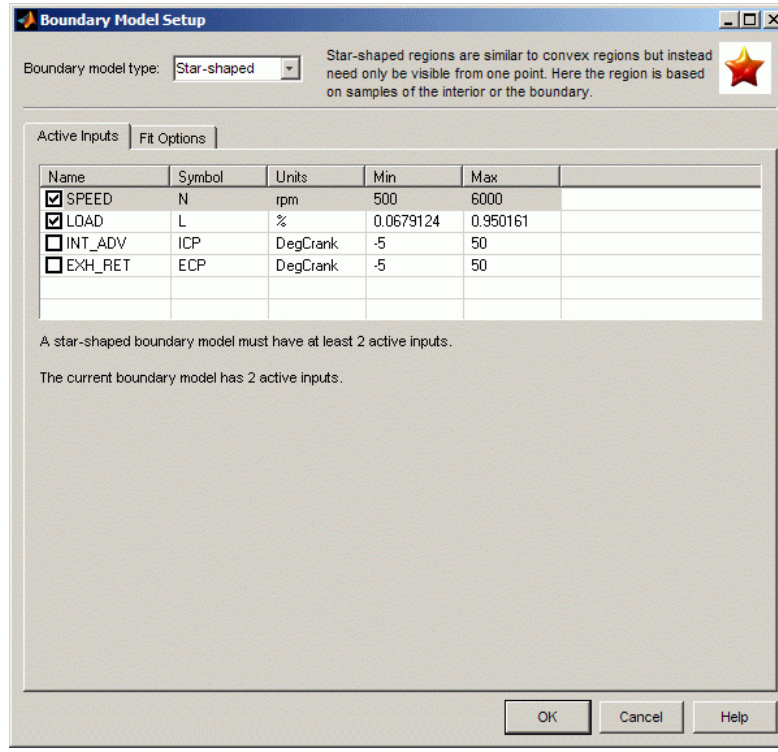

4 Leave the **Constraint type** set to `Star shaped` and leave only the `SPEED` and `LOAD` check boxes selected. Click **OK** and the boundary constraint is calculated.

A `Star shaped(N,L)` child node appears under the `Global` node and the view switches to the new model.

5 Select **View > Current View > 3D Slice** (or use the toolbar button) to examine the shape of the new boundary model. Drag the axes to rotate the plot. Use the drop-down menus to change the variables plotted. In any view, you can change the position of the plotted slice by altering the variable values in the edit boxes. You can split the views as in the Design and Data Editors using the buttons in the title bars, or the right-click or **View** menus.

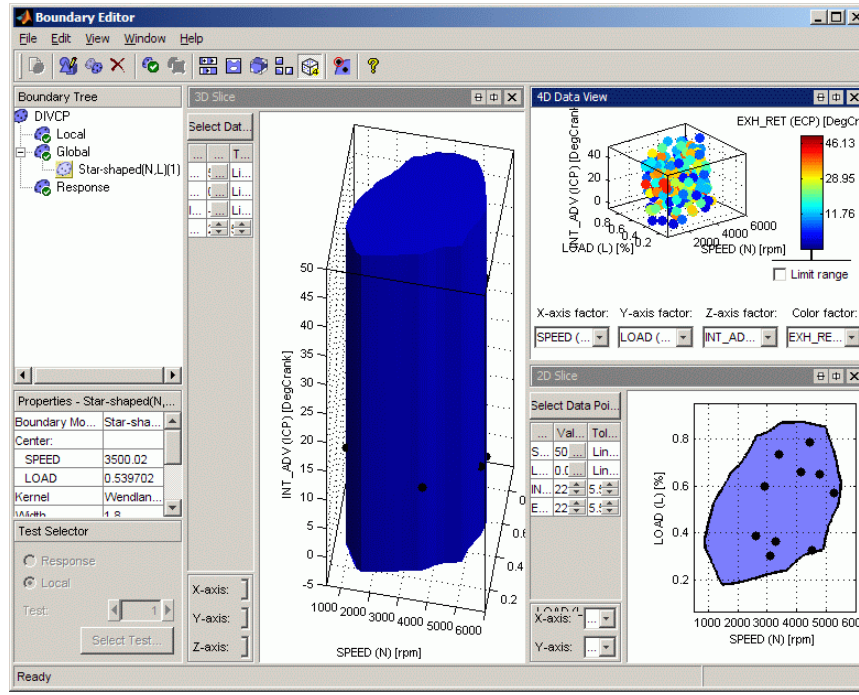

6 Click Show Boundary Points in the toolbar. Boundary points are outlined in red. In some projections and at some plot resolutions, points can falsely appear to be outside the boundary. You can alter the number of points plotted in the Value edit boxes, but high resolutions can be time-consuming to plot. You can also check particular points: in the 3D, 2D, and 1D Slice views you can click (and hold) points to see the values of the global variables at that point and the `Distance`. A distance of 0 means the point is exactly on the boundary, and negative values show the distance inside the boundary. You can use this function to check that enough points are inside or close to the boundary.

7 In a 2D or 3D Slice view, click the button **Select Data Point**. A dialog appears where you can select a data point. When you click **OK** the slice is plotted at the location of the selected data point. You can also double-click points to move the slice to that location.

8. Select **View > Current View > Pairwise Projections** (or use the toolbar button). Here you can see pairwise projections to view the boundary across all combinations of factors.

9. On one of the pairwise plots, click and drag to select a small area of points. The area you select is colored yellow across all plots, so you can view how those points are distributed across factors. With some detailed surfaces, areas in the pairwise plots can appear as discrete patches, so this feature is useful for tracking regions across factors. This can help you decide whether the boundary is capturing enough (or too much) detail of the surface.

10. Select the Global node, and click New Boundary Model in the toolbar. You can build other boundary models to compare and combine for the most useful model.

    From the Global node you will automatically get a global child node, so you do not need to select a level. The Boundary Model Setup dialog box appears.

11. This time leave all four inputs selected and click **OK**. A new global boundary constraint is calculated.

12. Compare the two global boundary models by selecting the Pairwise view, and selecting each model in the tree in turn.

13. Select one of the models, and select **Edit > Add to Best** (also in the toolbar).

14. Select the other global model and select **Edit > Add to Best**.

15. Now select the parent `Global` node to see the combination of the two constraints. Compare with both the child global nodes in the pairwise view.

    For more information see “Combining Best Boundary Models”.

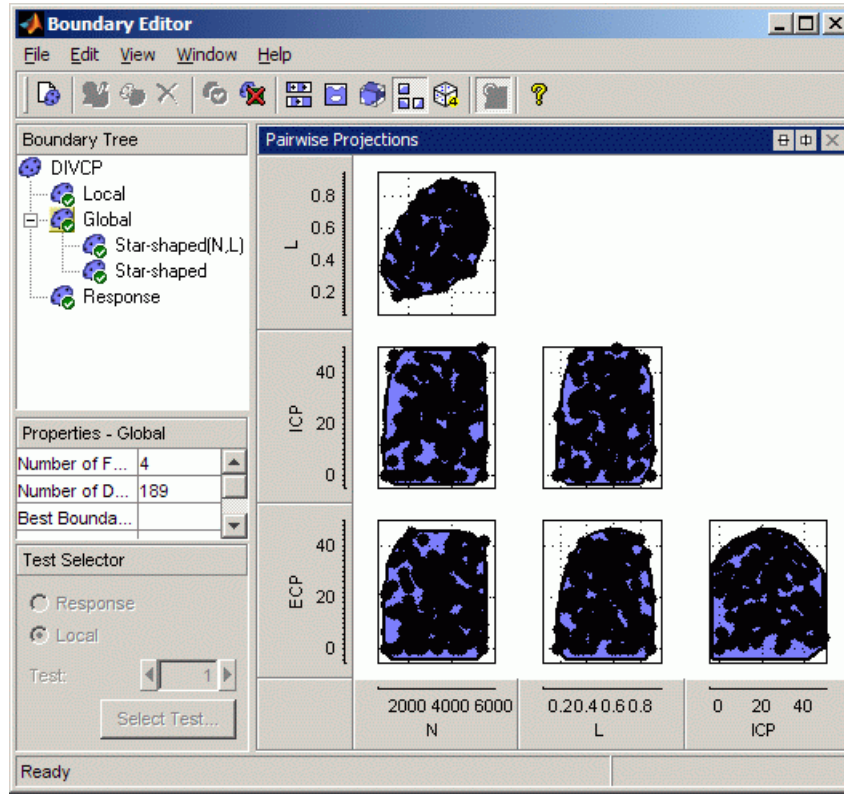

16 Close the Boundary Editor to return to the Model Browser. Once calculated, the boundary constraints remain part of the test plan unless you delete them.

Selecting Global and Two-Stage Models

In this section...
“Inspect the Global Models” on page 2-49
“Create Multiple Models to Compare” on page 2-51
“Create a Two-Stage Model” on page 2-55
“Adding New Response Models” on page 2-57

Inspect the Global Models

When you are satisfied with the local fits, inspect the global models in turn. You should check trends of global models. Do the trends go in the right direction? Previous engineering knowledge can be applied. The following steps suggest useful plots for investigating trends.

1 Expand the local model node (`PS22`) in the model tree and click `knot`.

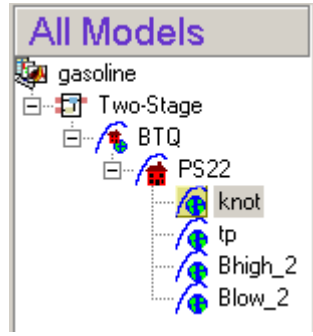

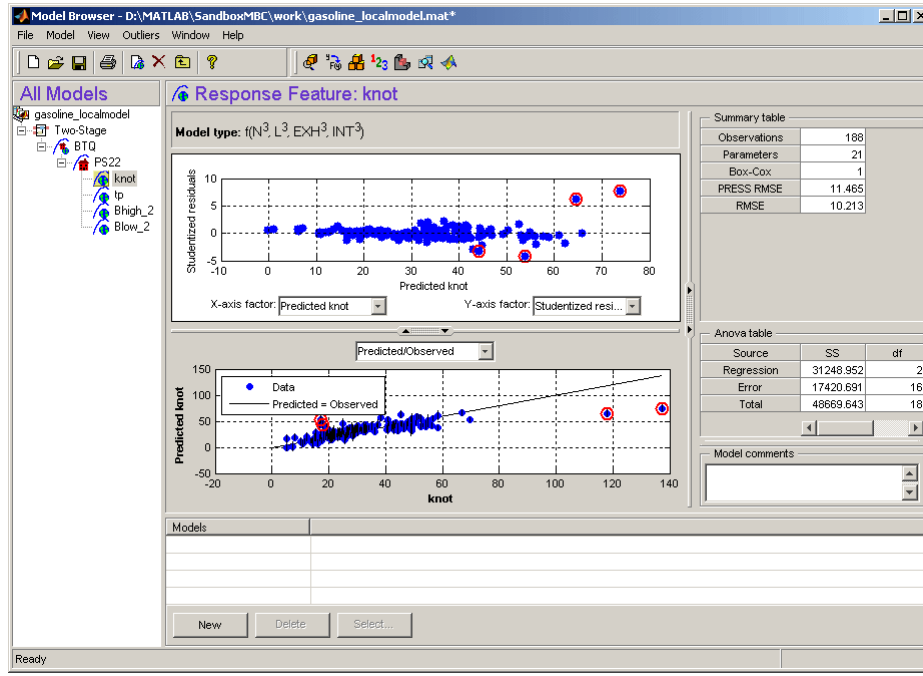

2 Right-click outliers (or any point) to see a plot of the test. You can inspect the shape of the torque/spark curve and see the values of the global variables. This can help you identify problem tests, perhaps on the edge of the stable operating region.

Also you can use the global model view scatter plots to plot predicted values against variables.

Note After identifying problem tests at the global level, return to the local level to decide whether to remove outliers or the whole test.

---

3 Select **Model > Evaluate > Fit Data** to open the Model Evaluation window.

- Select **View > Cross Section** to see the trends of your current model e.g., maximum torque (`max` model), or MBT (`knot` model). Try the **Response Surface** view. Apply engineering knowledge to check trends.

- Be aware that `Bhigh2` and `Blow2` must be negative (to ensure maximum torque occurs at MBT). To check for this, select **View > Response Surface** and click the `Table` **Display type** — here cells outside the boundary model are highlighted yellow to help you focus on the area of interest.

4. Also you can use the global model view scatter plots to plot `Predicted Bhigh2` or `Predicted Blow2` against `Speed` to look for suspicious positive values.

Create Multiple Models to Compare

1. Click the **Build Models** toolbar button.

2. Click **Browse**, locate the `mbctraining` directory and click **OK**.

3. Select the `DIVCP` template and click **OK**.

   Choose `PRESS RMSE` in the Model Selection dialog box and click **OK**, and a selection of child model types are built for `knot`. `PRESS RMSE` is the criterion used for automatically selecting the best model out of the child nodes.

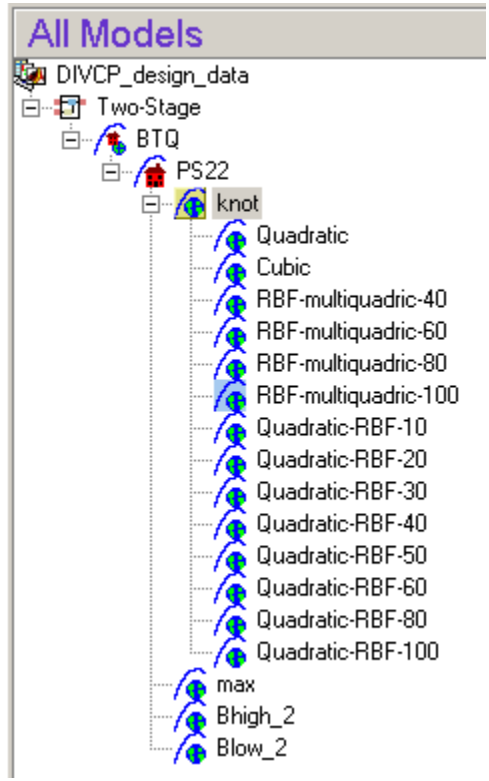

You can save your own template of any collection of child models. From the parent model node (e.g., `knot`) select **Model > Make Template**.

4. Now `knot` has several child models. Inspect each model in turn.If you remove outliers that have an RBF center (marked with a star) be sure to refit the model (click the toolbar button Update Model Fit) to reselect widths and centers.

5. Return to the `knot` model node and look at the statistics reported in the list of child models at the bottom. From any parent model node you can see a list of statistical comparisons for all the child nodes in the lower list pane, along with information such as the number of parameters. Use this information to help you decide which model is best.

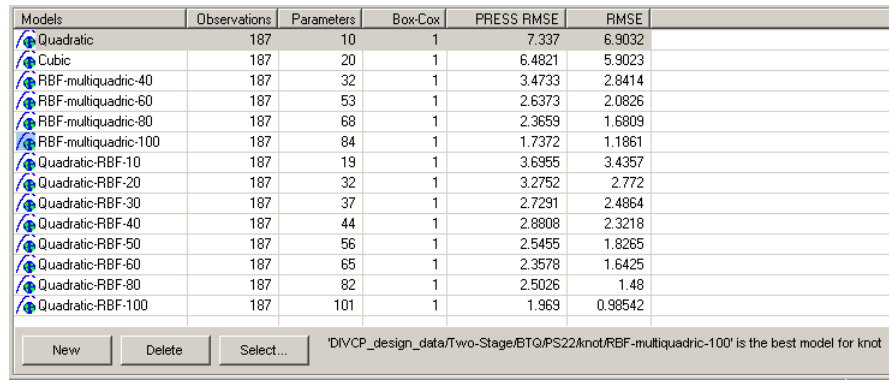

Models	Observations	Parameters	Box-Cox	PRESS RMSE	RMSE
Quadratic	187	10	1	7.337	6.9032
Cubic	187	20	1	6.4821	5.9023
RBF-multiquadric-40	187	32	1	3.4733	2.8414
RBF-multiquadric-60	187	53	1	2.6373	2.0826
RBF-multiquadric-80	187	68	1	2.3659	1.6809
RBF-multiquadric-100	187	84	1	1.7372	1.1861
Quadratic-RBF-10	187	19	1	3.6955	3.4357
Quadratic-RBF-20	187	32	1	3.2752	2.772
Quadratic-RBF-30	187	37	1	2.7291	2.4864
Quadratic-RBF-40	187	44	1	2.8808	2.3218
Quadratic-RBF-50	187	56	1	2.5455	1.8265
Quadratic-RBF-60	187	65	1	2.3578	1.6425
Quadratic-RBF-80	187	82	1	2.5026	1.48
Quadratic-RBF-100	187	101	1	1.969	0.98542

New | Delete | Select... | 'DIVCP_design_data/Two-Stage/BTQ/PS22/knot/RBF-multiquadric-100' is the best model for knot

- Look for lower RMSE values to indicate better fits.
- Look for lower PRESS RMSE values to indicate better fits without overfitting.

  PRESS RMSE is a measure of the predictive power of your models. It is useful to compare PRESS RMSE with RMSE as this may indicate problems with overfitting. RMSE is minimized when the model gets close to each data point; 'chasing' the data will therefore improve RMSE. However chasing the data can sometimes lead to strong oscillations in the model between the data points; this behavior can give good values of RMSE but is not representative of the data and will not give reliable prediction values where you do not already have data. The PRESS RMSE statistic guards against this by testing how well the current model would predict each of the points in the data set (in turn) if they were not included in the regression. To get a small PRESS RMSE usually indicates that the model is not overly sensitive to any single data point.

  If the value of PRESS RMSE is much bigger than the RMSE then you are overfitting - the model is unnecessarily complex.

  PRESS RMSE can be the most helpful single statistic you can use to search for the best fit relative to the number of terms in the model. However you should not rely on any single statistic, but use a variety of criteria and especially the graphical tools available for comparison of models in the Model Evaluation tool when you click Select. You can also use other diagnostic statistics to help you select models. For detailed guidance on how to understand the selection tools see and the "Model

Selection Guide" in the Model Browser documentation. (In the Help Browser you can right-click and select **Back** to return to previous pages).

- You can add other statistics here. Select **Model > Summary Statistics**. In the dialog, select the check box for **AICc** and click **OK**. A column of **AICc** values appears in the model list. This can be a useful statistic for comparison of models as it penalizes over-parameterized models. Over-fitting models can later cause problems with optimization when you use the models to create calibrations. See "Using Information Criteria to Compare Models".

6 From the `knot` model node, click **Select** in the **Models** pane at the bottom to open the Model Selection window. Try the different views and compare the child models by selecting them in the **Model List** at the bottom.

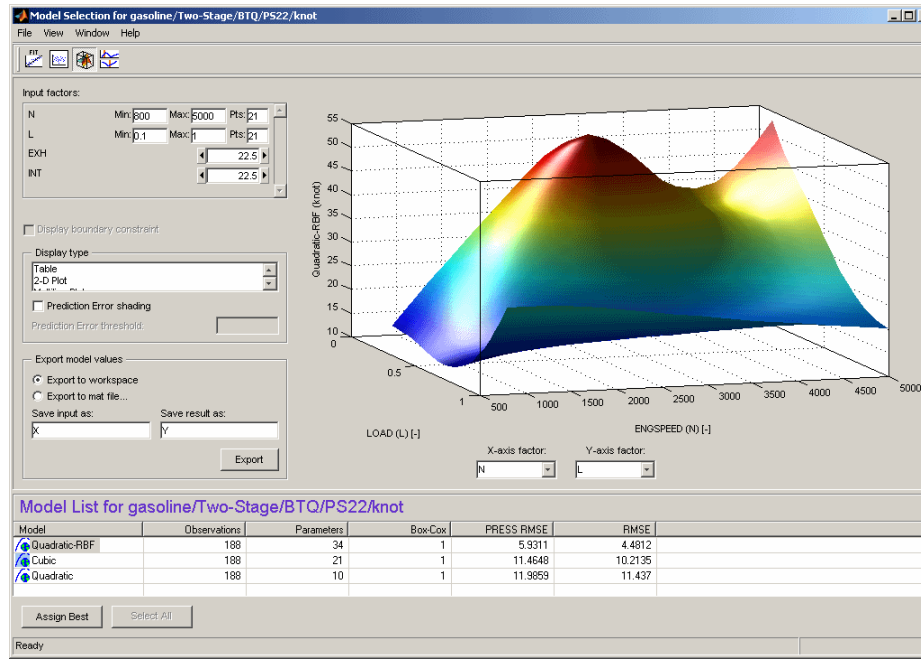

7 Select one of the child models as best (click the button **Assign Best**) and close the Model Selection window to return to the Model Browser. Try other model types if you are not satisfied with the quality of the fit. You could work through the modeling tutorial for more guided examples of how to select models, see Chapter 5, "Tutorial: Model Quickstart ".

8 When you have selected a best child model, select **File > Clean-up Tree** to discard all but the model you chose as best. The process of searching for good model fits often results in a large selection of models, and cleaning up the tree reduces file size. It can also save time when saving the file, especially if you have made a change such as removing an outlier at local level, that causes all response feature models to be refitted.

9 Select another global model such as `Bhigh_2` and click Build Models in the toolbar. Select the `DIVCP` template to automatically build the same selection of child model types for `Bhigh_2`. This can help you quickly build multiple models to compare. The Model Selection dialog appears, where you can choose a criterion such as `PRESS RMSE` or `AICc` for automatically selecting the best model out of the child nodes.

10 Repeat this process of searching for good global model fits for the other three global models.

Create a Two-Stage Model

1 When you have selected best models for each global model, return to the local model node (`PS22`) in the model tree, and click **Select** in the **Response Features** model list pane at the bottom to calculate the two-stage model and open the Model Selection window.

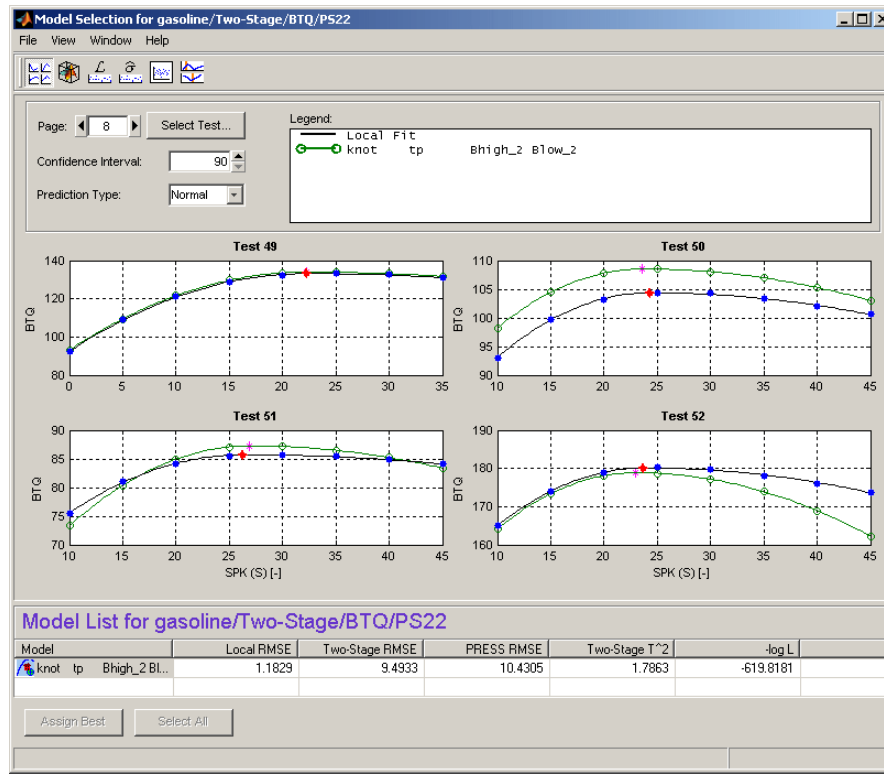

Look through the plots of the new two-stage model against the local fits and the data.

When you close the Model Selection window and accept the new model as best, the two-stage model is copied to the `BTQ` response node in the model tree.

2 You can choose to calculate maximum likelihood estimation (MLE) at this point. This process refits, taking proper account of the correlation between different response features. Try calculating MLE, then select the BTQ node and return to the Model Selection window to compare the MLE model with the univariate model (click **Select All** in the Model List pane).

For more guidance on creating multiple local, global and two-stage models to search for the best fit, work through the step-by-step examples in Chapter

5, “Tutorial: Model Quickstart ”, especially the section “Creating Multiple Models to Compare” on page 5-43.

Adding New Response Models

To complete the model building you should follow the same process as for modeling torque to create and select good models of the other responses: exhaust temperature and residual fraction (instructions below). These models are required for the optimization problems. Remember you can look at the example finished project, `Gasoline_project.mat`, in the `mbctraining` directory, to see how the example models have been constructed. You will use these example models in CAGE for the optimization section.

Exhaust Temperature Model

1. Select the test plan node in the model tree to return to the test plan view.
2. Double-click the **Responses** icon in the block diagram.
3. The Response Model Setup dialog appears.
   a. Select `EXTEMP` from the list of signals.
   b. Click the Local Model **Set Up** button. Select `Polynomial` from the **Model class** drop-down menu and click **OK**.
   c. Select `BTQ datum` from the **Datum** drop-down menu. It can be useful to plot the position of MBT on other models.

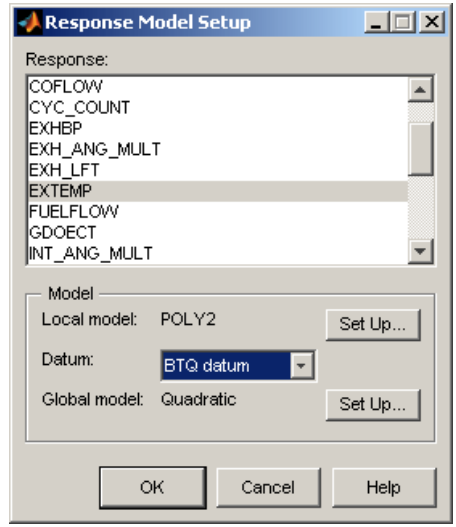

d Click **OK**.

A new set of local and global models is calculated for the exhaust temperature response.

The view switches to the new local node (`POLY2`).

1. You can copy the outliers you selected for the torque models. Select **Outliers > Copy Outliers From**. The Copy Outliers dialog box appears. Select the `BTQ` local model node in the tree (`PS22`, under `BTQ`), and click **OK** to copy the outlier selections to the EXTEMP local node.

2. Examine the fit in the same way as you did the torque fits.

   - Try different local models. Click **New** at the `EXTEMP` response node.
   - Try different global models as you did for the torque response features. Use **Build Models** and the `DIVCP` model template, or click **New** to add child nodes and try different model types one at a time. If you use **Build Models** at the local level you can apply a template to all response feature models at once.
   - Try a new response feature. Click **New** at one of the local nodes you have created and enter -10 for the value to use MBT minus 10 degrees of spark as a new response feature.

See Chapter 5, "Tutorial: Model Quickstart " for simple worked examples of adding new local, global, response feature and two-stage models.

3. When you are satisfied with the fits, return to the local model node and click **Select** to calculate the two-stage model. If you have added new response features, there will be more than one two-stage model to choose from in the Model Selection window.

4. Try calculating MLE and return to the Model Selection window to compare the MLE model with the univariate model and select the best.

Residual Fraction Model

1. Repeat the steps to create a model of residual fraction (`RESIDFRAC`).

   Make sure you use the `BTQ datum` datum link model so you can view MBT. As a result of this, one of the response features, `FX_0`, models residual fraction at MBT. You will use this in your optimizations in CAGE.

2. The process of searching for good model fits often results in a large selection of models. Remember to discard all but the models you chose as best, by selecting **File > Clean-up Tree**.

Look at the example finished project, `Gasoline_project.mat`, in the `mbctraining` directory, to see how the example models have been constructed. You will use these example models in CAGE for the optimization section.

Using Validation Data

After you have created models, you can use the validation data.

To import and filter the validation data set,

1. Select the top project node in the Model Tree.
2. Double click the `Data Object` in the Data Sets pane. The Data Editor appears.
3. Click the **Storage** button in the toolbar.

   Click **Store current filters** and **Store current test filters** in the toolbar. This allows you to reuse these filters in another data object without having to recreate them.
4. In the Model Browser, click **New Data** in the toolbar.
5. Click the Open File icon in the toolbar to load data from a file.

   The Data Import Wizard appears to select a file.
6. Use the Browse button to find and select the `DIVCP_Validation_DoE_Data.xls` data file in the `mbctraining` folder. Double-click to load the file, and click **Next**.
7. The Data Import Wizard displays a summary screen. Click **Finish** to accept the data.
8. You need to define test groupings as before.
   a. Select **Tools > Change Test Groupings** (or use the toolbar button)
   b. In the Test Groupings dialog box, clear the check box **One test/record**.
   c. Locate and double-click `GDOECT` in the **Variables** list box. `GDOECT` appears in the left list.25 tests are defined. Close the Define Test Groupings dialog box.
9. To apply the same filters as before, click the **Storage** button in the toolbar.

In the Storage dialog box, select the filter and test filter objects in turn and click **Append stored object** in the toolbar. This applies these filters to the current data object. Return to the Data Editor window and check the filters in the Filter List and Test Filter List View.

10 Close the Data Editor.

To attach the data set to your test plan for validation:

1 At the test plan level, select **TestPlan > Validation Data**. The Select Data for Validation wizard appears.

2 Select the validation data set (observe the number of tests), and click **Next**.

3 By default all tests are selected on the next screen, so click **Finish** to use all the tests to validate models in this test plan.

The validation data set appears in the right information pane for the test plan. Validation RMSE is automatically added to the summary statistics for comparison in the bottom list view of response models in the test plan.

You can now use the validation data to validate all models except response features. You can see validation statistics in the following places:

- Model List — `Validation RMSE` appears in the summary statistics in the lower list of models at the test plan, response and one-stage nodes
- At the local node view:
  - Pooled Statistics — `Validation RMSE` — The root mean squared error between the two-stage model and the validation data for all tests
  - Diagnostic Statistics > Local Diagnostics — Local model `Validation RMSE` for the currently selected test (if validation data is available for the current test—global variables must match)
- Summary Table — `Validation RMSE` for one-stage models

View validation plots in the following places:

- Plots of `Validation residuals` — For local models

- From any model node except response features, you can select **Model > Evaluate > Validation Data** to open the Model Evaluation window and investigate the model with the selected validation data.

  Check the model trends in the cross section view. Check the fit against the validation data. Is the fit acceptable? Do you need more data? Can you improve the fit with the existing data? Pay attention to the boundary model on the plots. Yellow areas are outside the boundary. Focus on the model trends only within the boundary model.

Exporting the Models

You can export models to file, to CAGE, to the workspace or to Simulink software. This capability makes models easy to share across engineering groups.

You can use **File > Export** for all of these options. The models exported depend on the model node you have selected in the tree. To export all models in the test plan, select the test plan node.

Note For this example, you use CAGE to import the models.

You use the example models in the CAGE part of the Model-Based Calibration Toolbox product to produce optimized calibration tables. Proceed to the next section for instructions.

Optimized Calibration

In this section...

"Problem Definition" on page 2-64

"Benefits of Automated Calibration" on page 2-65

Problem Definition

This section describes the creation and optimization of calibration tables for the gasoline case study. The Model Browser section of this case study covers creating the design for the experiment and creating and evaluating models from the resulting data. You can export your models directly to CAGE; or to Simulink software or to a file, ready to be imported into CAGE for model-based calibration generation. An example file is provided.

The aim of this case study is to produce optimized tables for

- Intake cam phase
- Exhaust cam phase
- Spark timing schedules

  as a function of load and rpm, subject to the following constraints

  - Constrain solutions to lie within the boundary constraint model
  - Constrain cam phase solutions so they do not change by more than 10^o between table cells (that is, no more than 10^o per 500 RPM change and per 0.1 load change).
  - Constrain residual fraction <= 25% at each drive cycle point (to ensure stable combustion). Residual fraction is the percentage of burned gas mass in the cylinder at intake valve close, relative to the total mass in the cylinder at intake valve close. Constraining maximum residual fraction is a simple and reasonable way of ensuring stable combustion. Residual fraction = 100 * Burned Gas Mass from Last Cycle / (Burned Gas Mass From Last Cycle + Fresh Air Mass)

CAGE is intended for model-based calibration, although you can still create tables without reference to models if you want. For this case study, you use

models produced in the Model Browser to generate calibrations in CAGE. You cover the following steps:

1. Load models of engine responses, decide on optimization strategy and define additional models. See "Importing Models into CAGE" on page 2-67.
2. Set up tables. See "Setting Up Calibration Tables to Fill" on page 2-70.
3. Define optimization objective and constraints. See "Setting Up the Optimization" on page 2-73.
4. Set up an operating point set for the optimization. See "Defining Variable Values" on page 2-78.
5. Run the optimization and view the results. See "Running the Optimization" on page 2-82
6. Duplicate and modify the optimization to create a sum optimization using the previous results as starting points. See "Setting Up the Sum Optimization" on page 2-85
7. Fill tables from optimization results. See "Filling Tables with Optimization Results" on page 2-92.
8. Use models and optimized tables to fill a spark estimator table. See "MBT Spark Estimator Problem" on page 2-94.

For guidance, you can look at the example finished project: `Gasoline_optimization.cag`.

Benefits of Automated Calibration

- You can move the table-filling process away from the test bed.
- You can regenerate calibrations when objectives, constraints, or calibration table layouts change, without additional testing.
- You can explore tradeoff possibilities interactively.
- You can produce initial calibrations using engine simulation software, before hardware is available.

CAGE can provide both automatic and interactive calibration optimization. You can trade off multiple objectives, deal with multiple constraints, and you can examine optimizations point-by-point or drive-cycle-based. You can use built-in optimization routines or write your own. You can fill groups of tables simultaneously, and optimize table values and breakpoint settings. CAGE can provide solutions for these example applications:

- Control problems
  - Injection timing and duration
  - EGR valve
  - Spark timing
  - Dual-independent variable valve timing
  - Emissions-constrained BSFC optimization over drive cycles
- Estimation problems
  - Torque
  - Emissions
  - Air flow and manifold pressure
  - Intake valve temperature
  - Borderline spark

Importing Models into CAGE

1. Start CAGE by typing `cage` at the MATLAB command line.
2. Select **File > Import From Project**. The CAGE Import Tool appears.
3. Click the **Import From Project File** button. Locate the example model file created in the Model Browser, `Gasoline_project.mat`, in the `mbctraining` directory, and click **Open**.
4. You can select from a list of models in the file. Select these models by **Ctrl**+clicking in the list:
   - `BTQ`
   - `EXTEMP`
   - `knot` (**Type** = `Datum`)

     In this case the `knot` model is duplicated because the datum model was used twice during modeling. The datum model tracked the maximum of the torque model, that is, MBT (the spark angle at maximum brake torque). This datum model was also used when modeling exhaust temperature and residual fraction, because it can be useful to see MBT on model plots for other factors. This is called a datum link model. You need only one copy of the `knot` model.
   - `FX_0` (Check that the **Location** specifies `RESIDFRAC`, *not* `EXTEMP`)

     The response feature model `FX_0` is `RESIDFRAC` at MBT. `RESIDFRAC` was constructed using a datum link (in this case `knot`, i.e., MBT) so `FX_0` is residual fraction evaluated at the datum point (MBT).

   Check that your selection matches those shown.

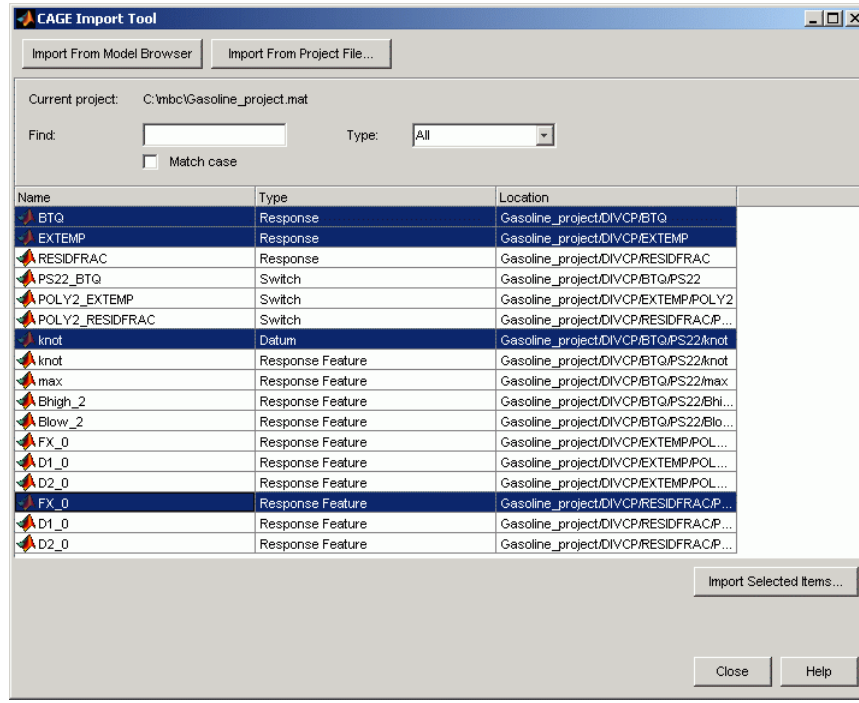

5 Click the **Import Selected Items** button .

6 In the following Import dialog,

- Double-click to edit the **CAGE Model Name** for the `knot` model to rename to `MBT`.
- Double-click to edit the **CAGE Model Name** for the `FX_0` model to rename to `RESIDFRACatMBT`.
- Click **OK** to import the models.

7 Close the Import Tool and click the Models button to select the Models view in CAGE. You should see the `BTQ`, `EXTEMP`, `MBT`, and `RESIDFRACatMBT` models in the list. The selected model is displayed in the other panes.

The objective is to produce optimized tables for spark and cam timings, subject to constraints on operating region, residual fraction, and the gradient of the

cam phasers over the calibrated tables. You want the optimization to search for the timings that give the best torque and minimum fuel consumption, subject to the constraints.

You could fix the spark timing to be MBT spark (the spark angle that produces maximum brake torque) by using the MBT spark model as the spark input to the other models in the optimization. With spark fixed at MBT, you can set up an optimization that allows the two valve timings to vary, exploring the space via a gradient descent algorithm to locate the optimal timings. Remember, this is a simplified problem and running an engine at MBT is knock-limited and so is not possible at all operating points. For this example, you will not use the MBT model as you will use spark as a free variable in the optimization. You will return to the MBT model later.

Note you can also:

- Export models directly from the Model Browser to CAGE when both are open
- Use the CAGE Import Tool to import models directly from the Model Browser
- Use the CAGE Import Tool to import models and other calibration items (tables, optimizations, tradeoffs, data sets, features) from any project file created in CAGE or the Model Browser. This can help you use existing projects to speed up the setup of new sessions.
- Export models to a file which you can import to CAGE by selecting **File > Import > Model**

Setting Up Calibration Tables to Fill

In this section...
"Altering Variable Ranges" on page 2-70
"Setting Up Tables" on page 2-70

Altering Variable Ranges

Editing variable ranges before creating tables will produce easy-to-read values (e.g., 0.1, 500, 0.2, 1000, etc.) for the table breakpoints when CAGE automatically spaces the normalizer values across the variable ranges.

1. Click the Variable Dictionary button to switch to the Variables view.
2. Select `N`, edit the range to `500` for the minimum, `5000` for the maximum, and press **Enter**.
3. Select `L`, edit the range to `0.1` for the minimum, `1` for the maximum, and press **Enter**.

Setting Up Tables

1. Set up a new 2-D table. Select **File > New > 2D Table.**
2. Name the table `SPK`.
3. Select `L` and `N` from the **Y** and **X Input** drop-down menus, , as shown.

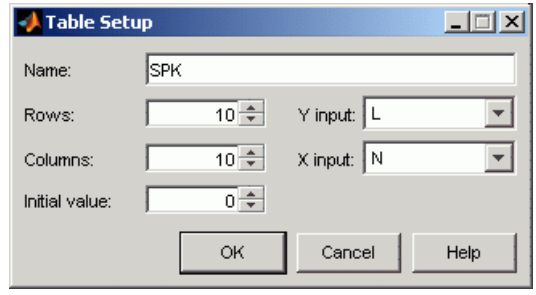

4 Leave the number of rows and columns at `10`, and leave `0` for the initial value.

5 Click **OK** to create the table.

CAGE switches to the Tables view, where you can see the new table and its normalizers in the Tables tree on the left. CAGE has automatically initialized the normalizers by spacing the breakpoints evenly across the range of the input variables N (speed) and L (load). Note that the normalizers appear as calibratable items in their own right, and as descendants (child nodes) of their tables.

Once you have created a table, you can duplicate it to create more tables that share the same breakpoints.

6 To set table limits, click to select the `SPK` table, then select **Table > Properties**. In the Table Properties dialog box, enter `0` and `50` in the **Table value limits** edit boxes. Click **OK**.

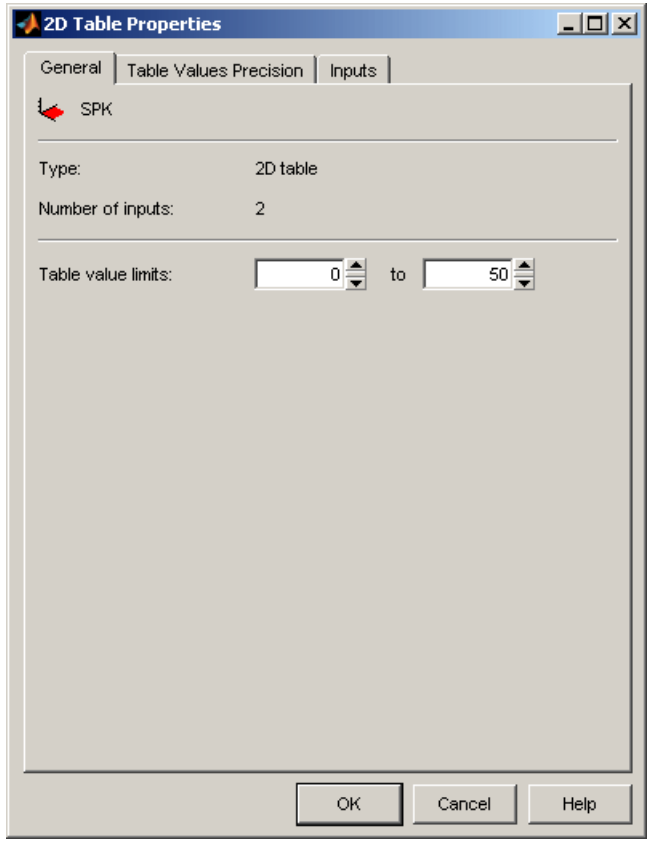

7 Right-click `SPK` in the tree and select **Duplicate SPK**.

8 Click to select the new table, press **F2**, and rename the table `INTCAM`.

9 Select **Table > Properties**. In the Table Properties dialog box, edit the lower limit to `-5`. Leave the upper limit at `50`, and click **OK**.

10 Right-click `INTCAM` in the tree and select **Duplicate INTCAM**.

11 Click to select the new table, press **F2**, and rename the table `EXHCAM`.

Now you have tables for optimized spark and cam timings, ready to fill with optimization results.

Setting Up the Optimization

CAGE provides a flexible optimization environment. You can define the objectives, the constraints, and the points where the optimization is carried out.

The objective is to maximize the weighted sum of torque over a set of [N, L] points subject to a set of constraints. This is a constrained single objective optimization problem.

You solve this problem in two parts:

1. Run an initial optimization to explore the problem point-by-point in order to obtain good starting points for the sum optimization.
2. Use the solutions from the initial optimization as the start points to run the sum optimization over the drive cycle to find optimal settings of spark (SPK) and cam (ICP, ECP) timing. Use these results to fill calibration tables.

CAGE has several built-in optimization routines and the capacity for you to write your own; in this case study you use `foptcon`. This is a modified version of `fmincon` from the Optimization Toolbox product. In CAGE, you can use the algorithm to minimize or maximize.

1. Select **File > New > Optimization**. The Optimization Wizard appears.

2 `foptcon` is selected by default, and this is the optimization algorithm you will use for this example. Note that this algorithm specifies a single objective in the **Objectives** column. Click **Next**.

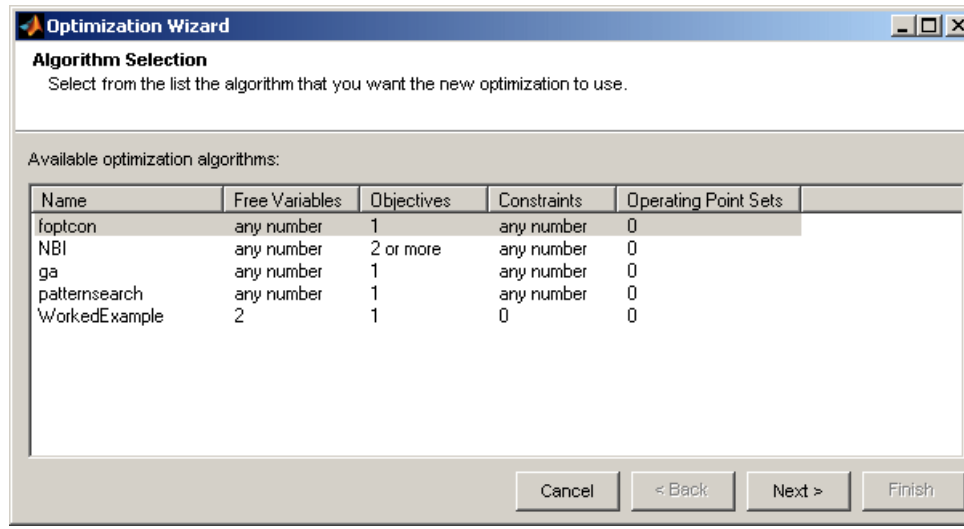

3 Increase

a The number of free variables to 3

b The number of constraints to 2

Click **Next**.

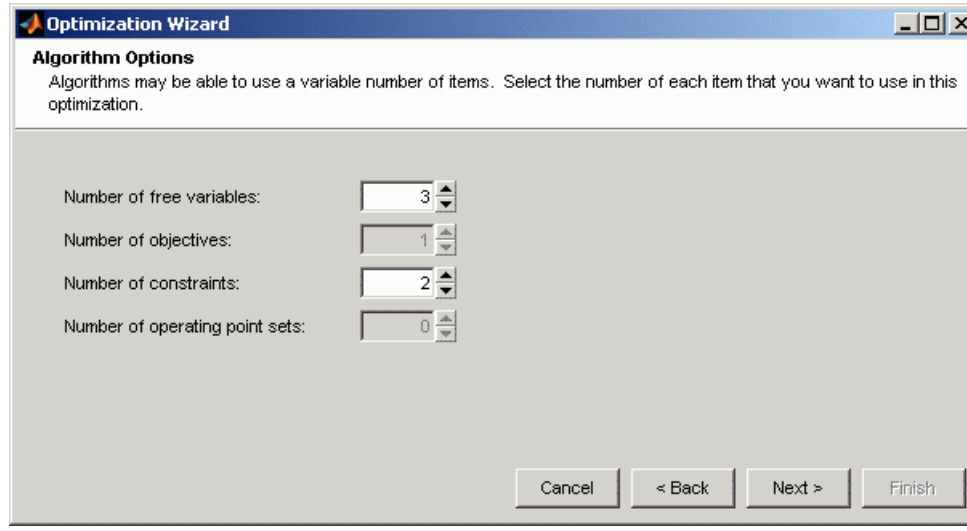

4 Select `S`, `ICP` and `ECP` in turn for the free variables, for each click the button to select them, and click **Next**.

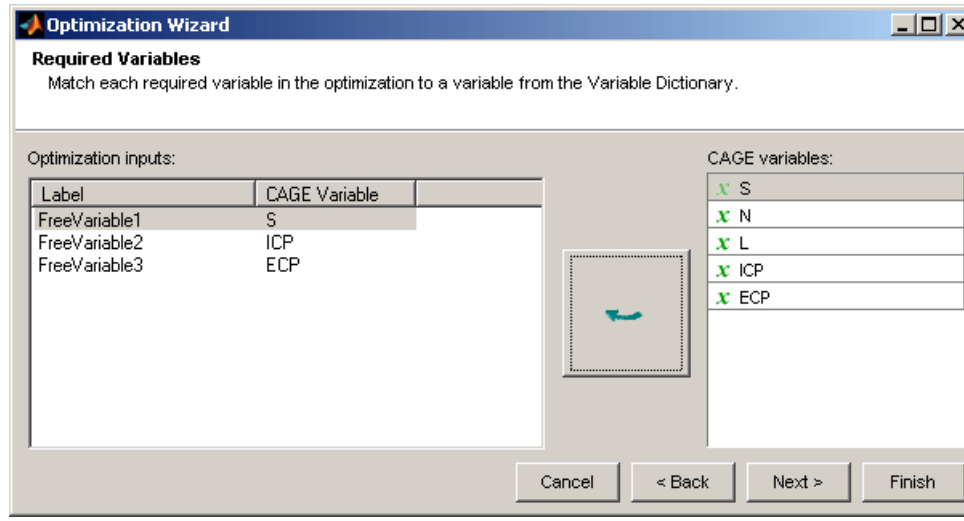

5 Select `BTQ` from the list of models on the right and click to select it for the objective. Select the **Maximize** radio button, and click **Next**.

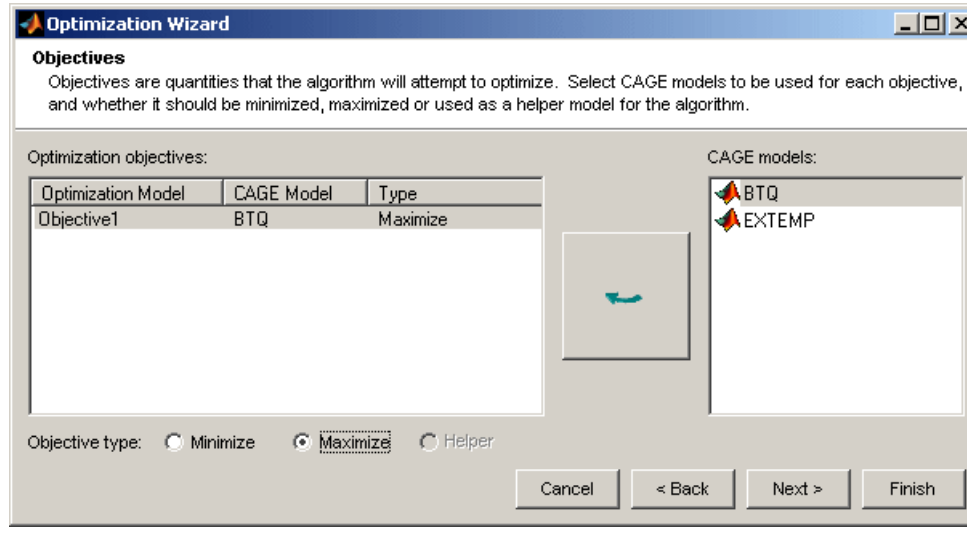

6 Select `RESIDFRACatMBT` from the list of models on the right and click to select it for the first constraint. With the `RESIDFRACatMBT` constraint selected, enter `25` in the edit box and press **Enter** (to constrain the optimization to stable combustion regions). Ensure that the expression reads `RESIDFRACatMBT <= 25`

7 Similarly, select `MBT` from the list of models on the right and click to select it for the second constraint.

Select the option button **Boundary of model** to use the boundary model to constrain the optimization.

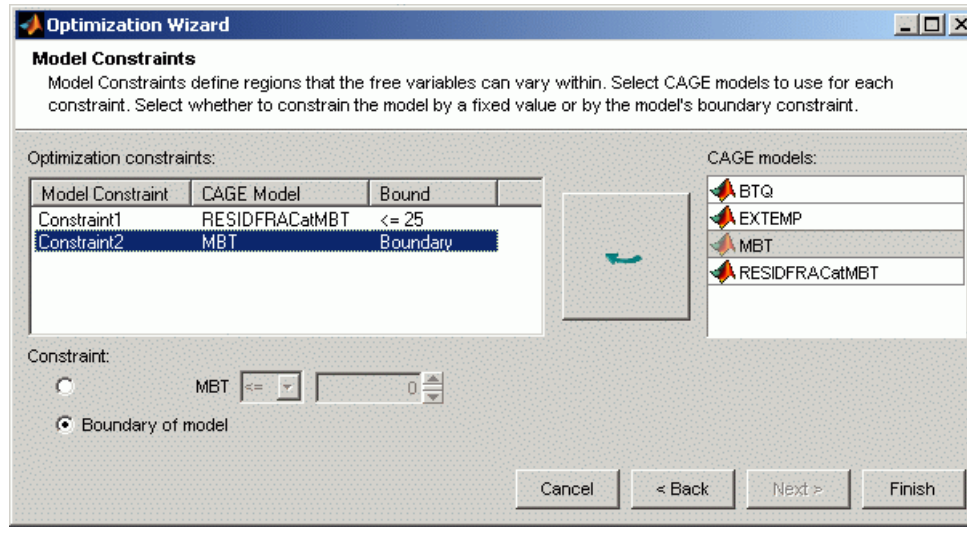

Click **Finish**.

CAGE switches to the Optimization view. Look at the information displayed. Here you can examine and edit the objective and constraint.

Defining Variable Values

You need to define the set of points where you want the optimization to run. In this case, you run the optimization over the table breakpoints in the tables. You can edit values manually in the **Input Variable Values** pane in the Optimization view, or you can import values from existing optimization output values, tables or data sets.

Edit set points of variables as follows.

1. Click the **Variable Dictionary** button to go to the Variable Dictionary view.
2. Click to select the variable `ECP` and edit the **Set Point** to `5`.
3. Click to select the variable `ICP` and edit the **Set Point** to `40`.

To import table breakpoints and initial values to your optimization, you will first create a feature to fill a SPK table from the MBT model. You can then use this table to initialize the starting values for your optimization.

To create the feature table,

1. Click the **Tables** button to return to the Tables view, and right-click to duplicate the `SPK` table. Rename the new table `SPK_Initial`.
2. Select **File > New > Feature**. CAGE switches to the Feature view.
3. Select **Feature > Graphical Strategy Editor**. Three Simulink windows appear.
4. In the CAGE project library window, double click the Tables block. The Tables library window opens.
5. Drag the `SPK_Initial` table block from the Tables library to the New Feature window.
6. Click to select `SPK_Initial`, then **Ctrl**+click the blue `New_Feature` block to connect the two.
7. Double-click the blue `New_Feature` block to import the strategy into CAGE.

8 In the Feature view, click the **Select Model** button.

9 In the Select Model dialog, select `MBT` and click **OK**.

10 Click **Fill Feature** in the toolbar. The Feature Fill Wizard appears.

11 Click **Next** three times, then click the **Fill Tables** button. Click **Finish** when the process has finished.

12 Expand the New_Feature node in the tree and select the `SPK_Initial` node to view the filled table. In the comparison pane you can see where the table limits cause the values to diverge from the feature model at low load.

13 Click the **Optimization** button to return to the Optimization view.

14 Select **Optimization > Import from Table Grid**.

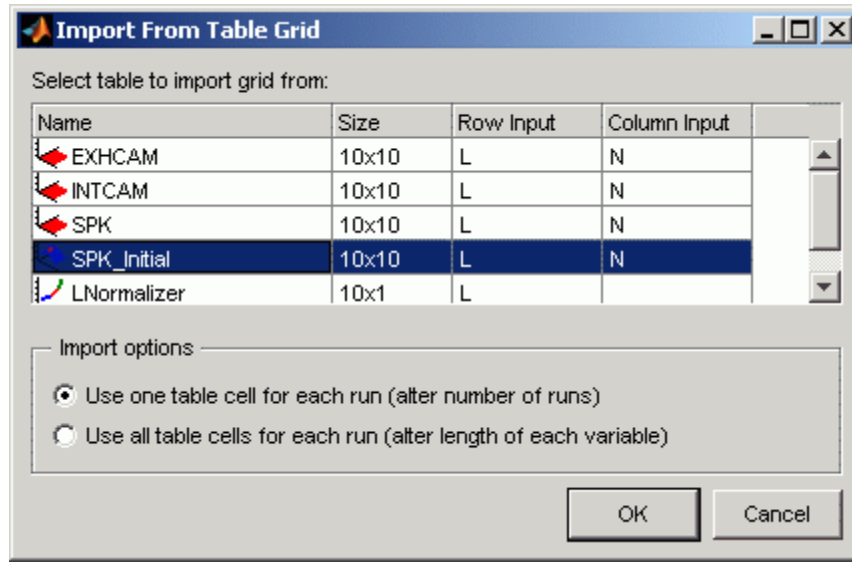

Select the `SPK_Initial` table and click **OK** to import the table grid into the optimization variable values.

15 Select **Optimization > Import from Table Values**.

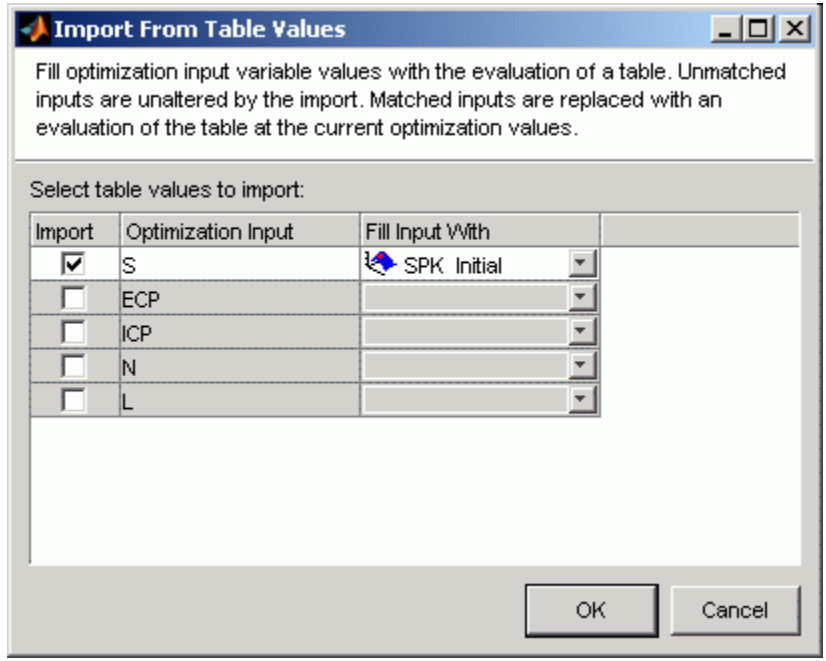

Select the check box to import values for `S`, and select `SPK_Initial` from the **Fill Input With** list. Click **OK** to import the table values of spark into the optimization initial values for `S`.

View the values in the **Input Variable Values** pane in the Optimization view. These values define the variable values at each point where you want the optimization to run. The fixed variable values are the table breakpoints, and the free variable initial values are the filled table values of spark, and the defaults (set point) for the other free variables.

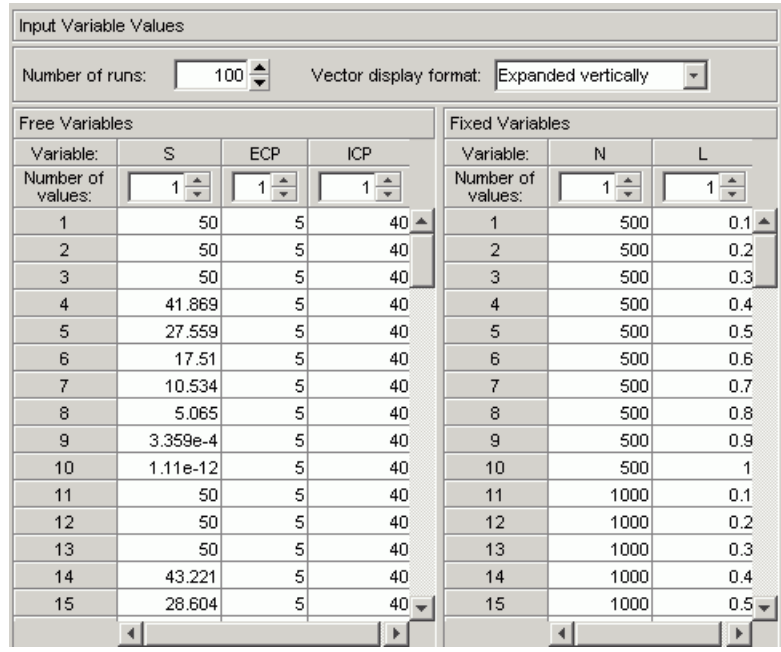
Input Variable Values
Number of runs: 100
Vector display format: Expanded vertically
Free Variables
Variable: S ECP ICP
Number of values: 1 1 1
1 50 5 40
2 50 5 40
3 50 5 40
4 41.869 5 40
5 27.559 5 40
6 17.51 5 40
7 10.534 5 40
8 5.065 5 40
9 3.359e-4 5 40
10 1.11e-12 5 40
11 50 5 40
12 50 5 40
13 50 5 40
14 43.221 5 40
15 28.604 5 40
Fixed Variables
Variable: N L
Number of values: 1 1
1 500 0.1
2 500 0.2
3 500 0.3
4 500 0.4
5 500 0.5
6 500 0.6
7 500 0.7
8 500 0.8
9 500 0.9
10 500 1
11 1000 0.1
12 1000 0.2
13 1000 0.3
14 1000 0.4
15 1000 0.5

Running the Optimization

You have defined objectives, constraints and a set of operating points. Your optimization is ready to run.

1. Click Run Optimization in the toolbar.

2. When the optimization is complete, the view switches to the new child node, `Optimization_Output`, in the Optimization tree under the `Optimization` node. View the results.

3. Look through the solutions at different operating points by clicking cells in the output table. Regions that do not meet the constraint are yellow in the plots and tables.

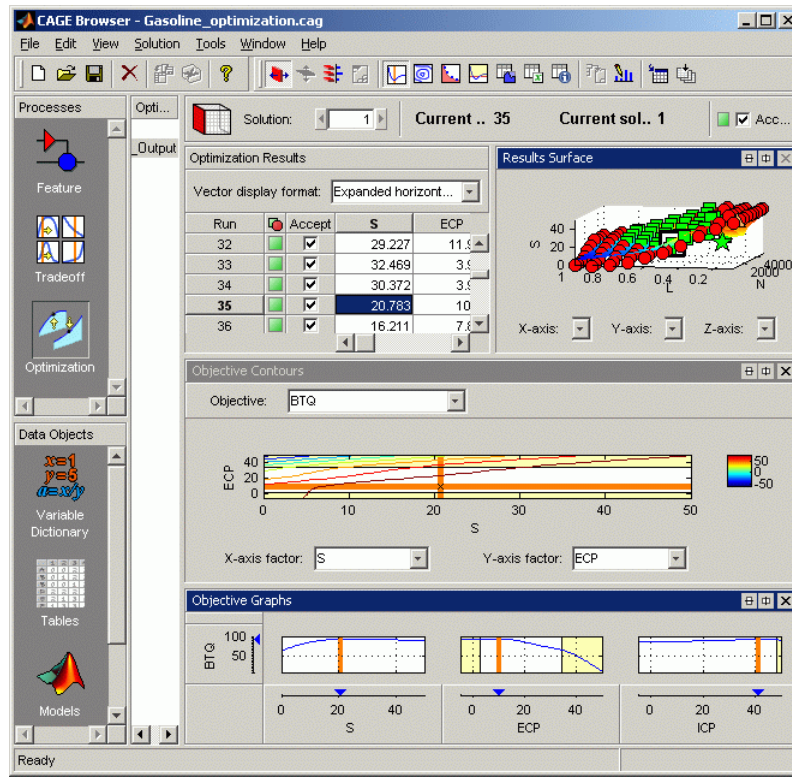

Because this is a single-objective optimization, there is only one solution at each point.

4 Observe the colored icons that indicate the Accept status of each solution plotted in the Results Surface view, and of each corresponding row in the Optimization Output Values table.

CAGE automatically selects the Accept check boxes in the table for solutions where the algorithm exit flag indicates success (>0). These solutions show a green square icon in the table and surface. CAGE also highlights unsuccessful solutions for you to investigate.

Click in the table or surface view to select solutions, and other plots (such as the Objective Contours and Objective Graphs) update.

5 Right-click the title bar of the Results Surface view and select **Current View > Results Contour**. Observe the Accept icons for the solutions.

6 Examine the orange triangle Accept status solutions. This indicates the algorithm exit flag is zero. This means the algorithm terminated because it exceeded limits on the amount of computation allowed (e.g., 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. Do not edit any check boxes yet.

The Objective Contours plot of `ECP` against `ICP` is useful for analyzing results. Also check the Objective Graphs for spark to ensure that maximum torque is obtained.

7 Export all acceptable results to tables to check surfaces. Click Fill Tables in the toolbar and use the wizard. If you need instructions, see “Filling Tables with Optimization Results” on page 2-92.

Check trends, and don’t worry about smoothness. You will apply table gradients constraints to smooth the table surfaces in the next optimization.

8 You can change the Accept selections using the check boxes for each solution, or by right-clicking points in the surface view. You can use these Accept selections to choose solutions within the table for use in filling tables, exporting to data sets and importing to other optimization starting values.

Reject point 22 (clear the check box). This is in the wrong region with respect to adjacent points in the table.

9 You may want to turn off the display of constraints in the Objective Contours plot before running the next sum optimization or viewing the output node for the sum optimization. This will save time displaying the output views. Right-click the contour plot and toggle **Display Constraints** off.

Setting Up the Sum Optimization

In this section...

"Setting Up the Optimization Initial Values and Objective" on page 2-85

"Creating Table Gradient Constraints" on page 2-88

"Running the Sum Optimization" on page 2-91

Setting Up the Optimization Initial Values and Objective

The aim of this case study is to produce optimized tables for

- Intake cam phase
- Exhaust cam phase
- Spark timing schedules

  as a function of load and RPM, subject to the following constraints:

  - Constrain solutions to lie within the boundary constraint model
  - Constrain cam phase solutions so they do not change by more than 10^o between table cells (that is, no more than 10^o per 500 RPM change and per 0.1 load change.)
  - Constrain residual fraction <= 25% at each drive cycle point

The optimization objective is to maximize the weighted sum of torque over a set of [N, L] points subject to the constraints described.

The initial point-by-point optimization was designed to obtain good starting points for the sum optimization. Now, you can duplicate and modify that initial optimization to run the sum problem. You use the solutions from the initial optimization as the start points to run the sum optimization over the drive cycle to find optimal settings of spark and cam timings. You use these results to fill calibration tables.

1 Right-click the initial optimization node, and select **Duplicate**.

2 Click to select the new optimization, press **F2**, and enter `Sum_optimization` as the new name.

3 Import initial values from the output of the first optimization. Select **Optimization > Import From Output**.

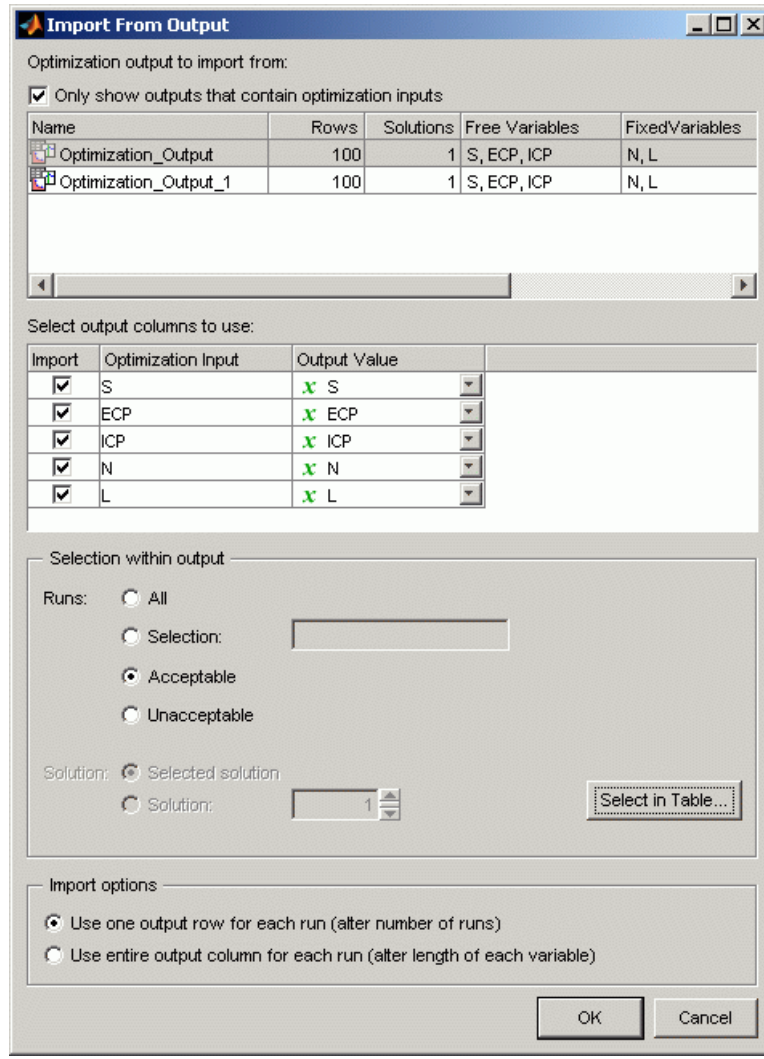

a Select the first `Optimization_Output` (there are two because the duplicate optimization also contains an output node).

b In the **Selection within output** pane, select the **Acceptable** option button to use only the solutions marked as acceptable.

c To view the selected Acceptable solutions, click the button **Select in Table**. Review—but do not change—the selected solutions.

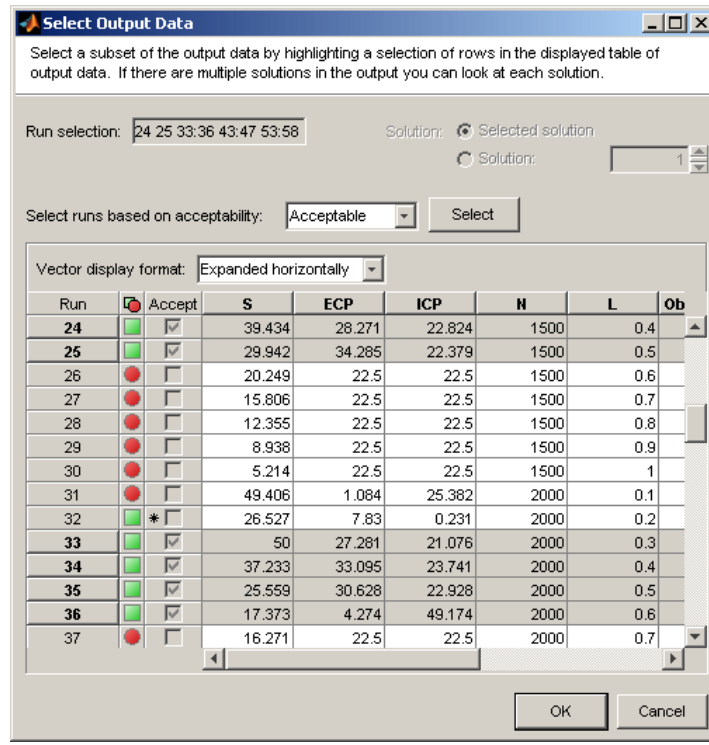

Click **OK** to dismiss the Select Output Data dialog box.

d Click **OK** to close the Import From Optimization Output dialog box and import the acceptable output values for all the free and fixed variables. These values are the starting values for the new optimization. Observe the new values in the Input Variable Values pane in the Optimization view.

4 Select **Optimization > Convert to Single Run**. In the Input Variable Values pane, there is now a single run with 42 values, instead of the previous 42 runs of single values.

5 To run a sum optimization over the drive cycle, convert the objective to a sum objective. Double-click `BTQ` in the Objectives pane to open the Edit Objective dialog box.

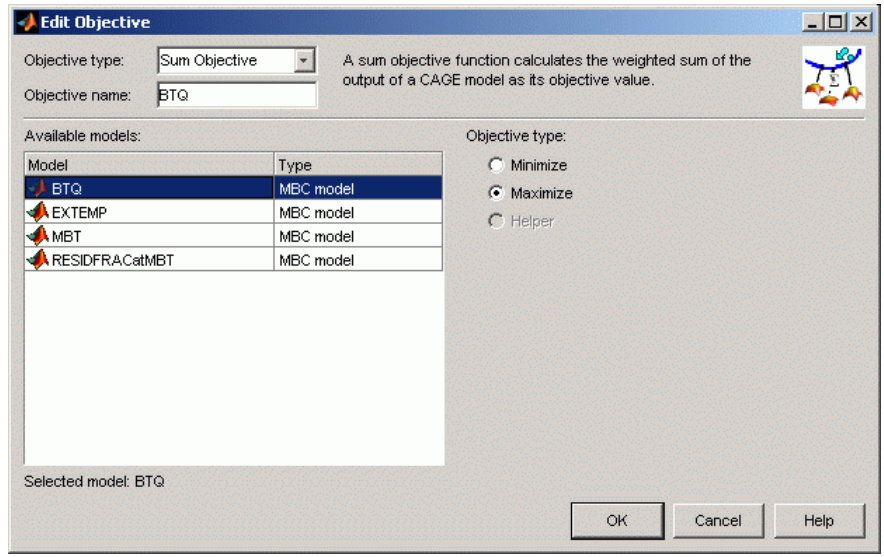

a Select `Sum Objective` in the **Objective type** drop-down menu.

b Select `BTQ` from the **Model** list, and verify that **Maximize** is selected for the objective type.

c Click **OK** to close the dialog box.

Now, your sum optimization is ready to run. Before you do this, see the next section to create table gradient constraints.

Creating Table Gradient Constraints

To meet engineering constraints it is essential that some calibration tables meet a level of smoothness. For example, large cam phase changes between adjacent table cells cannot be used because the controls cannot physically

move that fast. Achieving a level of smoothness is important for variable valve timing calibration problems. Unconstrained optimizations may find solutions where the cam timings change too rapidly across tables. You can apply table gradient constraints to make sure the resulting tables are smooth enough.

To apply a gradient constraint to restrict the change in `EXH` between adjacent table cells:

1. Right-click in the Constraints pane, and select **Add Constraint**. The Edit Constraint dialog box appears.

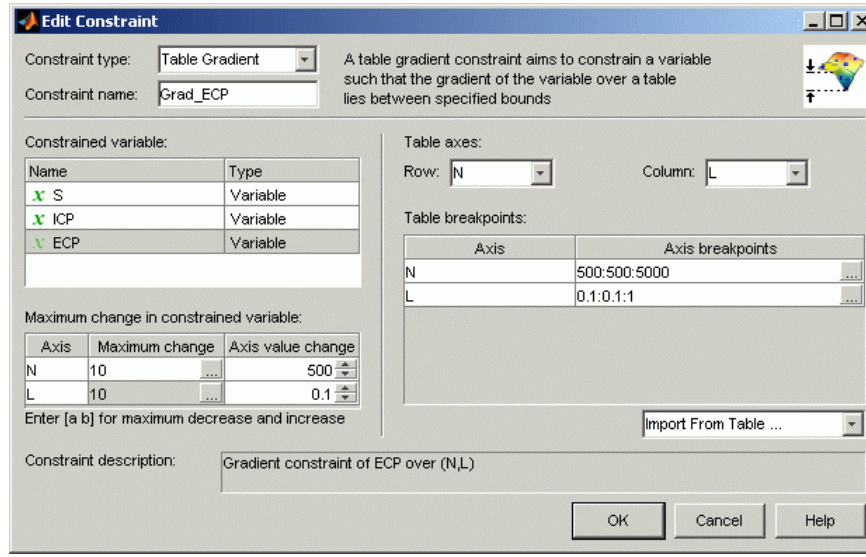

2. Select `Table Gradient` from the **Constraint type** drop-down menu.

3. Enter `Grad_ECP` in the **Constraint name** edit box.

4. Select the variable `ECP` to constrain.

5. Select the table axes `N` (from the **Row** drop-down) and `L` (from the **Column** drop-down).

6 Select `EXHCAM` from the **Import From Table** drop-down menu. The breakpoints of this table now appear in the breakpoints edit boxes: `500:500:5000` in N, and `0.1:0.1:1` in L.

7 Apply the gradient constraint across the breakpoints of the table. As shown in the following figure, enter `10` in the N axis **Maximum change** edit box, and verify `500` is in the **Axis value change** edit box. This specification constrains the maximum change to 10 degrees `ECP` per 500 change in `N` (the difference between cells, as shown in the axis breakpoints edit box).

8 Next, enter `10` in the L axis row **Maximum change** edit box, and verify that `0.1` is shown in the **Axis value change**, to constrain the maximum change to 10 degrees per 0.1 change in `L`.

9 Click **OK** to close the Edit Constraint dialog box.

Similarly, add another gradient constraint to restrict the change in `INT` between adjacent table cells as follows:

1 Right-click in the Constraints pane and select **Add Constraint**. The Edit Constraint dialog appears.

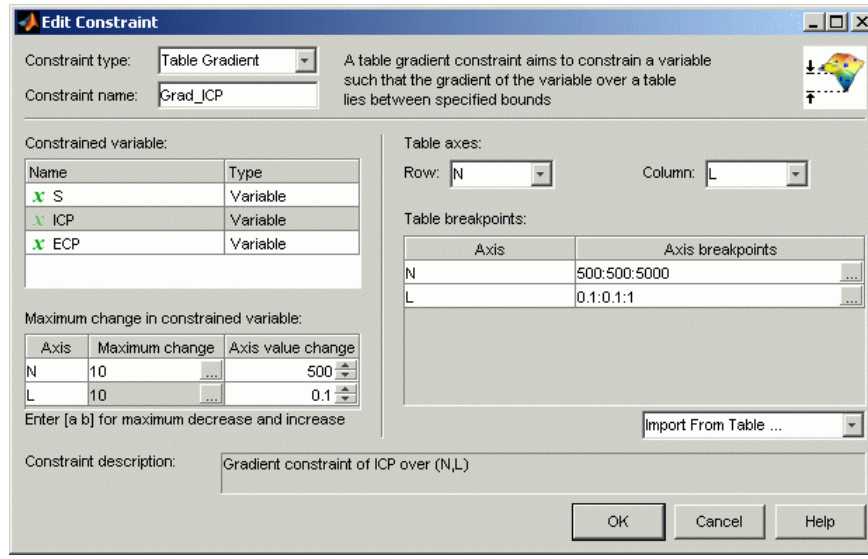

2. Select `Table Gradient` from the **Constraint type** drop-down menu.

3. Enter `Grad_ICP` in the **Constraint name** edit box.

4. Select the variable `ICP` to constrain.

5. Select the table axes `N` (from the **Row** drop-down) and `L` (from the **Column** list).

6. Select `INTCAM` from the **Import From Table** drop-down menu. The breakpoints of this table now appear in the breakpoints edit boxes: `500:500:5000` in N, and `0.1:0.1:1` in L.

7. Apply the gradient constraint across the breakpoints of the table. As shown, enter `10` in the N axis **Maximum change** edit box, to constrain the maximum change to 10 degrees `ICP` per 500 change in `N` (the difference between cells, as shown in the axis breakpoints edit box).

8. Next, enter `10` in the L axis row, to constrain the maximum change to 10 degrees per 0.1 load change.

9. Click **OK** to dismiss the Edit Constraint dialog.

Running the Sum Optimization

To run the optimization:

1. Click Run Optimization in the toolbar.

2. When the optimization is complete, the view switches to the updated child node, `Optimization_Output_1`, in the Optimization tree under the `Sum_optimization` node. (Recall that you duplicated an optimization, including its output node). View the results. For guidance see “Interpreting Sum Optimization Output”.

Remember you can check the example project, `Gasoline_optimization.cag` from the `mbctraining` directory, to view completed examples of the optimizations and filled tables.

Filling Tables with Optimization Results

1 Select **Solution > Fill Tables**.

The Table Filling Wizard appears.

2 **Shift**-click to multi-select the `SPK`, `INTCAM`, and `EXHCAM` tables and click the button to add the tables to the filling list. Click **Next**.

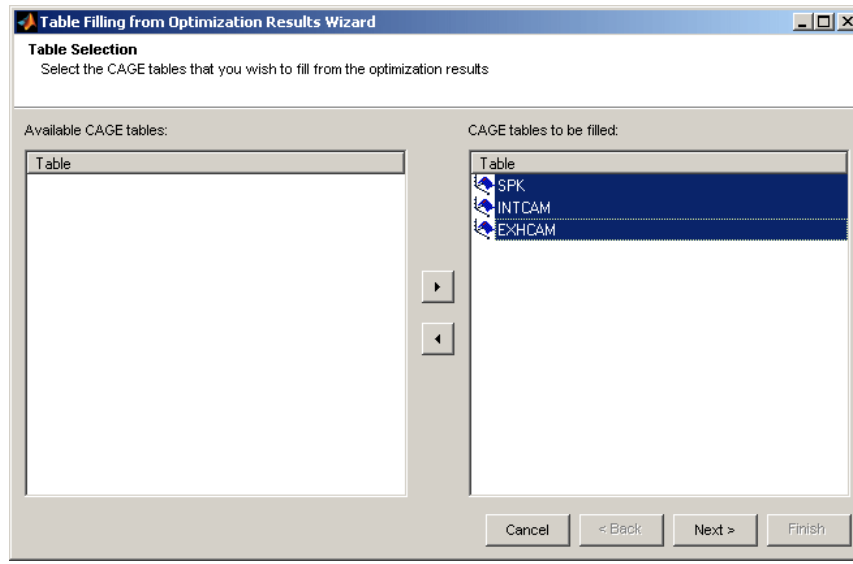

3 You need to select factors for table filling.

a Select the `SPK` table on the left, and double-click `S` in the right list of optimization results to select it for filling the table.

b The next table down (`INTCAM`) is automatically selected, so you can double-click `ICP` in the right list to fill this table.

c The next table down (`EXHCAM`) is automatically selected, so you can double-click `ECP` to fill this table.

Note that you could also choose the MBT model to fill a table, for example if you wished to compare your optimization results with the spark

timings specified by the MBT model. You will fill and optimize tables from models in more detail in the next section.

d Make sure the check box **Use acceptable solutions only** is selected.

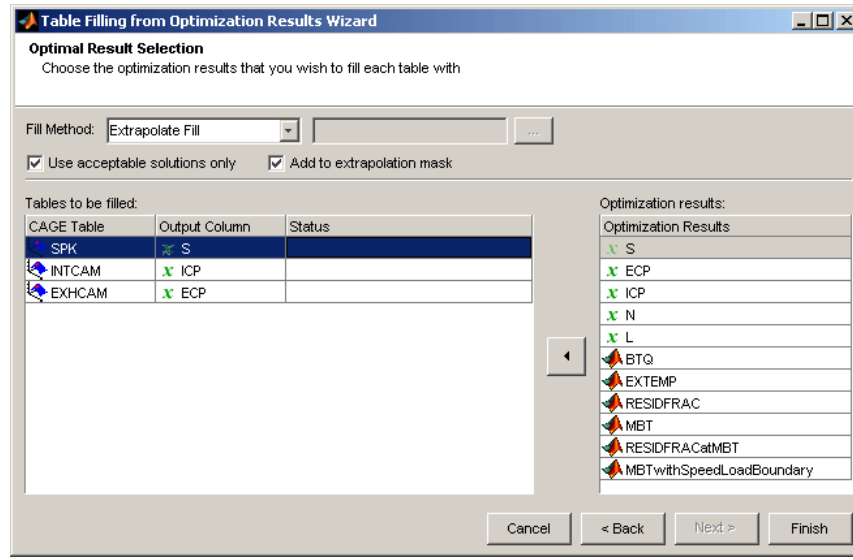

e Click **Finish** to fill the tables.

A dialog appears with the message that the tables have been filled successfully. Click **OK**.

4 Switch to the **Tables** view to view the filled tables. Click **Tables** in the **Data Objects** pane, and select the filled tables in turn.

MBT Spark Estimator Problem

In this section...

"What Is an Estimator?" on page 2-94
"View the Feature" on page 2-94
"View Variables" on page 2-95
"Edit and Import Boundary Model" on page 2-96
"Use the Feature Fill Wizard" on page 2-97
"Inspect Results" on page 2-101
"CAGE Import Tool" on page 2-104

What Is an Estimator?

The cam timings optimization results solved a control problem. You can also use CAGE to calibrate estimator problems. Here you will use an `MBT` model to produce a spark estimator feature which estimates MBT spark when each cam is on or off, using the cam timings found by the optimization.

View the Feature

1. Select **File > Open Project** and load the example project, `Gasoline_optimization.cag` from the `mbctraining` directory.

2. Click **Feature** in the **Processes** pane to go to the Feature view.

   The features `Exhaust_CAM`, `Intake_CAM`, and `MBT_Spark` are in the Feature tree. The `MBT_Start_Feature` was used to provide initial values for the optimization, as described in "Defining Variable Values" on page 2-78.

3. Select `MBT_Spark` and view the strategy by selecting **Feature > Graphical Strategy Editor**.

4. Examine the strategy model. Observe the Multiport Switch block which switches between MBT tables depending on whether each cam is on or off. The variables `Intake_On` and `Exhaust_On` are used to define whether cams are parked or not. All the MBT tables (with and without cams) share the same speed (`N`) and load (`L`) normalizers.

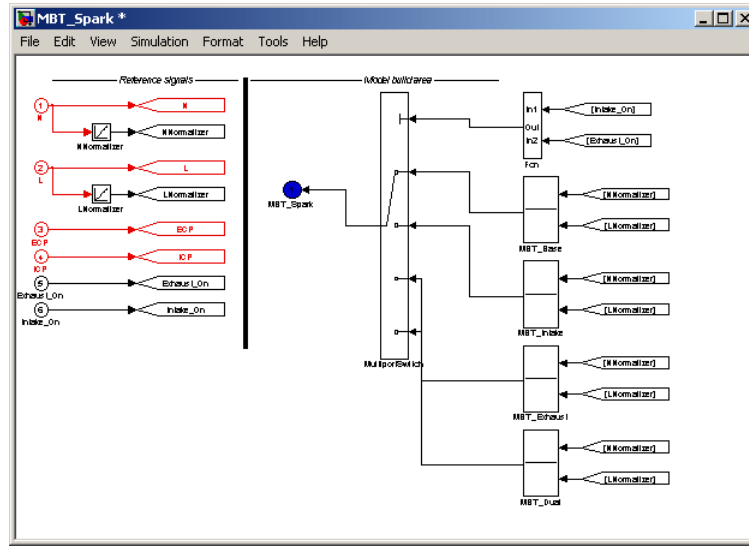

If you wanted to import this strategy model, you would double-click the `MBC_Spark` blue outport to parse the strategy into CAGE. This strategy has already been imported into this project so just close the model window.

5 Expand the `MBT_Spark` feature in the tree to see the MBT tables that have been created by importing the strategy: `MBT_Base`, `MBT_Intake`, `MBT_Exhaust`, and `MBT_Dual`. These tables share the same normalizers in speed and load.

6 Similarly view the strategies for the Exhaust and Intake CAM features. These features switch between parked cams and active cams depending on the threshold value. The CAM features define the cam inputs and switch between the optimal CAM tables (from the optimization results) and the parked values. You can use the linking functionality in the Feature Filling Wizard to connect features and tables to models.

View Variables

You need some variables and constants to define when the cams are parked so the strategy can switch between optimal cam timings and parked cams at a threshold value.

View how this is done:

1. Click **Variable Dictionary** in the **Data Objects** pane to switch to the Variable Dictionary view.
2. Click to select the new variable `Exhaust_On`.
3. Observe the **Minimum** is 0 and the **Maximum** is 1.
4. The variable `Intake_On` has the same values.

Also two constants define the parked cam positions:

1. Select `Exhaust_Parked`.
2. Observe the **Set Point** of this constant is `0`.
3. Select `Intake_Parked`; this has a **Set Point** of `0`.

Edit and Import Boundary Model

To edit the boundary model and import it to CAGE,

1. Open the Model Browser (enter `mbcmodel` at the MATLAB command line).
2. Load the example project. Select **File > Load Project**, locate and select `Gasoline_project.mat` in the `mbctraining` directory.
3. Select the `DIVCP` test plan node in the tree.
4. Select **TestPlan > Boundary Constraints** (or click **Edit boundary constraint** in the toolbar).

   The Boundary Editor opens.
5. Select the `Star shaped(N, L)` constraint in the tree, and click **Assign constraint as best** in the toolbar. This removes all the other constraints from the collection of best boundary models.
6. Close the Boundary Editor. In the Model Browser, observe the single constraint `Global/Star shaped(N, L)` listed under Boundary Constraint in the right information pane.

7 Return to CAGE and select **File > Import From Project**.

   The CAGE Import Tool appears. You can import directly from the Model Browser when it is open, and the CAGE Import Tool automatically displays the available items.

8 Select `knot` (the datum model) from the list.

9 Click the **Import Selected Items** button.

10 The Import dialog opens displaying the model you selected for import. Double-click the **CAGE Model Name** column cell to edit the name to `MBTwithSpeedLoadBoundary`, and click **OK** to import the model.

Use the Feature Fill Wizard

You can use the Feature Fill Wizard to fill and optimize the values in tables by reference to the model. 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.

1 Click **Feature** in the **Processes** pane to return to the Feature view, then select the feature node `MBT_Spark`.

2 Click in the toolbar, or select **Feature > Fill**. This opens the Feature Fill Wizard.

3 Select all four table check boxes to fill all tables.

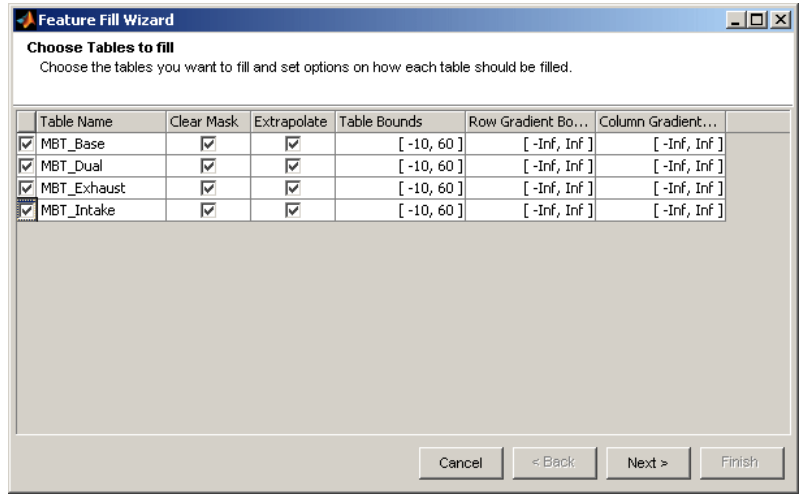

You could also explore setting gradient bounds to constrain table filling for smoothness.

This time leave the other settings at the defaults and click **Next**.

4 Choose filling model, constraint, and links.

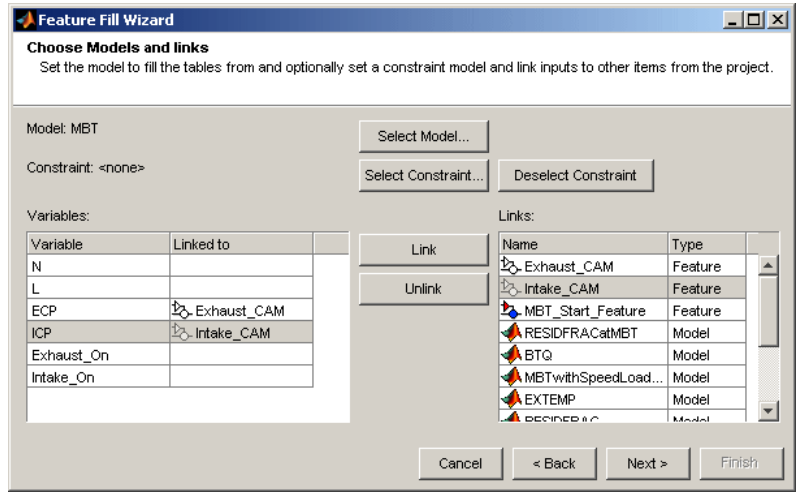

- Make sure `MBT` is the **Model** to fill the tables.

- Select `ECP` in the left **Variables** list and the `Exhaust_CAM` feature in the right **Links** list and click **Link**.
- Select `ICP` in the left list and `Intake_CAM` in the right list and click **Link**.
- Click **Select Constraint**. The Edit Constraint dialog box opens.

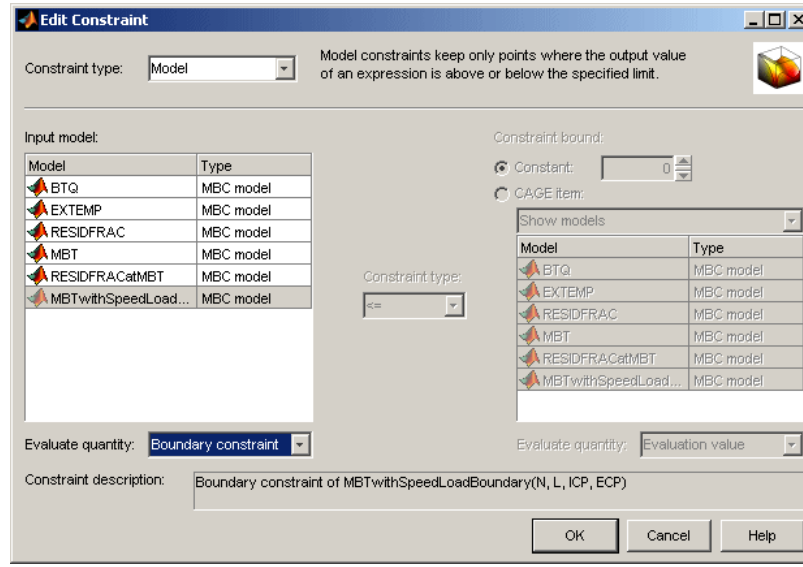

Select the model `MBTwithSpeedLoadBoundary`, select `Boundary constraint` as the **Evaluate quantity**, and click **OK**. This boundary model constraint ensures you only fill over speed/load points where the data was collected.

When you return to the Feature Fill Wizard, click **Next**.

5 Set values to optimize over.

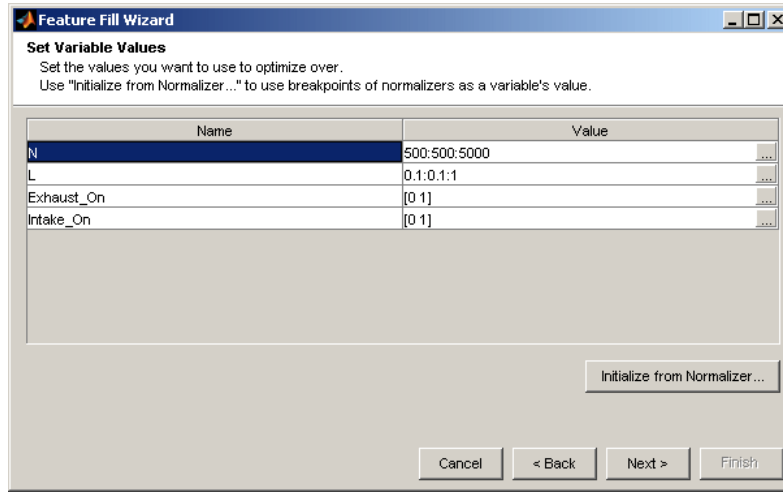

- Enter `0, 1` in the `Exhaust_On` **Value** edit box, and press **Enter**.
- Enter `0, 1` in the `Intake_On` **Value** edit box, and press **Enter**.

You use these values because the strategy includes the following tables.

Table	Intake_On Value	Exhaust_On Value
MBT_Base (cams parked)	0	0
MBT_Exhaust	0	1
MBT_Intake	1	0
MBT_Dual	1	1

The `N` and `L` normalizer values are automatically selected. You can edit normalizers manually, or you can click the **Initialize From Normalizer** button here to reach a dialog where you can select normalizers and 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. Leave the setting alone for now.

Click **Next**.

6 Fill tables and generate plots.

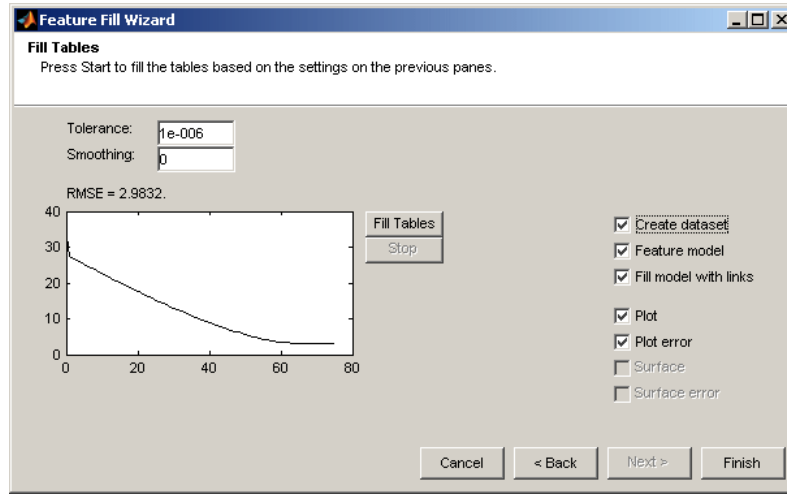

Click the **Fill Tables** button. Watch the graph as the optimization progresses.

When it is finished, select all enabled check boxes, and click **Finish**. Plots appear summarizing the feature fill data.

Inspect Results

Look at the filled tables, linked models and exported data set.

1 In the Feature view, select in turn the tables `MBT_Base`, `MBT_Intake`, `MBT_Exhaust`, and `MBT_Dual` in the feature tree.

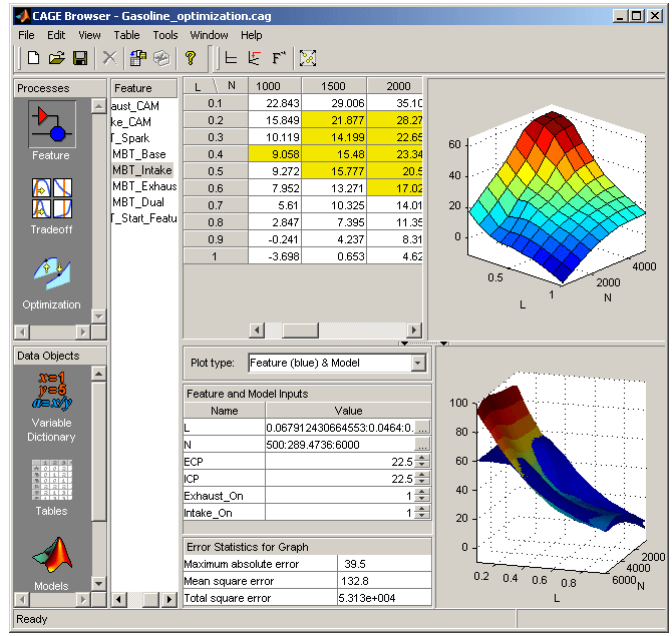

Observe the yellow mask area of cells in each table — the mask is defined by the limits from the `MBTwithSpeedLoadBoundary` boundary constraint model you selected. The rest of the table values are extrapolated after the mask cells are filled by the Feature Fill Wizard (specified by the **Extrapolate** check box). Table values are limited to [-10 60] as specified in the **Table Bounds** in the Feature Fill Wizard.

The lower comparison pane (if you have it open) does not change as you change table, because it displays a comparison between the whole feature and the model, not individual tables. Note if you use links the comparison pane is not showing a true comparison as the cam inputs are not constant. The comparison pane is showing the comparison using constant values for the cam timings.

2 Click **Models** to select the Models view. Look at the feature model and fill model `MBT_SPARK_Model` and `MBT_SPARK_FillModel`. If you can't view the whole Connections diagram of `MBT_SPARK_FillModel`, right-click and select **Zoom To Fit**.

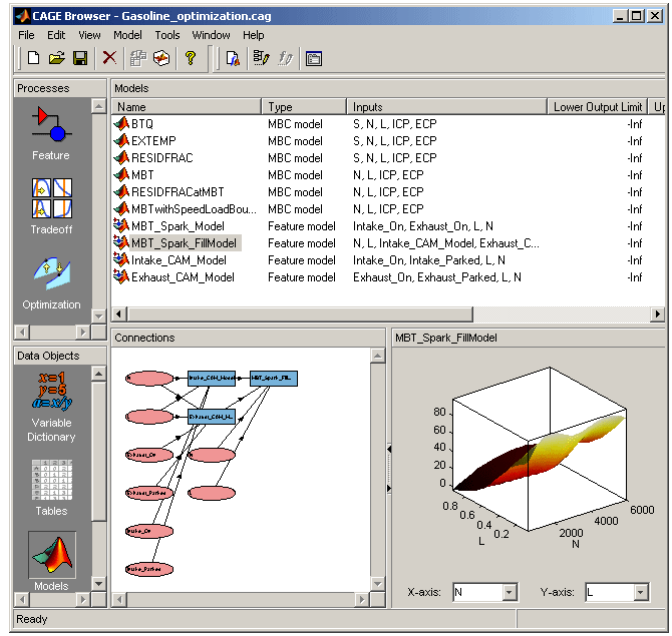

The linking functionality in the Feature Fill Wizard allows features and tables to be connected to models and any expression to be connected to feature inputs. These links can be made permanent by creating feature models and fill models that are static snapshots of the table on finishing the feature fill wizard (specified by the **Feature model** and **Fill model with links** check boxes on screen 4 of the Feature Fill Wizard).

Notice that the CAM features are converted to feature models (`Exhaust_CAM_Model` and `Intake_CAM_Model` are new feature models) and they are connected to `MBT_SPARK_FillModel`.

3 Click **Data Sets** to select the Data Sets view. Select the dataset `MBT_SPARK_FillResults` to study the gridded data, the model and feature values for the feature fill. Click View Data in the toolbar to see the data table view. All the links in the feature fill process are defined in this data set — try clicking column headers to see highlighted linked input columns.

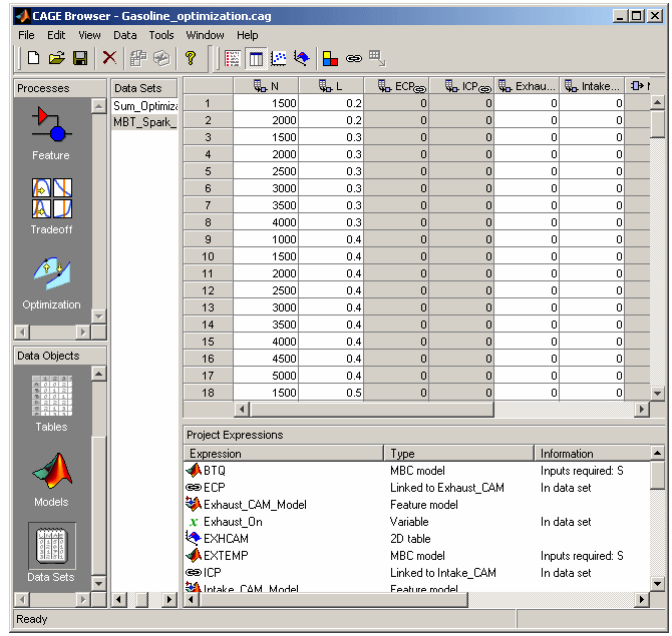

CAGE Import Tool

Suppose that you are calibrating a similar engine at a later date. You would build new models for this engine and then want to solve the same problems in CAGE. The CAGE Import Tool allows you to reuse the setup from your old CAGE session. All that is necessary to import new models on top of the existing ones and rerun the optimizations and feature fill problems. For example you could import new BTQ (replace BTQ), knot (replace MBT and MBTwithSpeedLoadBoundary) and EXTEMP models from the example file `Gasoline_Project.mat` as follows:

1. Select **File > Import From 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 CAGE Import Tool automatically shows the items available. Use the **Import From Project File** button to select a file.

If you are choosing a project file, a file browser dialog opens. Locate `Gasoline_Project.mat` and click **Open**.

3. The CAGE Import Tool displays the available items. Select the items you want to import from the list: BTQ, EXTEMP, and knot.

4. Click the **Import Selected Items** button.

5. The Import dialog opens displaying the items you selected for import.

   - Double-click the **CAGE Model Name** column cells to edit item names.
   - Choose to replace BTQ, EXTEMP and MBT. For knot, select `Replace` from the **Action** list, then double-click the **CAGE Model Name** column cells to open a dialog to select the correct item, MBT, to replace.
   - Click **OK** to import the items.

Now you can run the optimization again to generate new optimal CAM timings with new models. Export the optimization results to the INTCAM and EXHCAM tables, and use the Feature Fill Wizard to fill the MBT_SPARK strategy using the same steps as before.

3

Diesel Engine Calibration Case Study

This case study provides a step-by-step guide to using the Model-Based Calibration Toolbox product to solve a diesel engine calibration problem. This section includes the following topics:

- "Diesel Case Study Overview" on page 3-2
- "Design of Experiment" on page 3-4
- "Modeling" on page 3-17
- "Optimized Calibration" on page 3-31
- "Importing Models of Engine Responses into CAGE" on page 3-34
- "Defining Additional Variables and Models" on page 3-35
- "Setting Up Calibration Tables to Fill" on page 3-37
- "Setting Up the Optimization" on page 3-38
- "Setting Up Constraints" on page 3-42
- "Defining Variable Values" on page 3-47
- "Running the Optimization" on page 3-50
- "Setting Up the Sum Optimization" on page 3-52
- "Filling Tables with Optimization Results" on page 3-59

Diesel Case Study Overview

In this section...

"Introduction" on page 3-2

"Problem Definition" on page 3-2

Introduction

This case study is an example of a diesel engine control calibration, for a six-cylinder 9.0 L common-rail diesel engine with VGT (variable geometry turbo) and cooled EGR (exhaust gas recirculation). It is applied in an off-road application with a very narrow engine speed range from 1600 to 2200 RPM. The aim of the case study is to produce optimal SOI (start of injection), base fuel, VGT, and EGR calibration schedules as a function of commanded torque and RPM. It involves models for torque, peak pressure, equivalence ratio, exhaust temperature, VGT speed, and EGR mass fraction. The optimization setup in CAGE is based on an 8-mode off-road emission test, approximated to 7 mode points by neglecting the idle operating point of the engine.

The Model Browser part of the example takes you through the following steps:

1. Create a design for your experiment "Design of Experiment" on page 3-4
2. Create models from the collected data "Modeling" on page 3-17

The Model Browser section of the case study covers design of experiment, data handling, and model construction and export. In the later CAGE browser section of the case study you use the models to complete the optimization of the calibration tables, see "Optimized Calibration" on page 3-31.

Problem Definition

Produce optimal calibration tables in speed and torque for

Best injection timing	`soi`
Best fuel quantity	`basefuelmass`
Best fuel pressure	`fuelpress`

Best VGT	`grackmea`
Best EGR	`egrpos, egrmf`

Minimize mode-weighted brake specific fuel consumption, subject to constraints on

- Turbo speed (`vtgrpm`)
- Cylinder pressure (`pkpress`)
- Exhaust equivalence ratio (`exheqr`)

To solve this problem, you must first use the Model Browser part of the Model-Based Calibration Toolbox product to design an experiment for collecting data, and then create models based on that data. You will use the resulting models in the CAGE Browser part of the toolbox to produce optimal calibration tables.

Design of Experiment

In this section...

"Introducing Design of Experiment" on page 3-4
"Constraining the Design" on page 3-7
"Creating Candidate Designs" on page 3-14
"Data Collection" on page 3-15

Introducing Design of Experiment

Creating a design in the Model-Based Calibration toolbox comprises several steps. First, you need to enter the ranges and names of the variables being used and choose a default model. Then you can create an initial design and set up the constraints on the space. These constraints will be the same for all designs. From this constrained design, you can create a series of child designs adding varying numbers of points and using different construction techniques. You can choose the final design by comparing the statistics of the various child designs, while considering how many test points you can afford to run.

Variables are

`measrpm`	Engine speed (rpm)
`basefuelmass`	Fuel quantity per injection (mg)
`fuelpress`	Fuel pressure (MPa)
`grackmea`	VGT rack position (%)
`egrlft`	EGR valve position (mm)
`soi`	Start of injection (deg ATDC)

You need to set up a test plan before you can make designs. This experiment is set up as a two-stage test plan with start-of-injection (SOI) sweeps at the local level and the other five variables at the global level.

Open the example session with the test plan set up as follows:

1 Start the Model Browser by typing `mbcmodel` at the MATLAB command line.

2 Select **File > Open Project**. Locate the example session with the test plan set up, `Diesel_testplan.mat`, in the `mbctraining` directory and double click to load the project.

3 Click the Two-Stage test plan node in the model tree to see the test plan diagram.

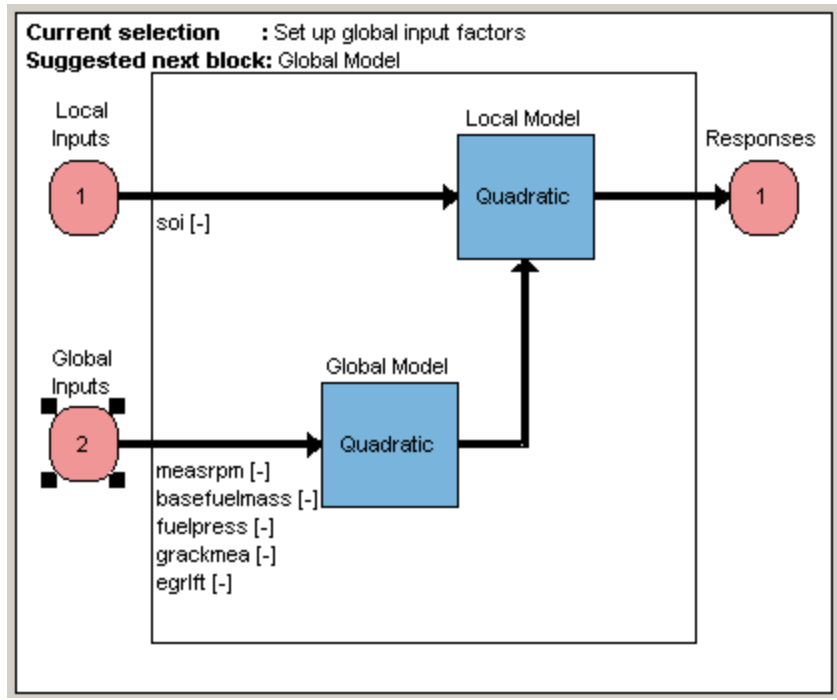

4 Double-click the Global Inputs block in the diagram to set the ranges of the inputs. You should set up the ranges before designing an experiment. You can enter the ranges in the min/max boxes to include the most extreme values you want to set for each variable. Check the ranges match those shown in the following example, then click **OK**.

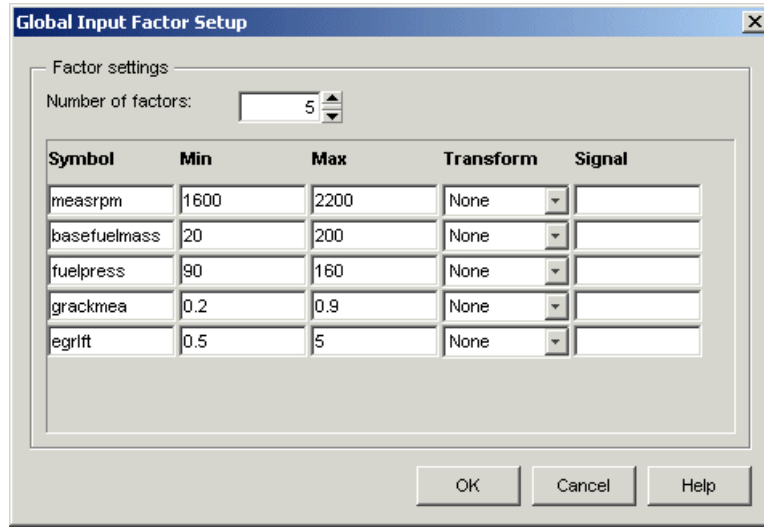

5 Double-click the Global Model block in the test plan diagram to view the model type. For this exercise, leave the model type at the default, which is a quadratic in all factors. Click **OK** to dismiss the dialog.

Remember that the statistical usefulness of different designs depends on the model type. For example, if you think you need cubic instead of quadratic in EGR, the number of points required rises dramatically and this has a highly adverse effect on the statistical quality of the designs.

Some possible models are

- Cubic polynomial, quadratic in fuel pressure: 41 terms
- Cubic polynomial, quadratic in fuel pressure and EGR: 31 terms

However, you do need to bear in mind that the final model will probably not be either of the possibilities listed here, because some terms will have been

removed, or it might even be an RBF (radial basis function). You choose the most suitable model you can in order to construct a design, then when you have collected the data you might find that a different model type produces the best fit.

Constraining the Design

These are the constraints you want to apply to the design space:

- `basefuelmass`
  - Maximum 200 at 1600 rpm, 175 at 2200 rpm
- `fuelpress`
  - Range 90 - 110 at 1600 rpm
  - Range 120 - 160 at 2200 rpm
- `grackmea`
  - Range 0.2 - 0.6 at 1600 rpm
  - Range 0.4 - 0.9 at 2200 rpm

The tables here are very simple: one output value defined at the min and max settings of RPM. The final constraint is a cube within the base fuel mass-fuel pressure-VGT space that moves and changes size as RPM is altered.

To add a constraint to a design,

1. First open the Design Editor by right-clicking the Global Model block in the test plan diagram and selecting **Design Experiment**.

2. Click the New Design button in the toolbar or select **File > New Design**. A new node called `Linear Model Design` appears.

   The new `Linear Model Design` node is automatically selected. An empty Design Table appears because you have not yet chosen a design, unless you have previous design views from other sessions. The Design Editor retains memory of view settings.

3. Select **Edit > Constraints** from the Design Editor menus.

4 The Constraints Manager dialog appears. Click **Add**.

5 The Edit Constraint dialog with available constraints appears. Make sure the default `1D Table` is selected from the **Constraint Type** drop-down menu. These are easier to set up than linear constraints, although working out the linear constraint numbers might be worthwhile for larger problems as it is faster.

6 You can select the appropriate factors to use. For the first constraint, choose `measrpm`, `basefuelmass`, and the inequality `<=` from the menus.

You can define the constraint by typing values in the edit boxes or by moving the large dots (clicking and dragging them) to define a boundary. For this constraint, you want to define two points.

7 Select the **Table Editor** tab and edit the **Number of breakpoints** to `2`, and click **Span Factor Range**.

8 On the **Graphical Editor** tab, click Move Points , then click and drag the right point (where measrpm =2200) down to basefuelmass =175. You can also enter the values in the **measrpm** and **basefuelmass** edit boxes, or in the table on the **Table Editor** tab.

9 Click to select the left point. Make sure the values are `1600` in the **measrpm** edit box and `200` in the **basefuelmass** edit box. The dialog should look like the following example.

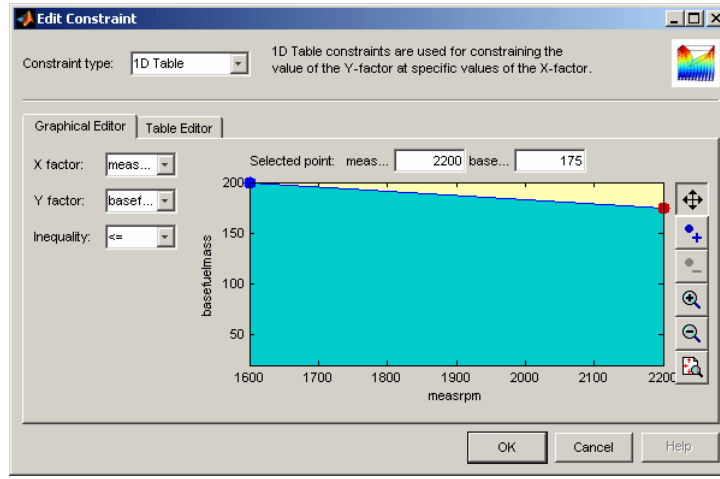

This constraint defines the range of `basefuelmass` in terms of RPM to within these bounds: maximum 200 at 1600 rpm, 175 at 2200 rpm.

10 Click **OK** to return to the Constraints Manager.

11 In the Constraints Manager, click **Duplicate** four times. This saves you setting up tables with only two points for the next constraints. Click to select the first new constraint, then click **Edit**.

You need to add constraints that define each of the following:

`fuelpress`

- Range 90 - 110 at 1600 rpm
- Range 120 - 160 at 2200 rpm

You achieve this by defining two constraints. In the first, the two table points define a `fuelpress` minimum of 90 at 1600 rpm and a minimum of 120 at 2200 rpm. In the second, the two table points define a `fuelpress` maximum of 110 at 1600 rpm and a maximum of 160 at 2200 rpm.

12 In the Edit Constraint dialog, change the Y factor to `fuelpress` and leave the X factor as `measrpm`.

13 Change the Inequality to >=.

14 Select the left point (where measrpm = 1600) and enter `90` in the **fuelpress** edit box.

15 Select the right point (where measrpm = 2200) and enter `120` in the **fuelpress** edit box. The dialog should look like the following.

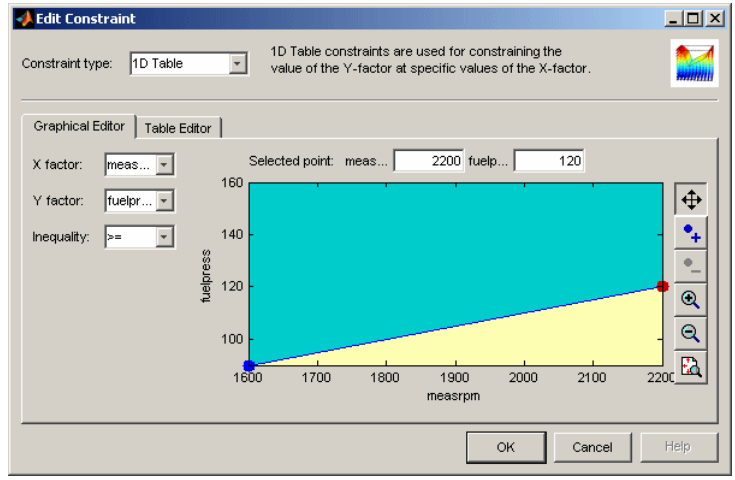

Click **OK** to return to the Constraint Manager.

16 Select the next constraint and click **Edit**. Edit the constraint to define a `fuelpress` maximum of 110 at 1600 rpm and a maximum of 160 at 2200 rpm, as shown.

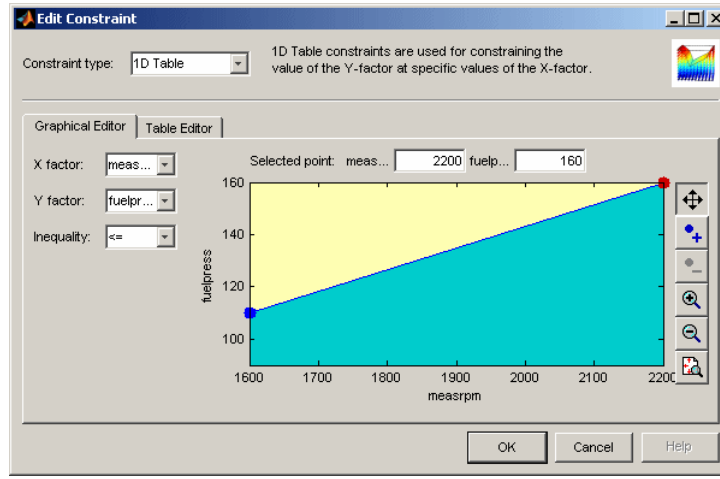

Click **OK** to return to the Constraint Manager.

17 Complete the other constraints in a similar way.

`grackmea`

- Range 0.2 - 0.6 at 1600 rpm
- Range 0.4 - 0.9 at 2200 rpm

This is achieved as shown in the following two constraints.

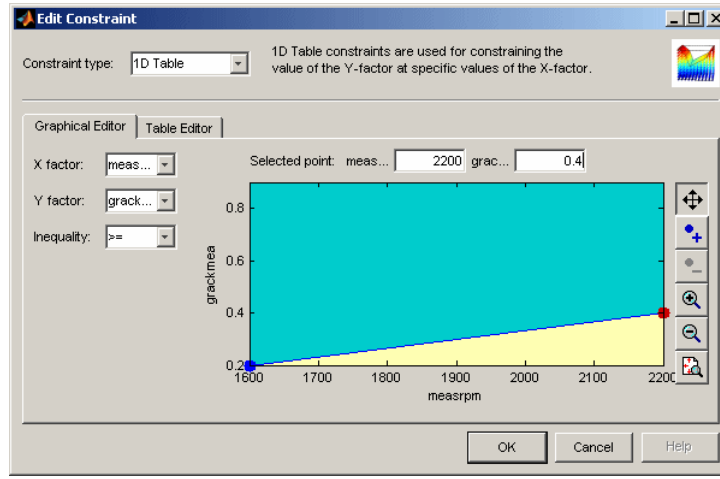

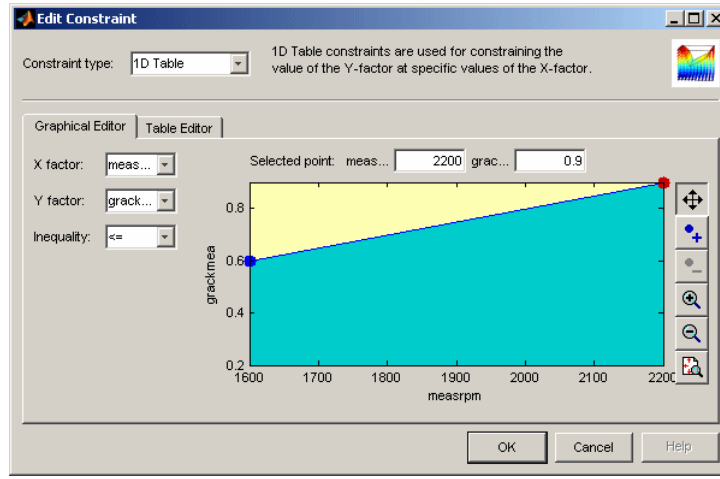

18 The Constraints Manager should contain all five constraints as shown.

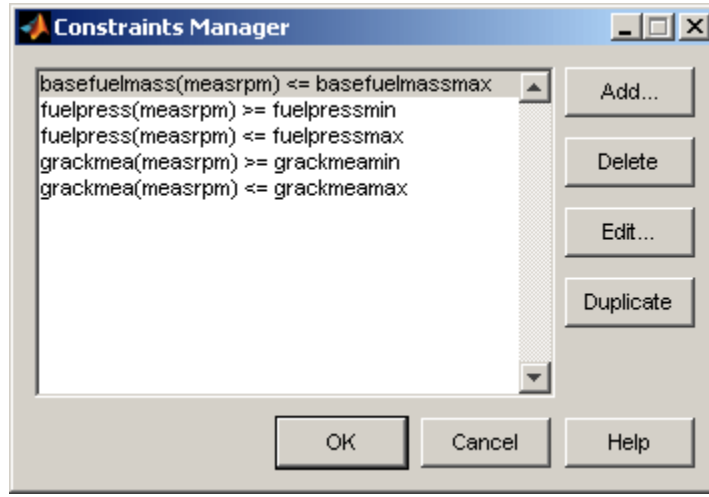

Click **OK** to return to the Design Editor.

19 Right-click a Design Editor view and select **Current View > 3D Constraints** to view the constrained design space.

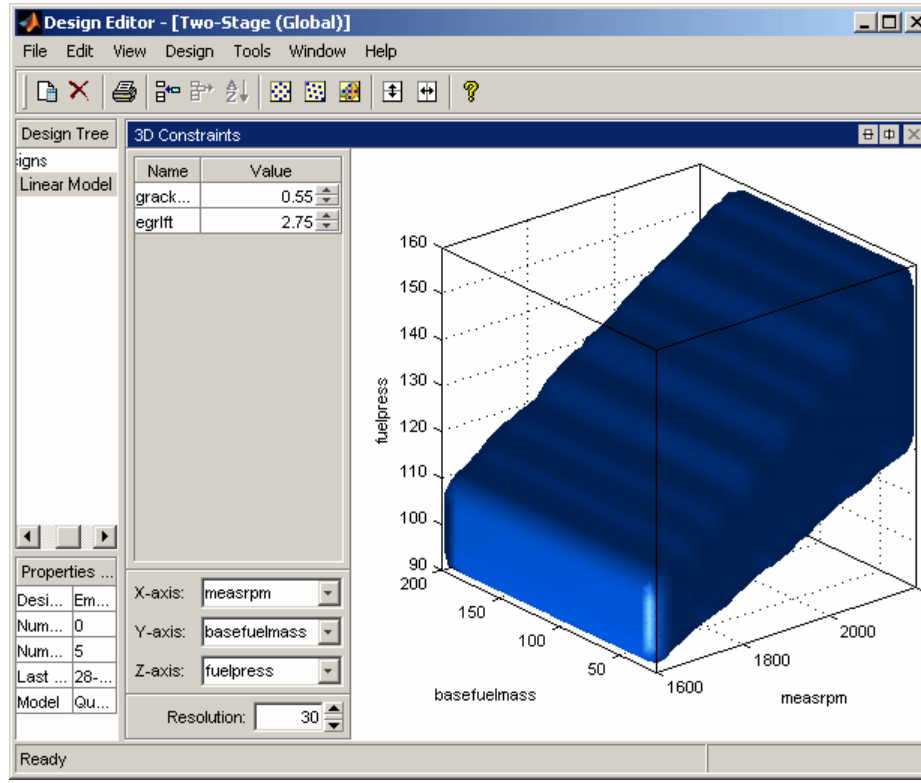

Creating Candidate Designs

Number of points

- How many do you have time for? When you consider the number of points, you need to remember that a sweep will be done at each point, and this will take some time.
- Do you need to allow time to fix problems or redo experimental points that can't be achieved due to combustion stability constraints?

Design type

- V-optimal: reduces average prediction error

V-optimal designs are often the preferred choice for engine testing. Optimal designs tend to push points to the edge, so they should give good coverage of the 1600 and 2200 RPM points while also allowing good modeling of the entire experimental region.

Create an optimal design with 65 points to compare to the example design.

1. Click Optimal Design in the toolbar. The Optimal Design dialog opens.
2. Enter `65` in the **Total number of points** edit box.
3. Select `V-Optimal` from the **Optimality criteria** drop-down menu and click **OK**.
4. The Optimizing Design dialog appears. Click **Accept** when iterations are not producing noticeable improvements; that is, the graph becomes very flat.
5. Examine the design points and compare to the constraint space by right-clicking the 3D Constraints view and selecting **Split View > 3D Design Projection**.

The final design used contained 65 points, for a quadratic in fuel pressure and EGR lift. V-optimal value = 0.302.

Data Collection

Data was generated by a Ricardo WAVE model using the experimental design. MATLAB and Simulink simulation tools control WAVE. Simulation tools support multiple WAVE processes retrieving test points from a central store. Average simulation time was 8 points (30 engine cycles each) per hour using four processors in parallel. Transient test results were then processed to extract steady-state results.

You can use the toolbox to import test data, view it, sort it into tests, verify ranges, filter out unwanted points, and select data for modeling. For details on any of these processes, see the examples in the gasoline case study section “Importing and Filtering Data” on page 2-18, and for comprehensive information on data handling in the toolbox, see “Data” in the Model Browser

documentation. See also the examples in the Chapter 7, “Tutorial: Data Editor”.

The example project provided (`Diesel_testplan.mat`) contains the filtered data attached to the test plan.

Modeling

In this section...
"Overview of Modeling Process" on page 3-17
"Building Models" on page 3-18
"Building and Evaluating Alternative Models" on page 3-20
"Creating Boundary Models" on page 3-26
"Export to CAGE" on page 3-29

Overview of Modeling Process

Outline of modeling steps you will cover in this example:

1. Build the models specified in the test plan.
2. Review local fits. Is the local model flexible enough?
3. Eliminate outliers.
4. Try alternative local models.
5. Review global fits. Is the global model flexible enough?
6. Eliminate outliers.
7. Try alternative global models.
8. Select best models.
9. Clean up tree.
10. Create boundary models.
11. Export models to file or CAGE.

See "Introduction to Two-Stage Modeling" on page 2-6 and "How Is a Two-Stage Model Constructed?" on page 2-27 for information about two-stage models (in the Help Browser you can right-click and select **Back** to return to this page).

Building Models

The file `Diesel_testplan.mat` contains a test plan where the model inputs and types are set up, and the data is loaded. You have set up the input ranges before designing experiments. Now you can build models of the responses.

1. Double-click the Responses block to select data signals for modeling and build response models.

2. The Data Wizard appears. You want to use all selected data, so click **Next** on the Select Data screen.

3. Select the check box to **Copy range** of the signals, then match up model input names with data signal names one pair at a time and click the button to associate each pair. As the input signal names are set up to match the data signal names, each model input factor you select is matched automatically in the data signals list, so you can then click the button to select the data signal and copy the range. If the names did not match you would have to select data signals manually.

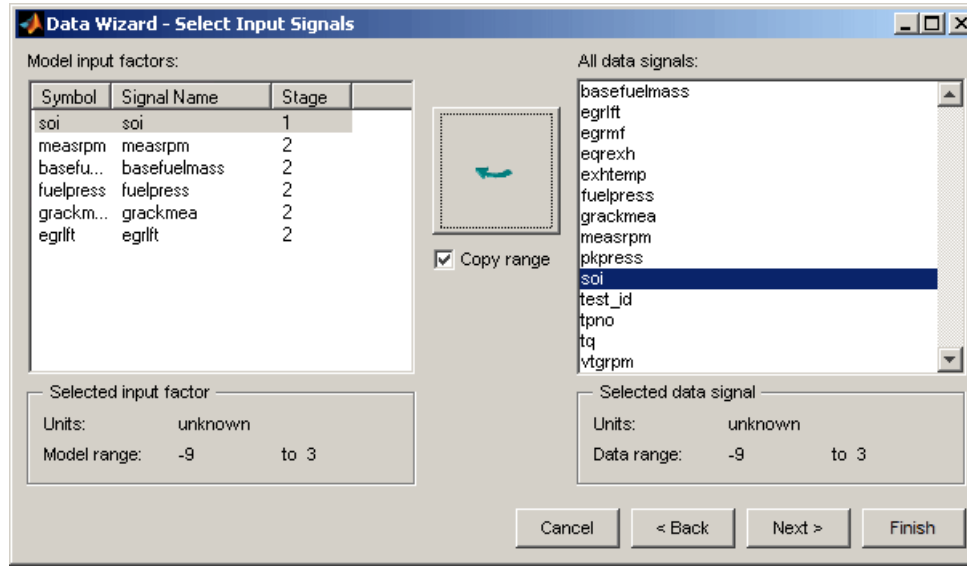

Click **Next**.

4 You can build one response at a time or set up several at once on the **Select Responses** screen of the Data Wizard. Click **Add** after selecting each of these responses:

- `egrmf`
- `eqrexh`
- `pkpress`
- `tq`
- `vtgrpm`

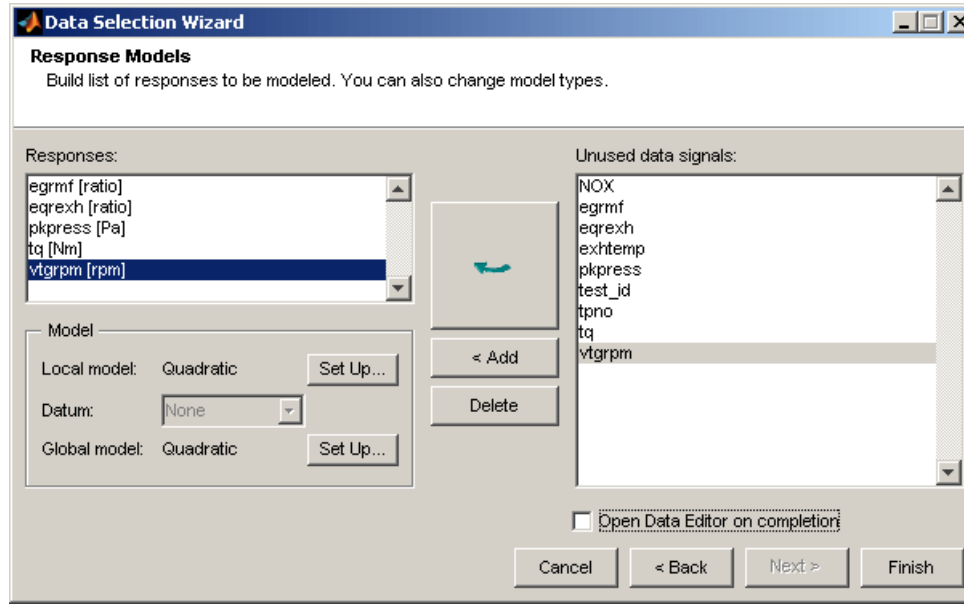

5 Clear the check box **Open Data Editor on completion**. If you left this check box selected you could choose to review the data selected for modeling in the Data Editor. Click **Finish**.

The models appear in the model tree in the All Models pane of the Model Browser.

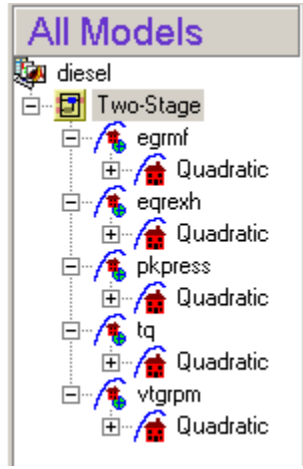

Building and Evaluating Alternative Models

1 Review the local fits. Start by selecting the local model node (`Quadratic`) under the `tq` response.

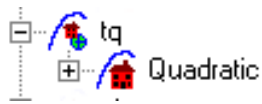

The local model view appears.

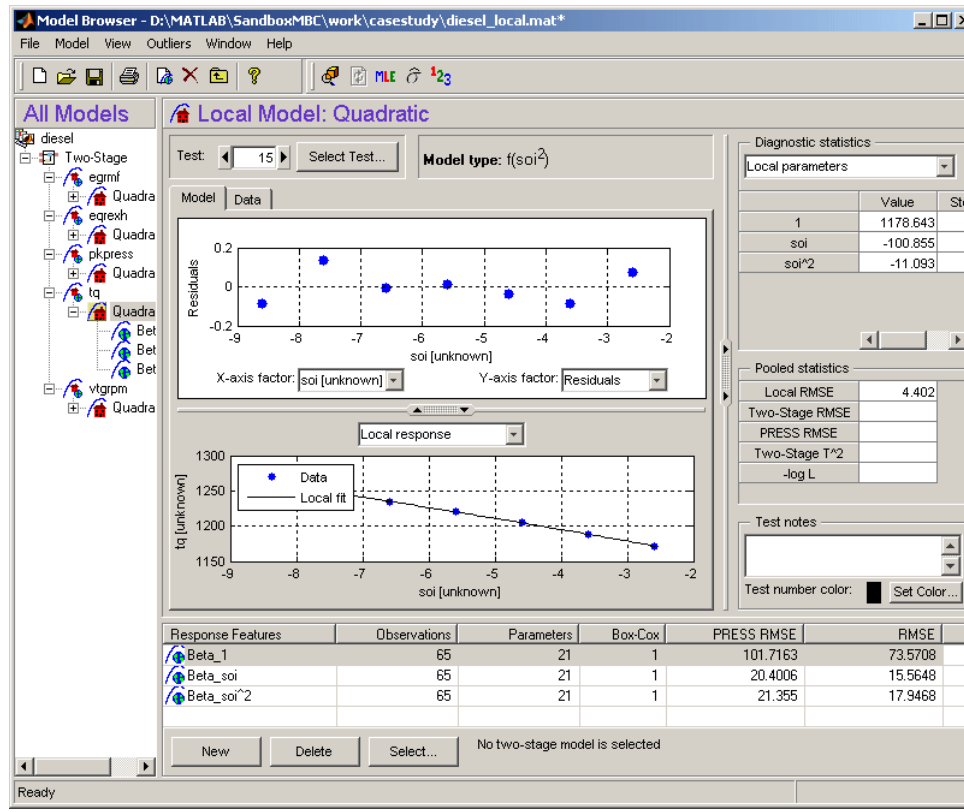

2 Look through the tests to inspect the fits. Use the **Test** controls.

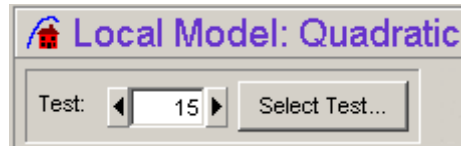

Look for trends in the residuals in the upper plot. Evenly distributed residuals are a good sign; trends in the residuals can indicate that the model is not flexible enough and you should try more flexible model types. A trend in magnitude of residuals indicates that transform or covariance modeling might be needed.

3 To quickly identify problem tests with high errors, click RMSE Plots in the toolbar (or **View** menu). Double-click tests with high error values in the RMSE Explorer to inspect them in the local model view. Strongly outlying residuals should be investigated.

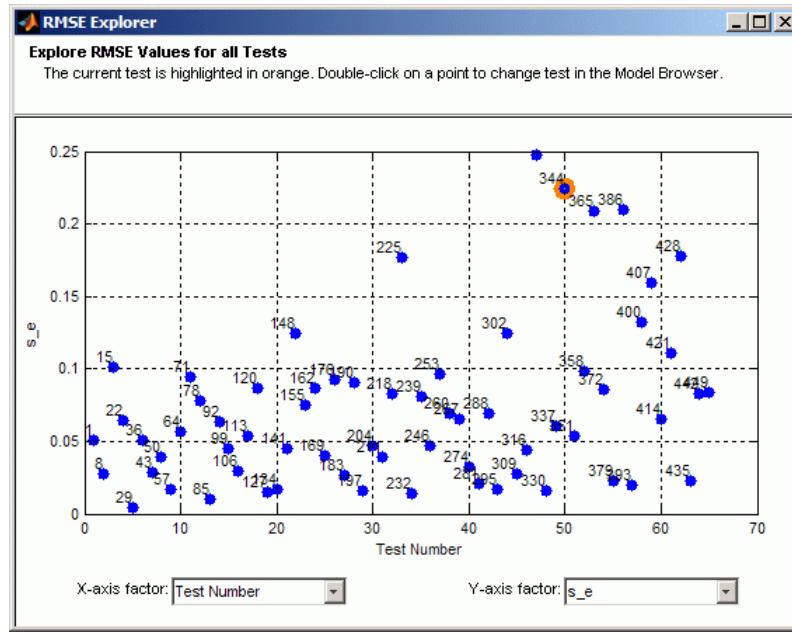

4 Consider removing outliers to improve fits if some points are badly distorting the fit. Engineering judgment is required to judge whether suspect data should be removed. Be careful not to remove outliers without good reason. If you keep removing points you can always get a better fit, but your aim is to achieve a model that predicts the data well. You can click the **Data** tab to inspect data variables for suspect tests.

5 Use the same principles to review the global fits. For example, expand the local model node (`Quadratic`) in the model tree and click `Beta_1`.

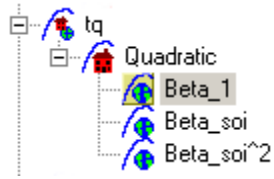

You can right-click outliers (or any point) in the plots to see a plot of the test and inspect the variables. For information on how the local and global models relate to each other, see "How Is a Two-Stage Model Constructed?" on page 2-27

6 At both local and global level, create alternative model types to compare.

a From the response node, create alternative local models to compare.

b From the local nodes, create alternative response feature models to compare.

c From the global nodes, create alternative child global models to compare.

To create models one at a time, click **New** from the response node to create a new local model node, or from the global node to create new child global model nodes, etc. Choose different model types from the Model Setup dialog, or you can change any existing model by selecting **Model > Set Up**. There are step-by-step instructions for doing this in the gasoline case study, see "Create Multiple Models to Compare" on page 2-51. You could also work through the modeling tutorial for more guided examples of how to select models, see Chapter 5, "Tutorial: Model Quickstart ".

When you have built a selection of different model types as child nodes of, say, a global model node, you can click the parent model node and select **Model > Make Template**. Save the template to a suitable directory, then you can use **Build Models** (in the toolbar) to automatically build the same selection of child model types for any other model.

Remember that from any parent node, you can see a list of statistical comparisons for all the child nodes in the lower list pane, along with information such as the number of parameters.

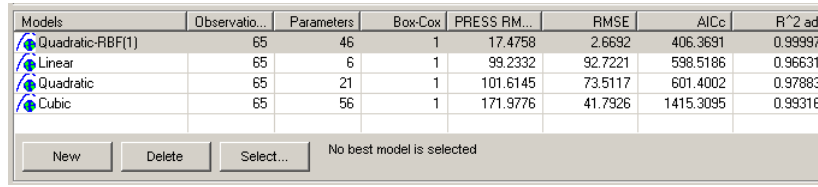

Review the fits graphically as well. Search for the best fit but be careful not to overfit. You can increase the number of terms in a model until every point is exactly on the line, but the predictive power of that model will be very low. PRESS RMSE can be the most helpful single statistic you can use to search for the best fit relative to the number of terms in the model. However, you should not rely on any single statistic, but use a variety of criteria and especially the graphical tools available for comparison of models in the Model Evaluation tool when you click **Select**. For detailed guidance see in the Model Browser documentation.

Make use of the Stepwise tool (in the toolbar for linear models) to automatically search for a good fit with the minimum number of useful model terms. You can set Stepwise to run automatically when you create models (for example, select `Min PRESS` from the **Stepwise** drop-down menu in the Model Setup dialog) or you can open the Stepwise window after the model is built. Remember that modeling is a tradeoff — too few parameters means the shape of the surface cannot be captured, while too many parameters gives a risk of overfitting.

7 Create two-stage models by clicking **Select** at the Local node. You must first select a best model for each response feature (global) model, if you have created alternatives. If you have created alternative response feature models, when you click **Select** you can choose the best combination of response features in the Model Selection window. There are many graphical tools available here, such as surface plots, contour plots, and movies.

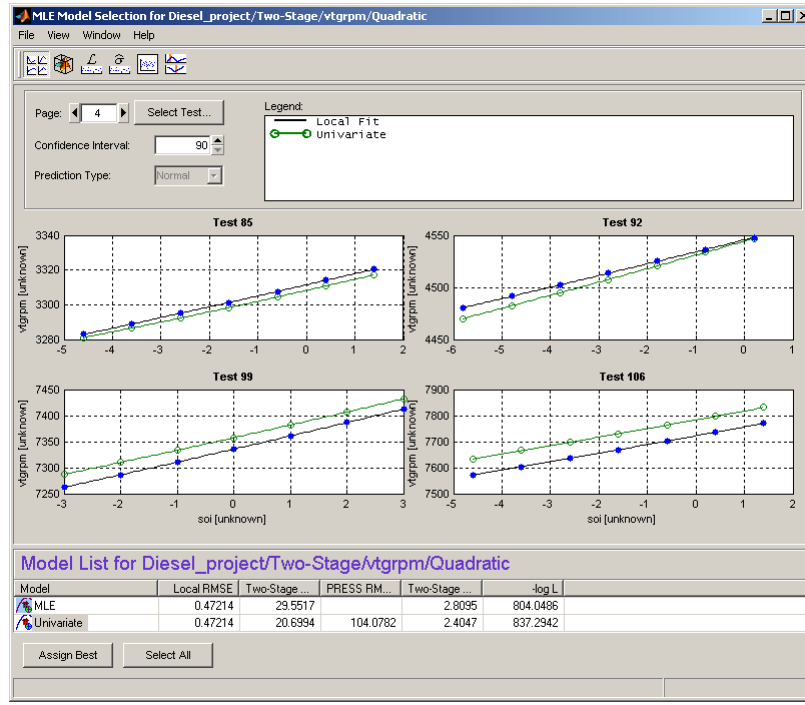

8 Try MLE (Maximum Likelihood Estimation). You can choose to calculate MLE in the dialog that appears immediately after building a two-stage model, or you can click **Cancel** and choose to try MLE later. You can use the MLE toolbar button to calculate an MLE model any time. This process refits, taking proper account of the correlation between different response features. Once you have an MLE model, you can click **Select** from the local node to compare it with the univariate two-stage model. You can choose the univariate model as best here if you want; this is the way to "go back" from MLE.

9 Searching for a good fit often results in large numbers of models. To discard all but the models you chose as best, select **File > Clean-up Tree**.

For guidance, look at the models in the example finished project, `Diesel_project.mat`, found in the `mbctraining` directory.

Creating Boundary Models

You can create a boundary model at the test plan node. A model describing the limits of the operating envelope can be useful when you are evaluating optimization results.

1. Select the test plan node in the model tree.
2. Select **TestPlan > Boundary Constraints**.

   The Boundary Constraint Editor opens.
3. Click Make Boundary Constraint in the toolbar. A dialog box opens where you can choose to build a boundary model of the response, local, or global inputs. First, make a model of the response boundary. Select **Response**, and click **OK**.

   A dialog box opens where you can select constraint inputs.
4. Set the **Constraint type** to `Convex Hull`. Leave only the `soi` and `measrpm` check boxes selected, as shown in the following figure.

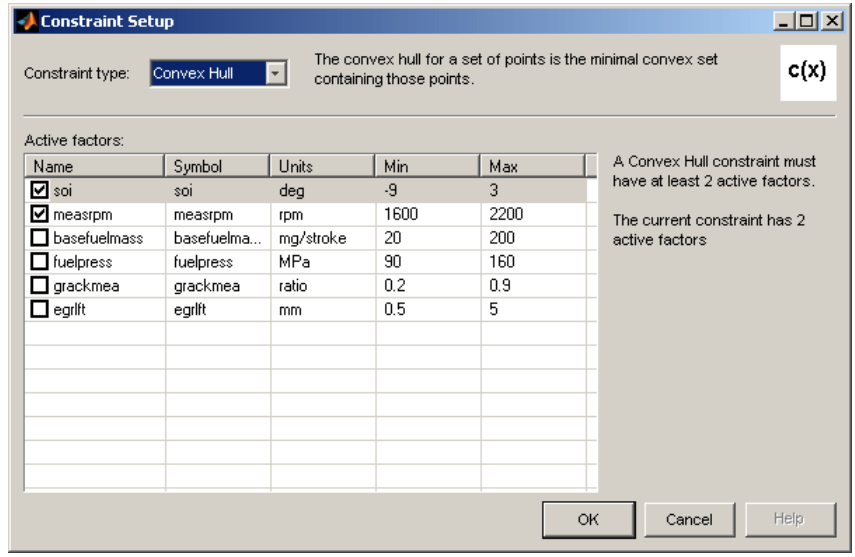

Click **OK**, and then click **OK** again in the next dialog box because there are no other parameters to set for this constraint.

A `Convex Hull(soi,measrpm)` child node appears under the `Response` node.

5 Select **View > Current View > Pairwise Projections** (or use the toolbar button). You can see pairwise projections to view the boundary across all combinations of factors.

6 Next, make a model of the global boundary. There is already a `Global` node with an empty child node under it. Select this node, `G_Boundary`, then select **Constraint > Set Up**.

A dialog box opens where you can select constraint inputs.

7 Set the **Constraint type** to `Convex Hull`. Leave only the `measrpm` and `grackmea` check boxes selected. Click **OK**.

The name changes to `Convex Hull(measrpm, grackmea)`, but the constraint is not complete yet.

8 Select **Constraint > Full Constraint Fit**. Click **OK** in the next dialog box as there are no other parameters to set for this constraint. The constraint model appears in the pairwise view. The tree should look as shown in the following figure.

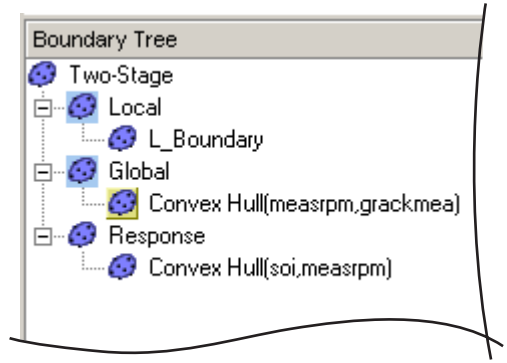

9 Click **Duplicate constraint** in the toolbar twice. Two new convex hull child nodes appear under the `Global` node.

10 Select the first new node, and select **Constraint > Set Up**. The Constraint Setup dialog box appears.

   a Leave only the `measrpm` and `fuelpress` check boxes selected. Click **OK**.

   b Select **Constraint > Full Constraint Fit**. Click **OK** in the next dialog box as there are no other parameters to set for this constraint.

11 Select the second new node, and select **Constraint > Set Up**. The Constraint Setup dialog box appears.

   a Leave only the `measrpm` and `basefuelmass` check boxes selected. Click **OK**.

   b Select **Constraint > Full Constraint Fit**. Click **OK** in the next dialog box as there are no other parameters to set for this constraint.

12 You can select a combination of models as best. Select each of the `Global` and `Response` models in turn, and select **Constraint > Add to Best** (also in the toolbar). Make sure both the Response and Global nodes are also added to best.

   You can delete the `Local` node as you are not interested in the local boundary.

   Make sure all remaining icons are blue, indicating they are included in the combination of best models, as shown in the following figure. For more information see .

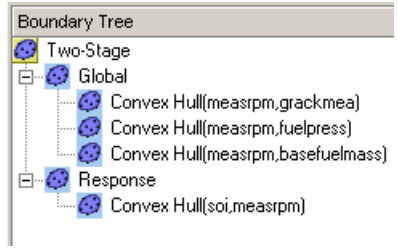

13 Now, select the root node `Two-Stage` to see the combination of all constraints in the pairwise view, as shown in the following figure.

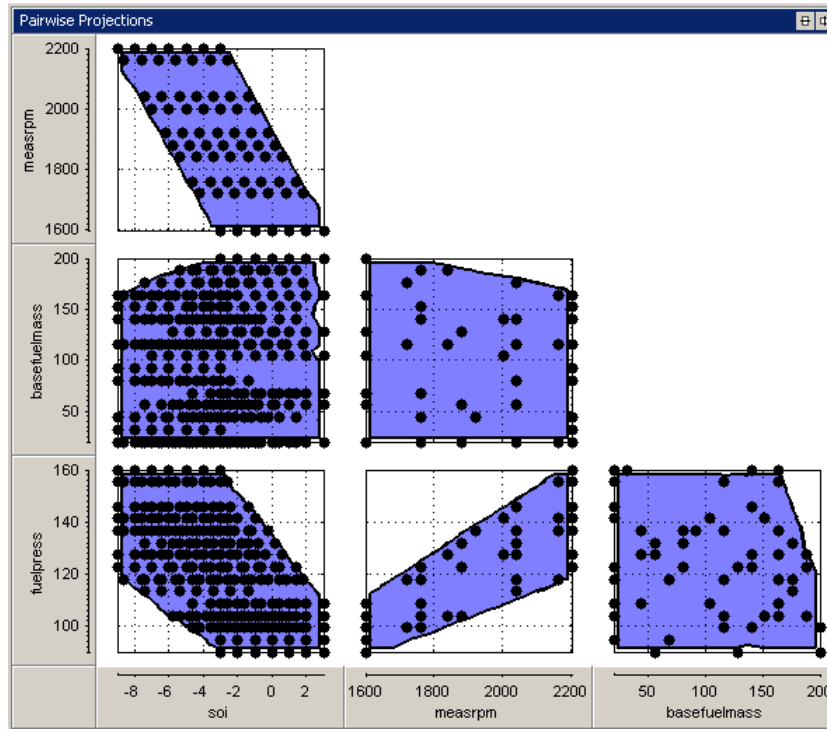

14 Close the Boundary Constraint Editor to return to the Model Browser. After they are calculated, the boundary constraints remain part of the test plan unless you delete them.

Export to CAGE

You can export your models to file, to CAGE, to the workspace or to Simulink. This capability makes models easy to share across engineering groups. You can do such exports in the Model Browser by selecting **File > Export**. The models exported depend on the model node you have selected in the tree. To export all models in the test plan, select the test plan node.

Note For this example, you use CAGE to import the models.

You use the example models in the CAGE part of the Model-Based Calibration Toolbox product to produce optimized calibration tables. Proceed to the next section for instructions.

Optimized Calibration

In this section...

"Problem Definition" on page 3-31

"Benefits of Automated Calibration" on page 3-32

Problem Definition

This section describes the creation and optimization of calibration tables for the diesel case study. You use the models created in the Model Browser section of this case study. An example file is provided.

The problem to solve is to produce tables in speed and torque for

Best injection timing	`soi`
Best fuel quantity	`basefuelmass`
Best fuel pressure	`fuelpress`
Best VTG	`grackmea`
Best EGR	`egrlft`

Also, to minimize brake specific fuel consumption, subject to constraints on

- Turbo speed (`vtgrpm`)
- Cylinder pressure (`pkpress`)
- Exhaust equivalence ratio (`exheqr`)

For this case study, you use models produced in the Model Browser to generate calibrations in CAGE. You cover the following steps:

1. Load models of engine responses. See "Importing Models of Engine Responses into CAGE" on page 3-34.
2. Define additional models and variables required by optimization strategy. See "Defining Additional Variables and Models" on page 3-35.

3 Set up tables. See “Setting Up Calibration Tables to Fill” on page 3-37.

4 Set up the optimization to provide feasible starting points for the sum problem. See “Setting Up the Optimization” on page 3-38.

5 Define constraints. See “Setting Up Constraints” on page 3-42.

6 Define optimization operating points. See “Defining Variable Values” on page 3-47.

7 Run the optimization and view results. See “Running the Optimization” on page 3-50

8 Duplicate and modify the optimization to create a sum optimization using the previous solutions as starting points. See “Setting Up the Sum Optimization” on page 3-52

9 Fill tables from optimization results. See “Filling Tables with Optimization Results” on page 3-59.

For guidance, you can look at the example finished project, `Diesel_optimization.cag`.

Benefits of Automated Calibration

- You can move the table-filling process away from the test bed.
- You can regenerate calibrations when objectives, constraints, or calibration table layouts change, without additional testing.
- You can explore tradeoff possibilities interactively.
- You can produce initial calibrations using engine simulation software, before hardware is available.

CAGE can provide both automatic and interactive calibration optimization. You can trade off multiple objectives, deal with multiple constraints, and you can examine optimizations point-by-point or over a drive-cycle. CAGE can provide solutions for these example applications:

- Example control applications

  Emissions-constrained BSFC optimization over drice cycles producing calibrations such as:

  - Optimal fuel injection timing schedule
  - Optimal fuel injection quantity schedule
  - Optimal EGR valve position and EGR mass fraction schedule
  - Optimal spark timing schedule
  - Optimal dual-independent variable valve timing schedules
- Estimation problems
  - Torque
  - Emissions
  - Air flow and manifold pressure
  - Intake valve temperature
  - Borderline spark

Importing Models of Engine Responses into CAGE

1. Start CAGE by typing `cage` at the MATLAB command line.
2. Select **File > Import > From Project**. The CAGE Import Tool appears.
3. Select **File > Select Project File** or click the toolbar button. Locate the example model file created in the Model Browser, `Diesel_project.mat`, in the `mbctraining` directory, and click **Open**.
4. You can select from a list of models in the file. Select `Response` from the **Type** drop-down list to filter the model list.
5. Select all the response models in the list by **Shift**+clicking.
6. Click the toolbar button **Import selected items**.
7. In the following Import dialog box, click **OK** to import the models.
8. Close the Import Tool and click the Models button to select the Models view in CAGE. You should see the models in the list. The selected model is displayed in the other panes.

You can also:

- Export models directly from the Model Browser to CAGE when both are open.
- Use the CAGE Import Tool to import models directly from the Model Browser.
- Use the CAGE Import Tool to import models and other calibration items (tables, optimizations, tradeoffs, data sets, features) from any project file created in CAGE or the Model Browser. This capability can help you use existing projects to speed up the setup of new sessions.
- Export models to a file, which you can import to CAGE by selecting **File > Import > Model**.

Defining Additional Variables and Models

Looking at the problem definition, you need to define some additional variables and models for this optimization. You want to minimize BSFC, but this is not a model output. You need to fill tables against speed and torque, but torque is not an input for the models. The solution to this is to create function models. You can then use a torque constraint and the BSFC function model in the optimization. To implement the torque constraint, you need to define a new variable for desired torque, `tq_desired`.

Follow these steps:

1. Select **File > New > Variable Item > Variable**. The Variable Dictionary view appears.
   a. Rename the new variable `tq_desired`.
   b. Set the range of this variable to be minimum `0` and maximum `1500`.
   c. Edit the set point to `600` Nm
2. Select **File > New > Function Model**.
   a. Enter `bsfc = 5400 / pi * basefuelmass / tq` and click **Next**.

      Assuming: base fuel mass [mg/inj], Tq [Nm], 6 cylinders, 4 stroke
   b. Select **Automatically assign/create inputs**. It is important to check for typos or this step can create unintended new inputs. Click **Finish**.

      The Models view appears.
3. This optimization requires a constraint on the air/fuel ratio (AFR). However, as you could not model AFR directly you need to create a function model that relates equivalence ratio to air/fuel ratio. Then you can constrain AFR. To do this, select **File > New > Function Model**.
   a. Enter `afr = 14.46/eqrexh` and click **Next**.
   b. Select **Automatically assign/create inputs** and click **Finish**.
4. Select **File > New > Variable Item > Variable**. The Variable Dictionary view appears.

a. Rename the new variable `afr_min`.

b. Set the range of this variable to be minimum `0` and maximum `100`

c. Edit the set point to `50`

Setting Up Calibration Tables to Fill

1. Set up a new 2-D table. Select **File > New > 2D Table**.
2. Name the table `soi_best`.
3. Select `measrpm` and `tq_desired` from the **Y Input** and **X Input** drop-down menus.
4. Enter 11 rows and 11 columns.
5. Enter 0 for the initial value.

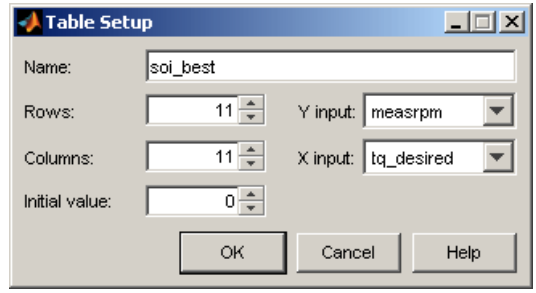

6. Click **OK** to create the table.

   Look at the **Tables** tree on the left. Note that the normalizers appear as calibratable items in their own right, and as descendants (child nodes) of their tables.

7. Once you create a table, you can duplicate it to create more tables that share the same breakpoints. Click `soi_best` in the tree, then select **Edit > Duplicate soi_best**. Repeat until you have four new tables, and rename them as follows:

   a. `basefuelmass_best`

   b. `fuelpress_best`

   c. `grackmea_best`

   d. `egrlft_best`

Setting Up the Optimization

CAGE provides a flexible optimization environment. You can define the objectives, the constraints, and the points where the optimization is carried out.

The objective is to minimize weighted BSFC over a set of points, subject to a set of constraints. This is a constrained single objective optimization problem.

You solve this problem in two parts:

1. Run an initial optimization to explore the problem point-by-point in order to obtain good starting points for the sum optimization.
2. Use the solutions from the initial optimization as the start points to run the sum optimization over the drive cycle to find optimal control settings. Use these results to fill calibration tables.

CAGE has several built-in optimization routines and the capacity for you to write your own. In this case study, you use `foptcon`. This algorithm is a modified version of `fmincon` from the Optimization Toolbox product. In CAGE, you can use the algorithm to minimize or maximize an objective function.

1. Select **File > New > Optimization**. The Optimization Wizard appears.

   `foptcon` is selected by default, and this is the optimization algorithm you will use for this example. This algorithm specifies a single objective in the **Objectives** column.

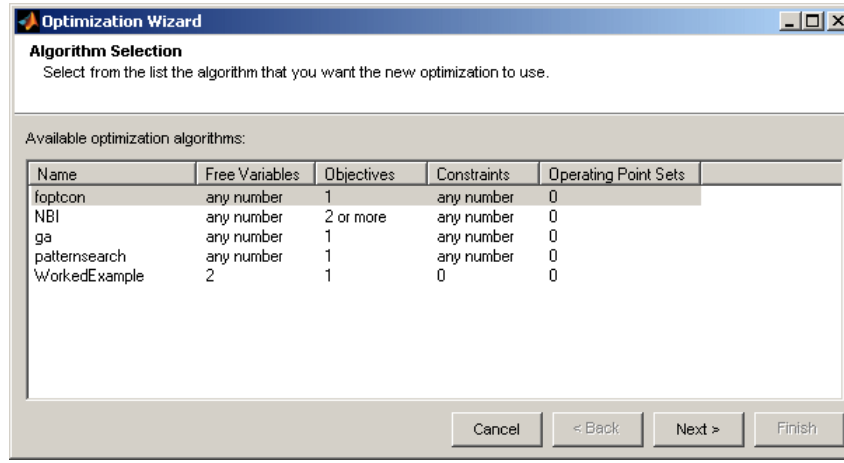

2 Click **Next**.

3 Here you set up your options as follows:

a Increase the number of free variables to 5

b Increase the number of constraints to 6

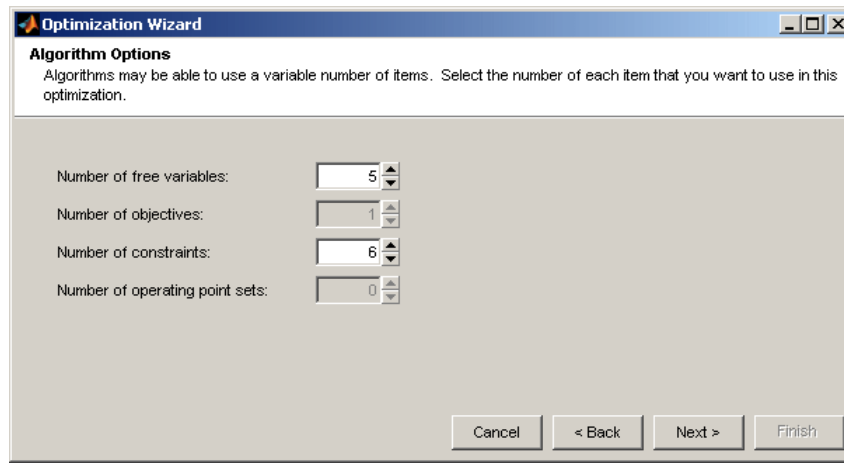

Click **Next**.

4 Select the free variables.

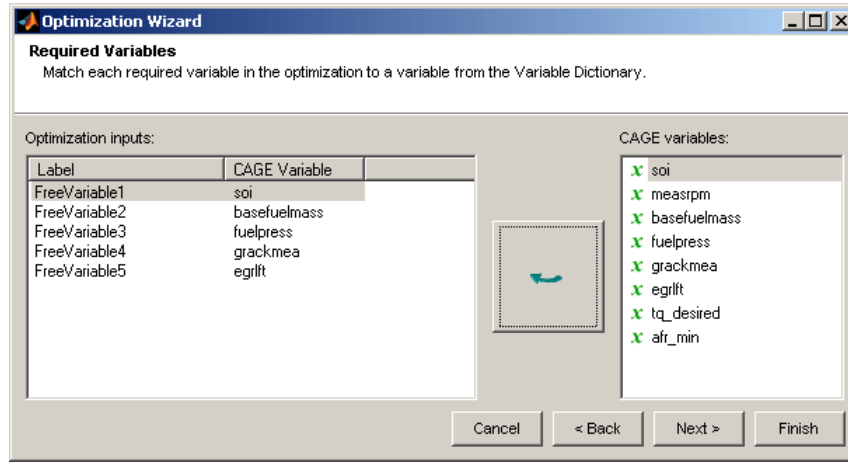

Select in turn each of the following free variables and click the button to select them:

a `soi`

b `basefuelmass`

c `fuelpress`

d `grackmea`

e `egrlft`

Click **Next**.

5 Select `bsfc` from the list of models on the right and click the button to select this model for the objective. Leave the objective type set to **Minimize**.

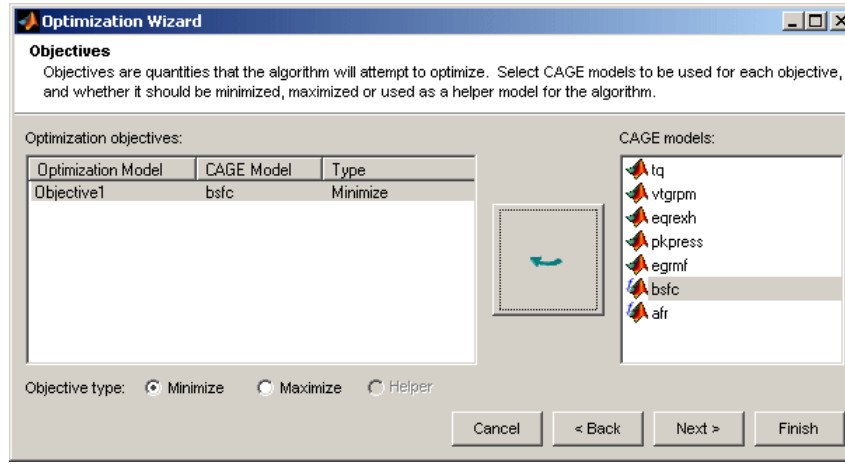

Click **Finish**.

The **Optimization** view appears. You have not yet set up your constraints, so you can add them here.

Setting Up Constraints

This case study problem has CAGE model constraints on the following quantities (which you will set up next):

- `tq_desired`
- `vtgrpm`
- `pkpress`
- `afr`
- `soi`
- `grackmea`
- `fuelpress`
- `basefuelmass`

In the Optimization view these constraints are shown as follows.

Constraints

Name	Description
tq_minus	tq(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft) <= tq_desired
tq_plus	tq(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft) >= tq_desired
afr	afr(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft) >= afr_min
vtgrpm	vtgrpm(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft) <= 128000
pkpress	pkpress(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft) <= 18000000
tq_Boundary	Boundary constraint of tq(soi, measrpm, basefuelmass, fuelpress, grackmea, egrlft)

You can edit constraint names to aid analysis in other Optimization views (right-click, or double-click to use the Edit Constraint dialog box).

Set up these constraints as follows:

1. Double-click to edit the first constraint to `tq <= tq_desired`. The Constraint Editor appears.
   a. Leave the Constraint type as `Model`.
   b. Edit the Constraint name to `tq_minus`.
   c. Select `tq` in the **Input model** list
   d. Select <= as the **Constraint type**

e Select the **CAGE item** radio button

f Select `Show variables` from the drop-down list.

g Select `tq_desired` in the variable list. The dialog should look as shown.

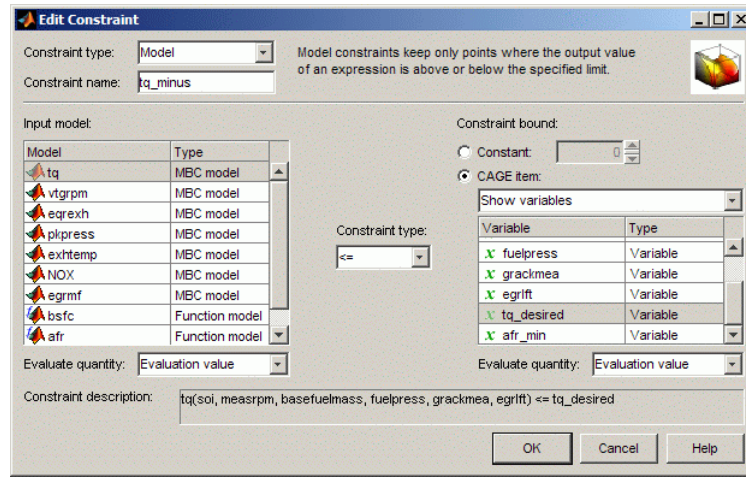

h Click **OK** to return to the Optimization view.

2 Similarly, double-click to edit the next constraint as shown: `tq >= tq_desired`.

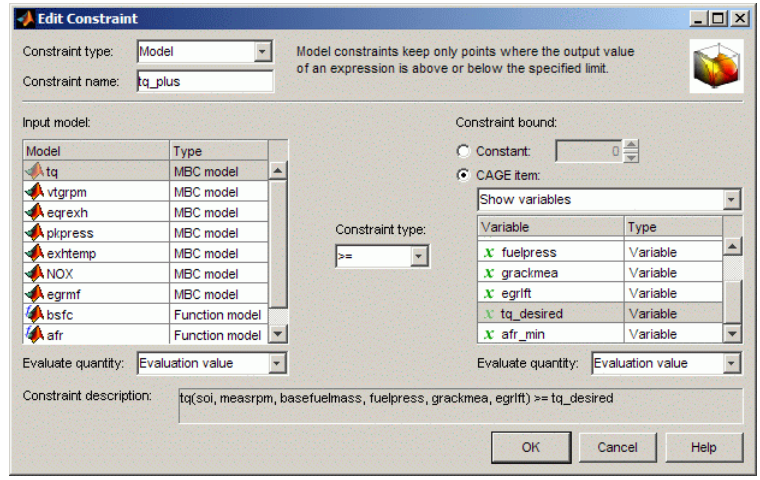

3 Repeat to set up the constraint `afr >= afr_min`, as shown.

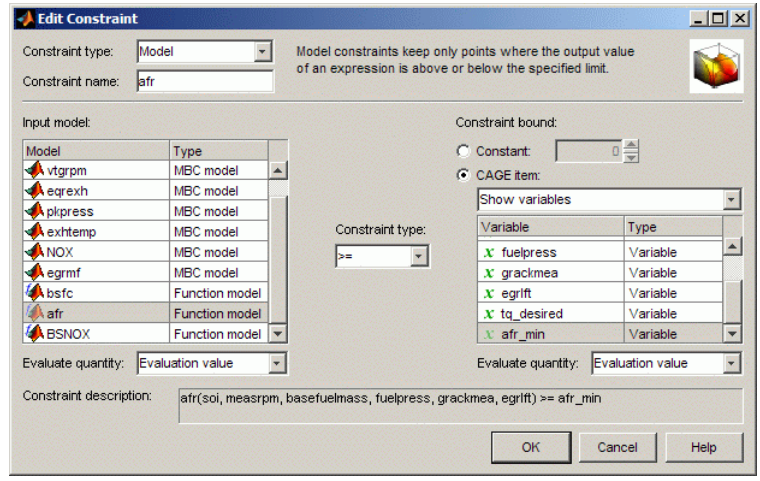

4 Repeat to set up the constraint `vtgrpm <= 128000`, as shown. Note, here you enter a value in the **Constant** edit box, rather than select a variable from the CAGE item list.

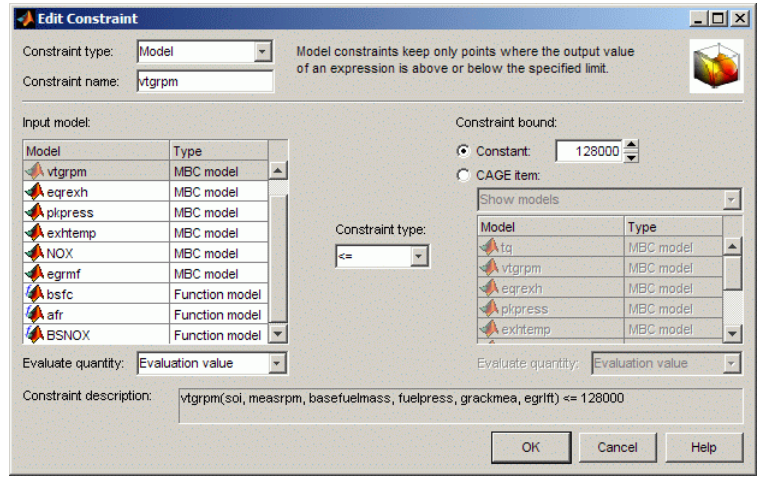

5 Repeat to set up the constraint `pkpress <= 18000000`, as shown. Note, here you enter a value in the **Constant** edit box.

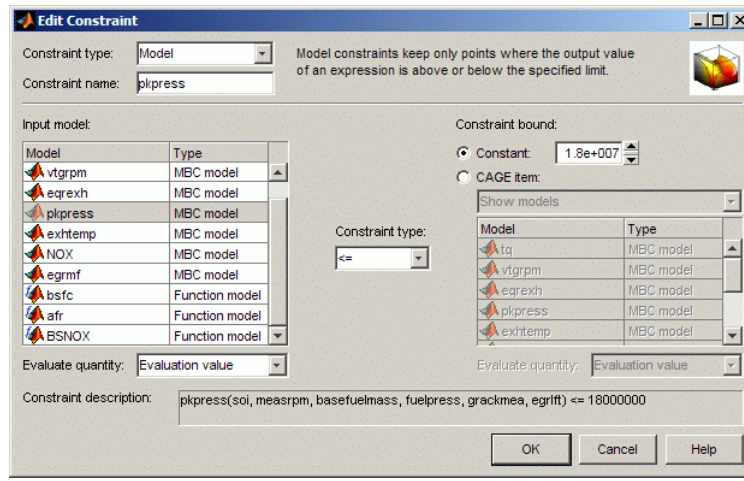

6 Repeat to set up a boundary model constraint, as shown in the following figure.

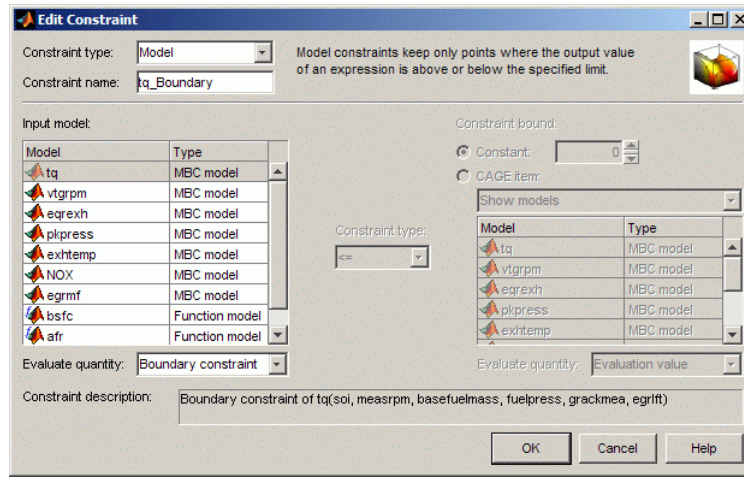

a Leave the Constraint type as `Model`.

b Edit the Constraint name to `tq_Boundary`.

c Select `tq` in the **Input model** list.

d Select `Boundary constraint` in the **Evaluate quantity** drop-down list.

e Click **OK** to return to the Optimization view.

Defining Variable Values

You need to define the set of drive cycle points where you want the optimization to run. To do this you use the **Input Variable Values** pane in the Optimization view.

1. Increase the **Number of runs** to 7. Click the buttons or enter the value in the box. The number of rows in the fixed and free variables panes increases to 7. The default values are the set point of each variable. Leave the initial values for the free variables at the defaults.

2. You can enter or copy values into the **Fixed Variables** pane to define the fixed variable values at each point where you want the optimization to run. You can copy all the variable values from a text file, or from the Help Browser copy each column in turn. Copy and paste the following values into the `measrpm` column in the **Fixed Variables** pane.

`measrpm`
2200
2200
2200
2200
1600
1600
1600

3. Copy and paste the following values into the `tq_desired` column in the **Fixed Variables** pane.

`tq_desired`
1263
947
632
126

1550
1163
775

4 Copy and paste the following values into the `afr_min` column in the **Fixed Variables** pane.

`afr_min`
25.5
27.75
30.0
0
22
22.5
23

The **Fixed Variables** pane should look as shown. You have specified the speed and torque values you want for each run, and also the minimum air/fuel ratio (AFR) at each N, TQ point.

Fixed Variables

Variable:	measrpm	tq_desired	afr_min
Number of values:	1	1	1
1	2200	1263	25.5
2	2200	947	27.75
3	2200	632	30
4	2200	126	0
5	1600	1550	22
6	1600	1163	22.5
7	1600	775	23

If you wished to constrain the free variables more than the range defined in the variable dictionary, you could select **Optimization > Edit Free Variable Ranges**. Leave the defaults for this problem.

Running the Optimization

Now you can run the optimization. You have set up objectives, constraints and a set of operating points.

1 Click Run Optimization in the toolbar.

2 When the optimization is complete, the view switches to the new child node, `Optimization_Output`, in the Optimization tree under the `Optimization` node. View the results.

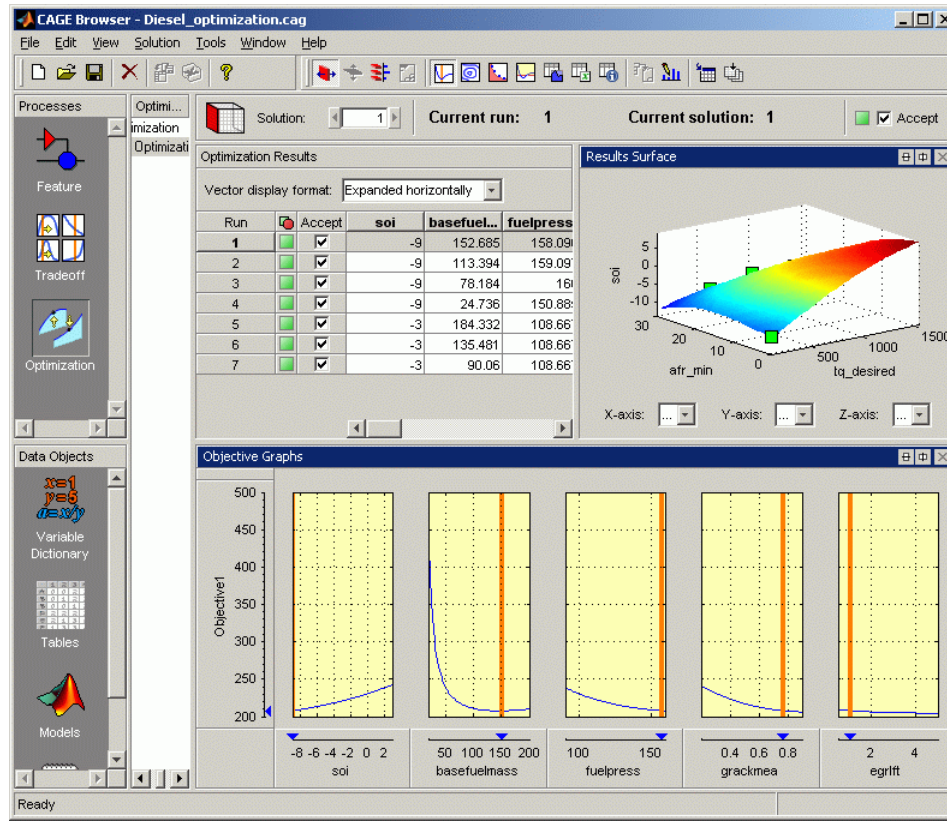

3 Look through the solutions at different operating points by clicking cells in the output table. Compare with the gasoline example, where regions that

do not meet the constraint are yellow in the graphs — in this case the torque inequality constraints cause all the graphs to be yellow, because these constraints produce a feasible contour in the free variable space, and any deviation off this contour is infeasible.

4 Split the view to display the constraint graphs, and scroll through them to see each constraint in relation to the solutions.

Setting Up the Sum Optimization

In this section...

"Setting Up Initial Values and Sum Objective" on page 3-52

"Creating the Brake Specific NOx Constraint" on page 3-54

"Setting Weights for the Sum Objective and Constraint" on page 3-55

"Set Parameters and Run Optimization" on page 3-57

Setting Up Initial Values and Sum Objective

The initial point-by-point optimization was designed to obtain good starting points for the sum optimization. Now, you can duplicate and modify that initial optimization to run the sum problem. You use the solutions from the initial optimization as the start points to run the sum optimization over the drive cycle to find optimal settings. You use these results to fill calibration tables.

1. Right-click the initial optimization node, and select **Duplicate**.
2. Click to select the new optimization, press **F2**, and enter `Sum_optimization` as the new name.
3. Import initial values from the output of the first optimization. Select **Optimization > Import From Output**.

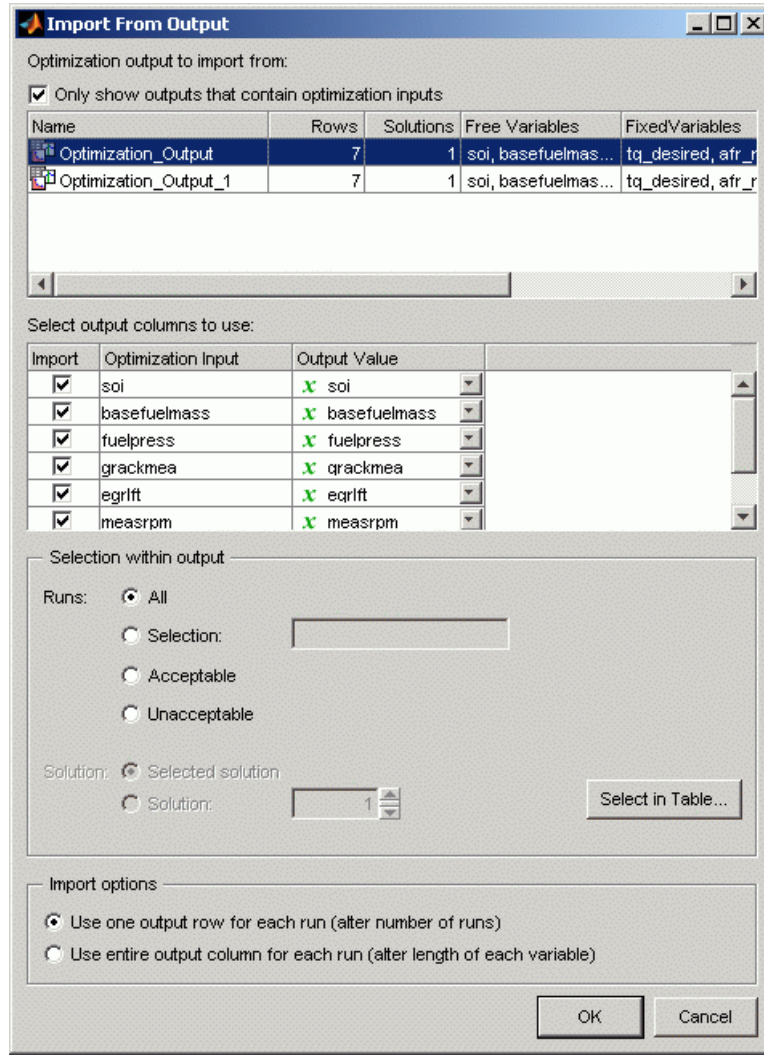

a Select the first `Optimization_Output` (there are two because the duplicate optimization also contains an output node).

b Click **OK** to close the Import From Output dialog box and import the output values for all the free and fixed variables. These variables are the starting values for the new optimization. Review the new values in the Input Variable Values pane in the Optimization view.

4. Select **Optimization > Convert to Single Run**. In the Input Variable Values pane, there is now a single run with 7 values, instead of the previous 7 runs of single values.

5. To run a sum optimization over the drive cycle, you must convert the objective to a sum objective. Double-click `Objective1` in the Objectives pane to open the Edit Objective dialog box.

   a. Select `Sum Objective` in the **Objective type** drop-down menu.

   b. Select `bsfc` from the **Model** list, and verify that **Minimize** is selected for the objective type.

   c. Click **OK** to close the dialog box.

Now your sum optimization is ready to run. Before you run it, however, see the next section to create a brake specific NOx constraint.

Creating the Brake Specific NOx Constraint

In this optimization, you will use brake specific NOx to constrain the results.

BSNOx is defined as:

$$BSNOx = \frac{\sum_{i=1}^{N} w_i NOx_i}{2\pi/60000 \sum_{i=1}^{N} w_i N_i TQ_i}$$

You implement this as a weighted sum constraint in CAGE. You create a function model that represents the following term of the BSNOx sum:

$$BSNOx_term = \frac{NOx}{2\pi/60000 \sum_{i=1}^{N} w_i N_i TQ_i}$$

The denominator of this term represents the weighted power (speed x torque) over the drive cycle, which is constant for this optimization. This constant can be calculated as follows:

$$
\begin{aligned}
&2\pi/60000 \sum_{i=1}^{N} w_i N_i TQ_i \\
&= 2\pi/60000 \begin{pmatrix} 2200\times1263 + 2200\times947 + 2200\times632 + 2200\times126 + \\ 1600\times1550 + 1600\times1163 + 1600\times775 \end{pmatrix} \\
&= 159.5573
\end{aligned}
$$

Therefore the function model to be created is

$$BSNOx_term = \frac{NOx}{159.5573}$$

Create the function model as follows:

1. Select **File > New > Function Model**.
2. Enter `BSNOX = 3600*1000*NOX/159.5573`, and click **Next**.

   The NOx model is in kilograms per hour, and you want to calculate BSNOx in grams per second, so multiply NOx by 3600*1000.
3. Select **Automatically assign/create inputs**. It is important to check for input errors or this step can create unintended new inputs. Click **Finish**.

   The Models view appears with the new function model. Now, you can use this function model to create a weighted sum constraint.

Setting Weights for the Sum Objective and Constraint

To create a weighted sum constraint:

1. Click the Optimization button to return to the Optimization view.
2. Right-click in the Constraints pane, and select **Add Constraint**. The Edit Constraint dialog box appears.

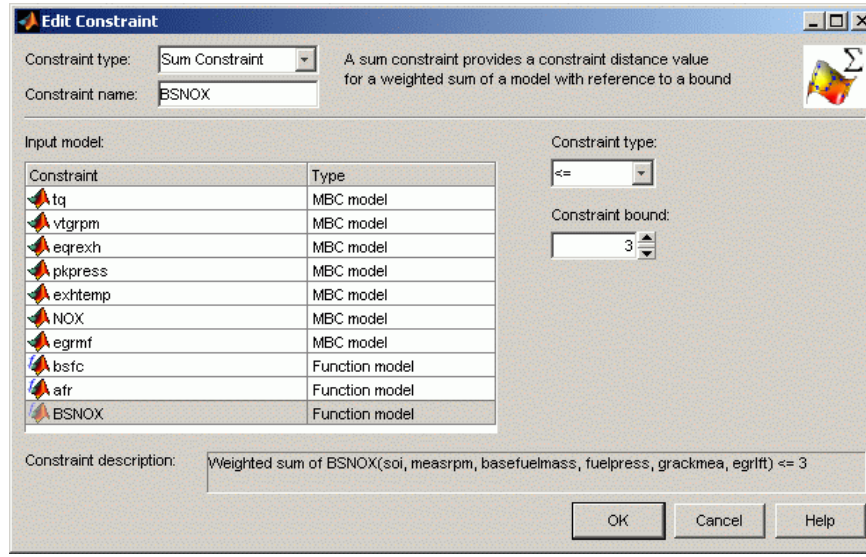

3 Select `Sum Constraint` from the **Constraint type** drop-down menu.

4 Enter `BSNOX` in the **Constraint name** edit box.

5 Select `BSNOX` from the **Input model** list.

6 Enter `3` in the **Constraint bound** edit box, and press **Enter**. Verify that the **Constraint description** shows Weighted sum of BSNOX <= 3, to constrain the weighted sum to less than 3g/kWh.

7 Click **OK** to create the constraint, and return to the Optimization view.

Set the weights in the Optimization view as follows:

1 Enter `7` in the **Number of values** edit box for the Objective1_weights column.

2 Enter, or copy and paste, the following values into the Objective1_weights column in the **Fixed Variables** pane.

Objective 1 Weights
0.5

Objective 1 Weights
0.15
0.15
0.05
0.05
0.05
0.05

These values sum to 1.

3 Similarly, enter 7 in the **Number of values** edit box for the Constraint1_weights column.

4 Copy the following values into the Constraint1_weights column in the **Fixed Variables** pane.

Constraint 1 Weights
0.15
0.15
0.15
0.1
0.1
0.1
0.1

These values sum to 0.85.

Set Parameters and Run Optimization

To alter parameters and then run the optimization:

1 Click Set Up and Run Optimization in the toolbar.

2 Change the following parameters in the Optimization Parameters dialog box:

- **Maximum iterations** = `500`
- **Maximum function evaluations** = `2000`
- **Function tolerance** = `1e-7`
- **Variable tolerance** = `1e-7`

Click **OK** to close the dialog box, and the optimization runs.

3 When the optimization is complete, the view switches to the updated child node, `Optimization_Output_1`, in the Optimization tree under the `Sum_optimization` node. (Recall that you duplicated an optimization, including its output node). View the results. For guidance see "Interpreting Sum Optimization Output".

Filling Tables with Optimization Results

You can use the optimization results to fill the tables you created.

1. From the optimization output node, select **Solution > Fill Tables**.

   The Table Filling Wizard appears.

2. **Shift**-click to multi-select all the following tables and click the button to add all your tables to the filling list:

   - `soi_best`
   - `basefuelmass_best`
   - `fuelpress_best`
   - `grackmea_best`
   - `egrlft_best`

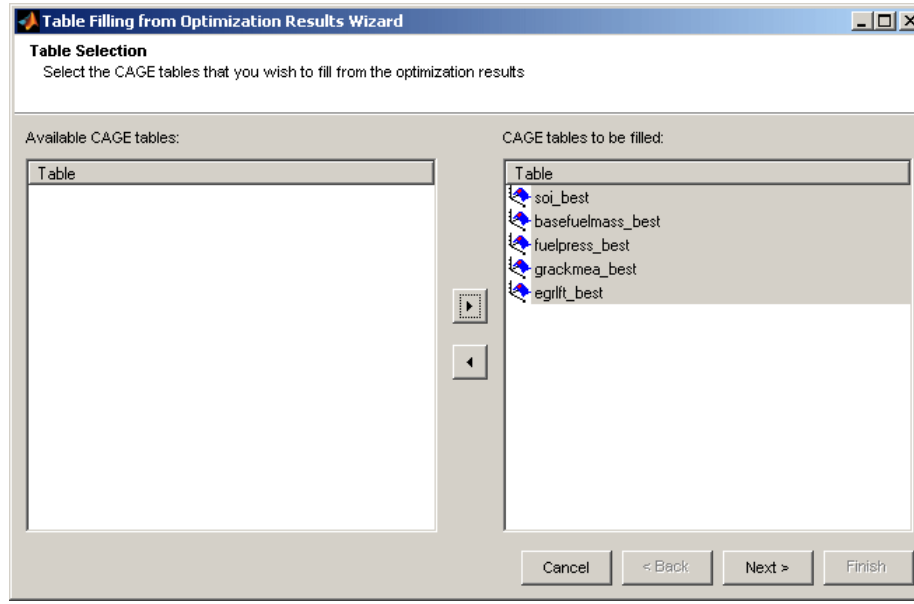

Click **Next**.

3 Match the following pairs of an output from the right list of optimization results with a table on the left.

- `soi` with `soi_best`
- `basefuelmass` with `basefuelmass_best`
- `fuelpress` with `fuelpress_best`
- `grackmea` with `grackmea_best`
- `egrlft` with `egrlft_best`

You can double-click in the right list to select an item, and CAGE automatically moves focus on to the next table in the left list. For example, with the `soi_best` table selected in the left list, double-click `soi` in the right list, then you can double-click `basefuelmass` for the next table, and so on.

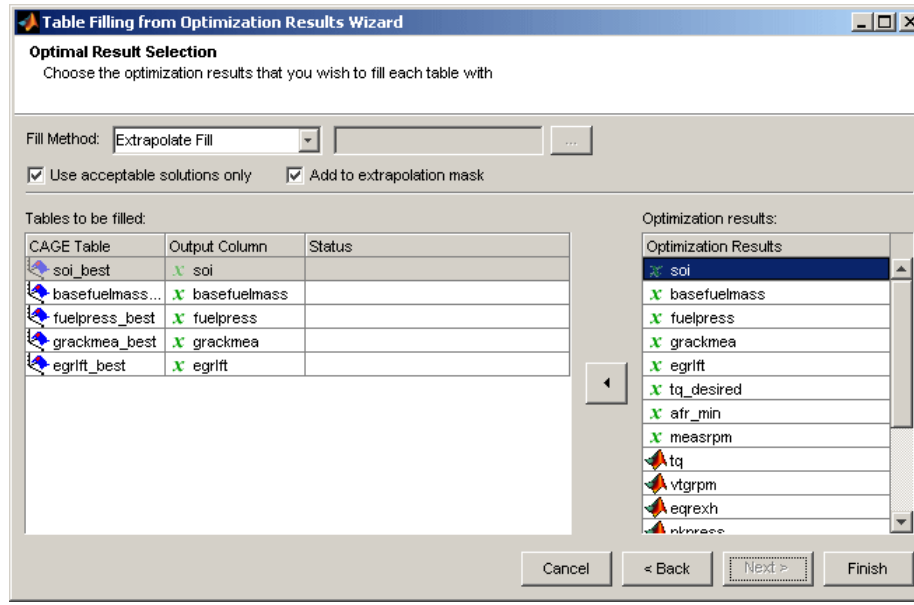

4 Click **Finish**.

A dialog appears indicating whether the tables have been filled successfully. Click **Close**.

5 Switch to the **Tables** view to view the filled tables. Click **Tables** in the **Data Objects** pane, and select the tables in turn to view the results.

Look at the example finished project, `Diesel_optimization.cag`.

4

Command-Line Interface to the Model-Based Calibration Toolbox Product

This section contains the following topics:

- "Introduction to the Command-Line Interface" on page 4-2
- "Processes You Can Automate" on page 4-3
- "Understanding Model Structure" on page 4-6
- "How the Model Tree Relates to Command-Line Objects" on page 4-9
- "Command-Line Demos" on page 4-14

Introduction to the Command-Line Interface

The Model-Based Calibration Toolbox product is a software tool for modelling and calibrating powertrain systems. The command-line interface to the Model-Based Calibration Toolbox product enables the design of experiments and modeling tools available in the toolbox to be accessible from the test bed.

You can use these commands to assemble your specific engine calibration processes into an easy to use script or graphical interface. Calibration technicians and engineers can use the custom interface without the need for extensive training. This system enables:

- Transfer of knowledge from the research and development engineers into the production environment
- Faster calibration
- Improved calibration quality
- Improved system understanding
- Reduced development time

See Model-Based Calibration Toolbox Demos for command-line examples.

Processes You Can Automate

The following description illustrates an example engine modeling process you can automate with the command-line Model-Based Calibration Toolbox product. You can assemble the commands for these steps into an easy-to-use script or graphical interface. This is a guideline for some of the steps you can use to model engine data.

1. Create or load a project — CreateProject; Load
2. Create a new test plan for the project using a template set up in the Model Browser — CreateTestplan
3. Create designs that define data points to collect on the test bed — CreateDesign. Work with classical, space-filling or optimal designs: CreateConstraint; CreateCandidateSet; Generate; FixPoints; Augment.
4. Create or load a data object for the project and make it editable — CreateData; BeginEdit
5. Load data from a file or the workspace — ImportFromFile; ImportFromMBCDataStructure

   You can instead specify the required data file when you call CreateData; you must still call BeginEdit before you can then make changes to the data.
6. Work with the data:
   - Examine data values — Value
   - Modify the data to remove unwanted records — AddFilter; AddTestFilter
   - Add user-defined variables to the data — AddVariable
   - Add new data — Append
   - Group your data for hierarchical modeling by applying rules — DefineTestGroups; DefineNumberOfRecordsPerTest
   - Export your data to the workspace — ExportToMBCDataStructure
7. Save your changes to the data, or discard them — CommitEdit; RollbackEdit

8. Designate which project data object to use for modeling in your test plan — AttachData

9. Create models for the data; these can be one- or two-stage models and can include datum models — CreateResponse

10. Work with your models:

    - Examine input data and response data — DoubleInputData; DoubleResponseData
    - Examine predicted values at specified inputs — PredictedValue; PredictedValueForTest
    - Examine Predicted Error Variance (PEV) at specified inputs — PEV; PEVForTest
    - Examine and remove outliers — OutlierIndices; OutlierIndicesForTest; RemoveOutliers; ; RestoreData
    - Create a selection of alternative models — CreateAlternativeModels
    - Choose the best model by using the diagnostic statistics —- AlternativeModelStatistics; DiagnosticStatistics; SummaryStatistics
    - Extract a model object from any response object (Model Object), then:
        - Fit to new data (Fit)
        - Create a copy of the model, change model type, properties and fit algorithm settings (CreateModel; ModelSetup; Type (for models); Properties (for models); CreateAlgorithm)
        - Include and exclude terms to improve the model (StepwiseRegression)
        - Examine regression matrices and coefficient values (Jacobian; ParameterStatistics)
        - If you change the model you need to use UpdateResponse to replace the new model back into the response in the project.

- For two-stage test plans, once you are satisfied with the fit of the local and response feature models (and have selected best models from any alternatives you created), you can calculate the two-stage model — MakeHierarchicalResponse.

- Now you can also examine the predicted values and PEV of the two-stage model.
- You can export any of these models to MATLAB or Simulink software — Export

This overview is not an exhaustive list of the commands available. For that, see “Function Reference” and “Commands — Alphabetical List” in the *Model-Based Calibration Toolbox Reference*.

Understanding Model Structure

In this section...
"Overview of Model Structure" on page 4-6
"Projects" on page 4-6
"Test Plans" on page 4-6
"Responses" on page 4-6

Overview of Model Structure

To use the Model Browser in the Model-Based Calibration Toolbox product, you must understand the structure and functions of the model tree to navigate the views. To use the command-line version of the toolbox, you must understand the same structure and functions, that is, how projects, test plans, and models fit together. The following sections describe the relationship between the different models that you can construct. The diagrams in the following section, "How the Model Tree Relates to Command-Line Objects" on page 4-9, illustrate these relationships.

Projects

- Projects can have one or more test plans.
- Projects can have one or more data objects.

Test Plans

- Test plans have no more than one data object.
- Test plans have response objects.
  - If a one-stage test plan, these are simply known as responses.
  - If two-stage test plan, these are hierarchical responses.

Responses

A response is a model fitted to some data. These are the types of responses:

- Hierarchical Response (Level 0)

  A hierarchical response (also known as a two-stage response) models a `ResponseSignalName` using a local response and one or more response features.

  A hierarchical response has one or more different local responses (accessible via the property `LocalResponses`) that provide different possible models of the `ResponseSignalName`. One of these must be chosen as the best, and that will then be the local response used subsequently. The response features of each of the local responses are available directly from those local response objects.

- Local Response (Level 1)

  The local response consists of models of the `ResponseSignalName` as a function of the local input factors. The local input factors are accessible via the `InputSignalNames` property.

  A local response has one or more response features (accessible via the property `ResponseFeatures`) containing the models fitted to those response features of the local model.

- Response (Level 1 or 2)

  - For two-stage test plans, response objects model the response features of local responses (`ResponseSignalName` corresponds to the name of the response feature). In this case, the response has a level value of 2.
  - For one-stage test plans, response objects simply model the `ResponseSignalName` as a function of the input factors. In this case, the response will have a level value of 1.

  All responses can have zero or more alternative responses (accessible via the property `AlternativeResponses`) that provide different possible models of the `ResponseSignalName`. These all retain the same level as the response for which they are an alternative. One of these must be chosen as the best and that will then be the response used subsequently.

See the illustrations in the following section, "How the Model Tree Relates to Command-Line Objects" on page 4-9, for examples of different responses and how they relate to each other.

Note that each response contains a model object (`mbcmodel.model`) that can be extracted and manipulated independently of the project. You can change the model type and settings, fit to new data, examine coeffficients, regression matrices and predicted values, and use stepwise functions to include or remove terms. You can change model type, properties and fit algorithm settings. See “Models”Model Object. If you change the model, you must use `UpdateResponse` to replace the new model type in the response object in the project. When you use `UpdateResponse` the new model is fitted to the response data. See UpdateResponse.

How the Model Tree Relates to Command-Line Objects

The tree in the Model Browser displays the hierarchical structure of models. This structure must be understood to use the command-line interface. The following examples illustrate the relationship between projects, test plans and responses in one-stage and two-stage models.

The following is an example of a two-stage model tree.

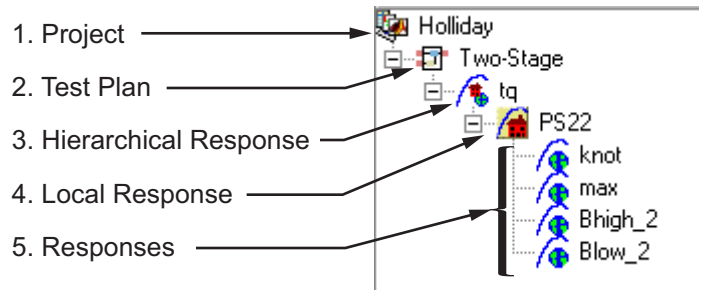

The elements of the tree correspond to the following objects in the command-line interface:

1 Project

2 Test Plan

3 Hierarchical Response

4 Local Response

5 Responses

The following example illustrates a project containing a one-stage test plan; in the command-line interface this corresponds to a project, one-stage test plan, and a response model.

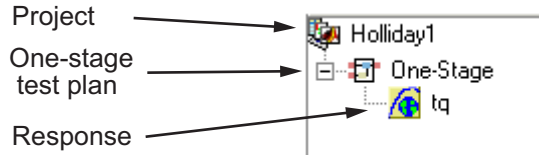

Hierarchical responses can have multiple local responses, as shown in the following example from the Model Browser. In the command-line interface these are accessible via the property LocalResponses for a hierarchical response object (`mbcmodel.hierarchicalresponse`). In this example, the local responses are `PS22`, `PS32`, and `POLY2`.

Only one of these local responses can be chosen as best (in this example, `PS22`, indicated by the blue icon) and used to construct the hierarchical response, together with the associated response features of the local response. Each local response object has a set of responses, accessible by the property ResponseFeatures(Local Response).

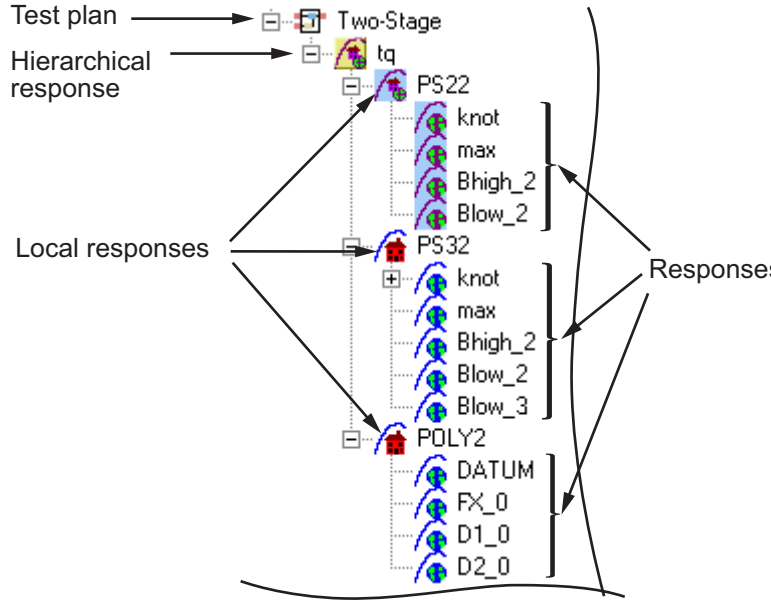

Responses can have zero or more alternative responses, as shown in the following model tree. You call the method CreateAlternativeModels on the command line to do the same.

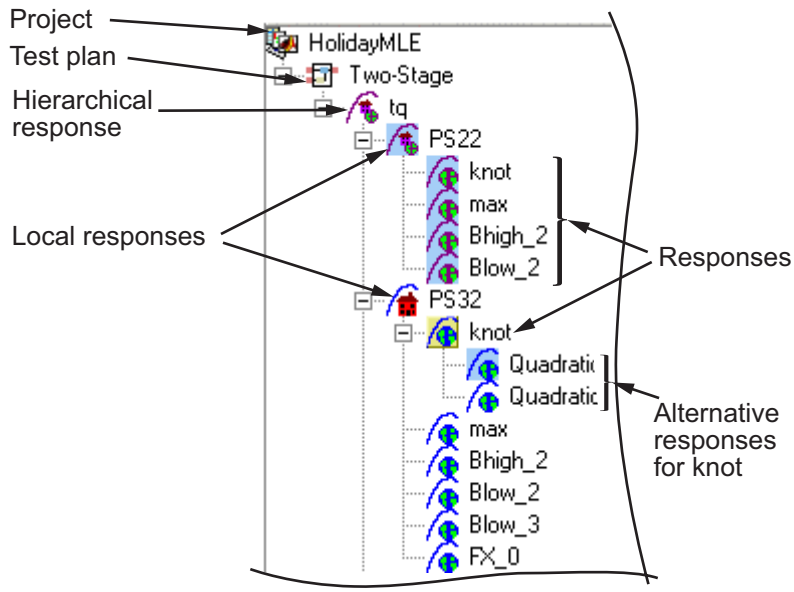

In this example, the alternative responses for the `knot` response are accessible via the property AlternativeResponses. You can create alternative responses for any response (including all one-stage responses).

You can use model templates to try alternative model types for several responses. The following example shows the results of using a model template for four alternative responses (`Linear-RBF`, `RBF-multiquadric`, `Cubic`, and `Quadratic`). The model template has been used to create alternative responses for the responses `knot` and `max`. You can call the method CreateAlternativeModels to do this in the command-line interface.

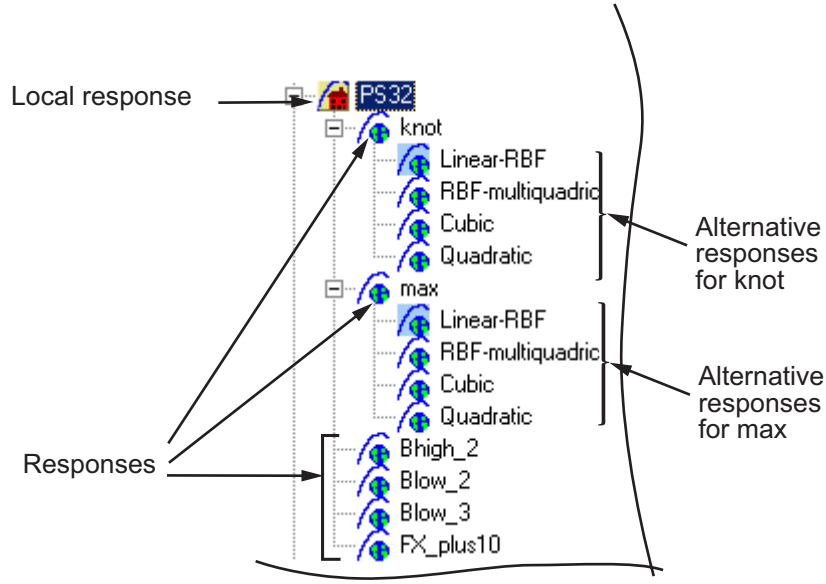

One of the alternative responses must be chosen as best for each response (call the method ChooseAsBest). In this example, when `Linear-RBF` is chosen as best from the alternatives for the `knot` response, then it is copied to `knot`.

Command-Line Demos

See Model-Based Calibration Toolbox Demos for a selection of command-line demos.

5

Tutorial: Model Quickstart

This section discusses the following topics:

- "Two-Stage Models" on page 5-2
- "Starting the Toolbox and Loading Data" on page 5-4
- "Setting Up the Model" on page 5-8
- "Verifying the Model" on page 5-24
- "Exporting the Model" on page 5-41
- "Creating Multiple Models to Compare" on page 5-43

Two-Stage Models

This tutorial gives you a quick introduction to the modeling end of the Model-Based Calibration Toolbox product. The instructions show you how to use the toolbox to create a statistical model of an automobile engine that predicts the torque generated by the engine as a function of spark angle and other variables. We provide example engine data and show you how to fit a statistical model to the data, examine and verify the fit, and export the model. In the normal modeling process, you would create many different models for one project and compare them to find the best solution. The tutorial also provides a quick guide to fitting and comparing multiple models.

Following is an explanation of how two-stage models are constructed and how they differ from one-stage models.

This tutorial is a step-by-step guide to constructing a single two-stage model for modeling engine brake torque as a function of spark, engine speed, load, and air/fuel ratio. One-stage modeling fits a model to all the data in one process, without accounting for the structure of the data. When data has an obvious hierarchical structure (as here), two-stage modeling is better suited to the task.

The usual way for collecting brake torque data is to fix engine speed, load, and air/fuel ratio within each test and sweep the spark angle across a range of angles. For this experimental setup, there are two sources of variation. The first source is variation within tests when the spark angle is changed. The second source of variation is between tests when the engine speed, load, and air/fuel ratio are changed. The variation within a test is called local, and the variation between tests, global. Two-stage modeling estimates the local and global variation separately by fitting local and global models in two stages. A local model is fitted to each test independently. The results from all the local models are used to fit global models across all the global variables. Once the global models have been estimated, they can be used to estimate the local models' coefficients for any speed, load, and air/fuel ratio. The relationship between the local and global models is shown in the following block diagram, as you will see in the Model Browser.

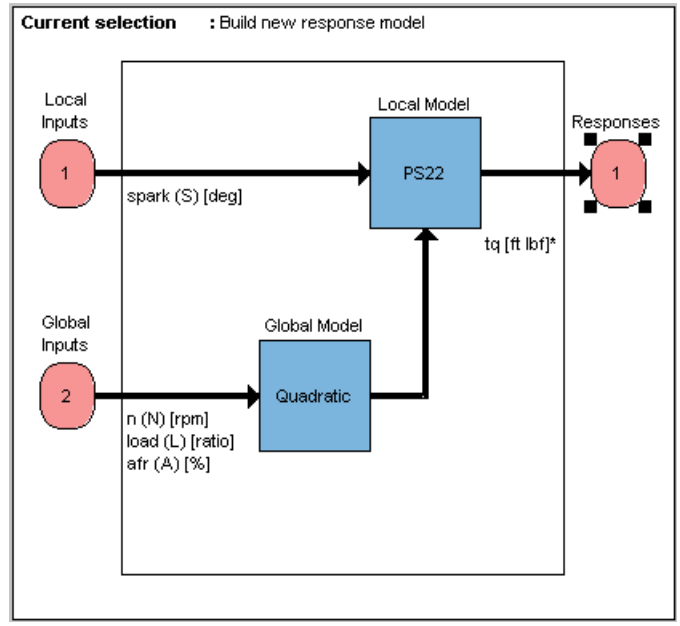
Current selection : Build new response model
Local Inputs
1
spark (S) [deg]
Local Model
PS22
Responses
1
tq [ft lbf]*
Global Inputs
2
n (N) [rpm]
load (L) [ratio]
afr (A) [%]
Global Model
Quadratic

Starting the Toolbox and Loading Data

1. Double-click the MATLAB icon to start MATLAB.
2. To start the Model Browser, enter `mbcmodel` at the command prompt in MATLAB.
3. If you have never used the toolbox before, the User Information dialog appears. If you want, you can fill in any or all of the fields: your name, company, department, and contact information, or you can click **Cancel**. The user information is used to tag comments and actions so that you can track changes in your files (it does not collect information for The MathWorks).

   Note You can edit your user information at any time by selecting **File > Preferences**.

   ---

4. When you finish with the User Information dialog, click **OK**.

   The Model Browser window appears.

In this window, the left pane, **All Models**, shows the hierarchy of the models currently built in a tree. At the start, only one node, the project, is in the tree. As you build models, they appear as child nodes of the project. The right panes change, depending on the tree node selected. You navigate to different views by selecting different nodes in the model tree. Different tips can appear in the **Tip of the Day** pane.

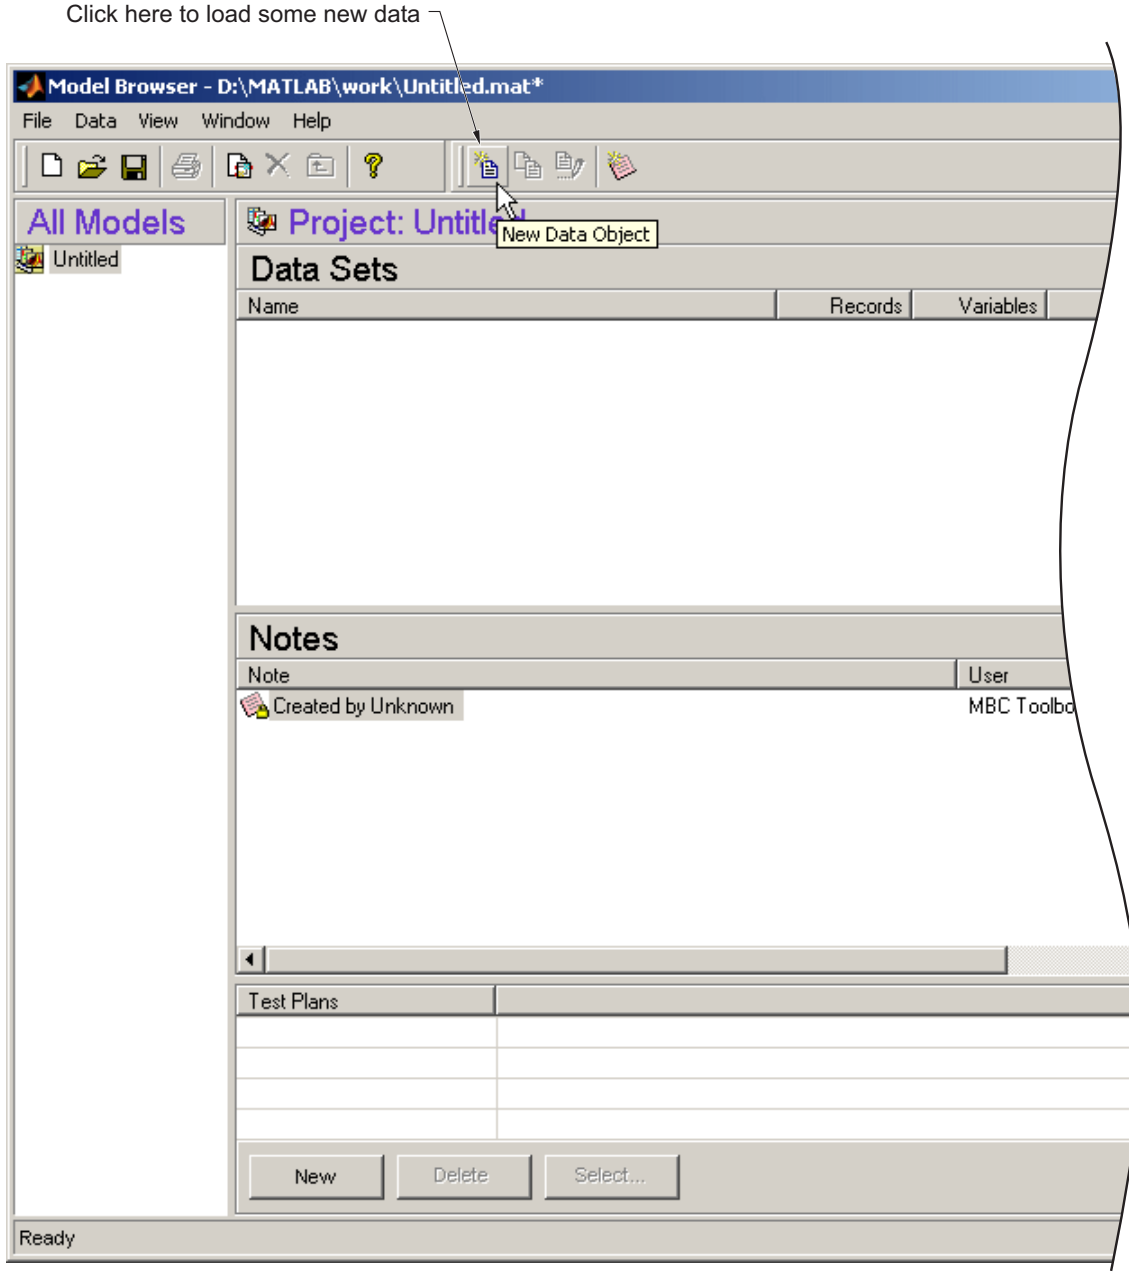
Click here to load some new data
Model Browser - D:\MATLAB\work\Untitled.mat*
File Data View Window Help
New Data Object
All Models
Untitled
Project: Untitled
Data Sets
Name
Records
Variables
Notes
Note
User
Created by Unknown
MBC Toolbo
Test Plans
New
Delete
Select...
Ready

Load the example data file `holliday.xls`:

1. Click the button on the toolbar, or select **Data > New Data**.

   This opens the Data Editor window.

2. Click the Import File icon on the toolbar ( ) or select **File > Import > File**.

3. Use the browse button to the right of the edit box in the Data Import Wizard to open a file browser and find the file `holliday.xls` in the `mbctraining` directory. Click **Open** or double-click the file.

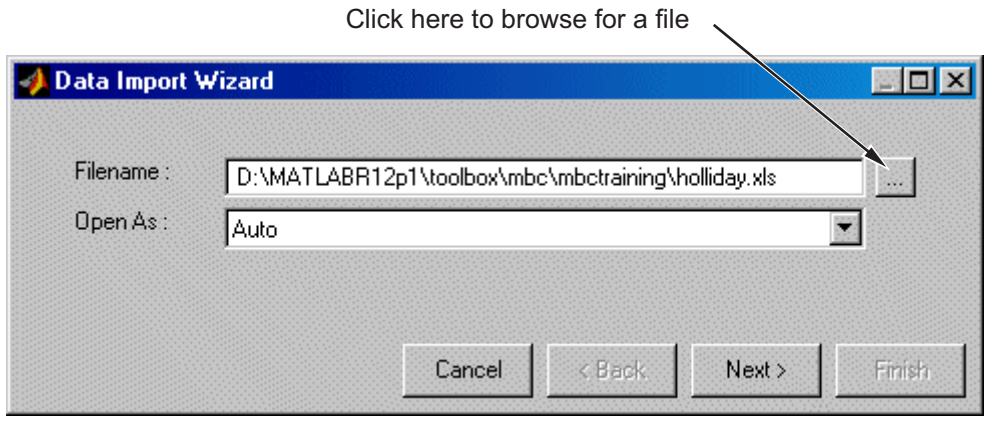

4. The file pathname appears in the Data Import Wizard. Click **Next**.

5. A summary screen displays information about the data. Click **Finish** to close the Data Import Wizard and return to the Data Editor.

   You can view plots of the data in the Data Editor by selecting variables and tests in the lists on the left side. Have a look through the data to get an idea of the shape of curve formed by plotting torque against spark.

   For more details on functionality available within the Data Editor, see “Data”.

6 Close the Data Editor to accept the data and return to the Model Browser. Notice that the new data set appears in the **Data Sets** pane.

This data is from Holliday, T., "The Design and Analysis of Engine Mapping Experiments: A Two-Stage Approach," Ph.D. thesis, University of Birmingham, 1995.

Setting Up the Model

In this section...
"Creating the Test Plan" on page 5-8
"Setting Up the Local Model" on page 5-10
"Setting Up the Global Model" on page 5-14
"Selecting Data" on page 5-18
"Specifying the Response Model" on page 5-22

Creating the Test Plan

Now you can use the data to create a statistical model of an automobile engine that predicts the torque generated by the engine as a function of spark angle and other variables.

Note It does not matter in which order you set up local and global models, as both must be completed before you set up the response model.

---

1 To create a new test plan, do one of the following:

- In the **Test Plans** list pane at the bottom, click **New**.

  Alternatively, click the New Test Plan button ( ) in the toolbar. Note that this button changes depending on which node is selected in the model tree. It always creates a child node (not necessarily a test plan node), as does the **New** button at the bottom.

  Or select **File > New Test Plan**.

The New Test Plan dialog box appears.

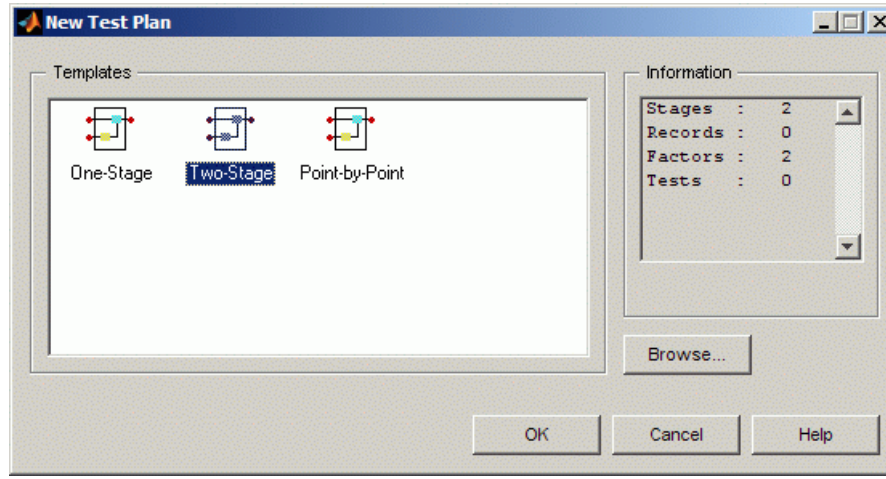

2 Click the two-stage test plan icon and click **OK**.

The default name of the new test plan, `Two-Stage`, appears in the Model Browser tree, in the **All Models** pane.

The view switches to the new test plan node, and the Model Browser window displays a diagram representing the two-stage model.

See also “Test Plan Level”.

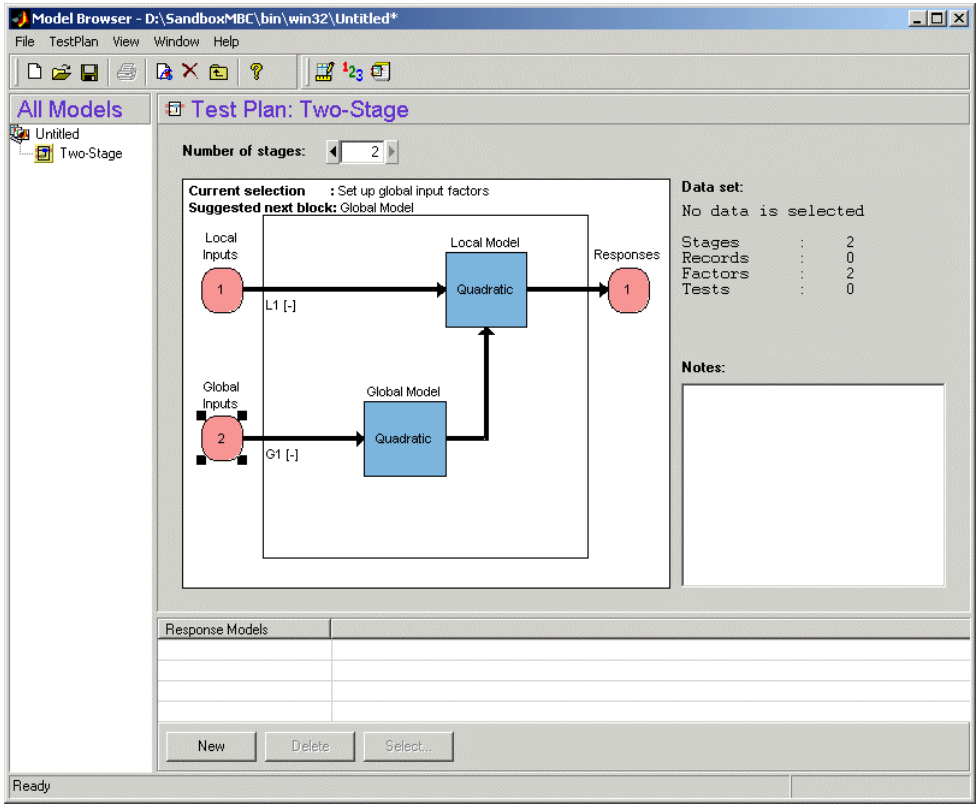

Setting Up the Local Model

Setting up the local model requires that you specify the model's inputs and type.

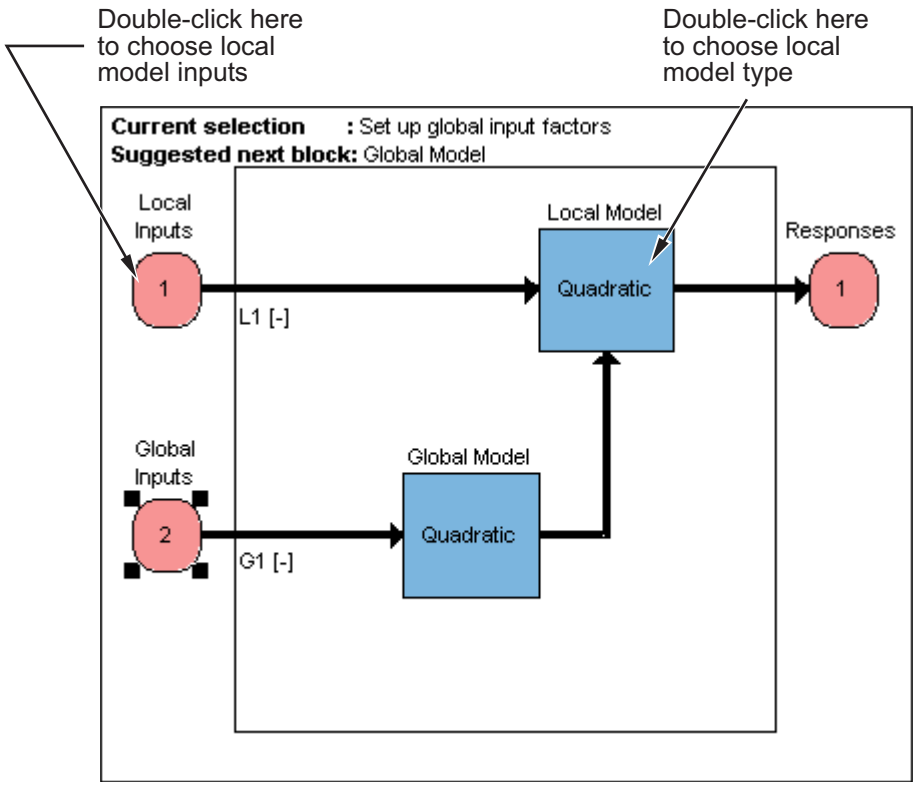

Specifying the Local Model Input

The model you are building is intended to predict the torque generated by an engine as a function of spark angle at a specified operating point defined by the engine's speed, air/fuel ratio, and load. The input to the local model is therefore the spark angle.

To specify spark angle as the input,

1. Double-click (or right-click) the Local Inputs icon on the model diagram to specify the local model input.

   The Local Input Factor Setup dialog box appears.

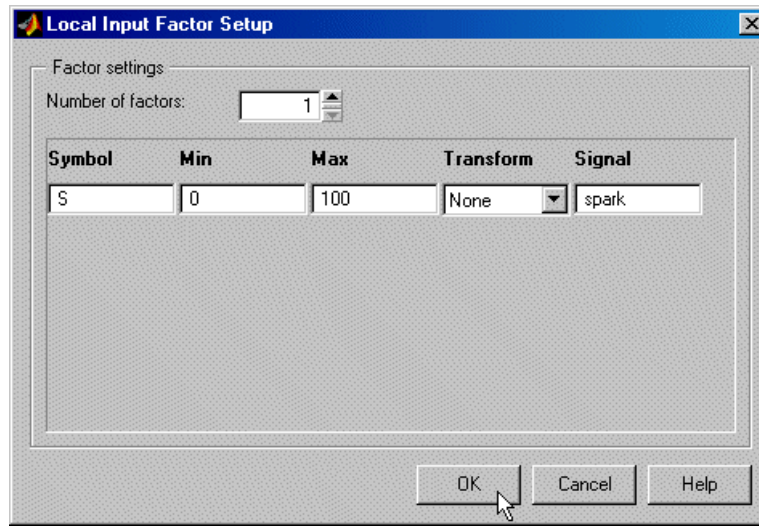

a Set **Symbol** to `S`.

b Set **Signal** to `spark`. This is optional and matches the raw data.

2 Click **OK** to dismiss the dialog box.

Notice that the new name of the local model input now appears on the two-stage model diagram.

Specifying the Local Model Type

The type of a local model is the shape of curve used to fit the test data, for example, quadratic, cubic, or polyspline curves. In this example, you use polyspline curves to fit the test data. A spline is a curve made up of pieces of polynomial, joined smoothly together. The points of the joins are called knots. In this case, there is only one knot. These polynomial spline curves are very useful for torque/spark models, where different curvature is required above and below the maximum.

To specify polyspline as the model type,

1 Double-click the local model icon in the model diagram.

The Local Model Setup dialog box appears.

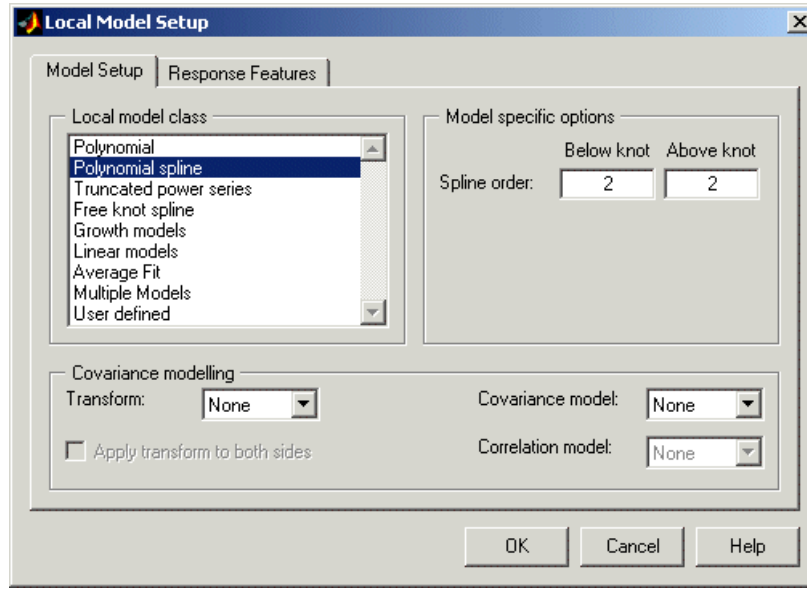

a Select `Polynomial Spline` from the **Local Model Class** list.

b Set **Spline Order** to `2` below and `2` above knot.

2 Click **OK** to dismiss the dialog box.

Notice that the new name of the local model class, `PS` (for polyspline) `2,2` (for spline order above and below knot) now appears on the two-stage model diagram.

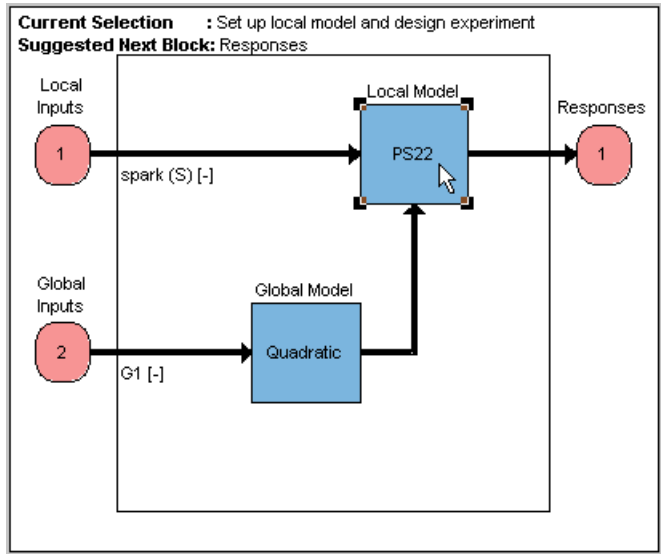

Setting Up the Global Model

Setting up the global model is similar to setting up the local model. You must specify the model (or curve) type and the inputs used to create the model.

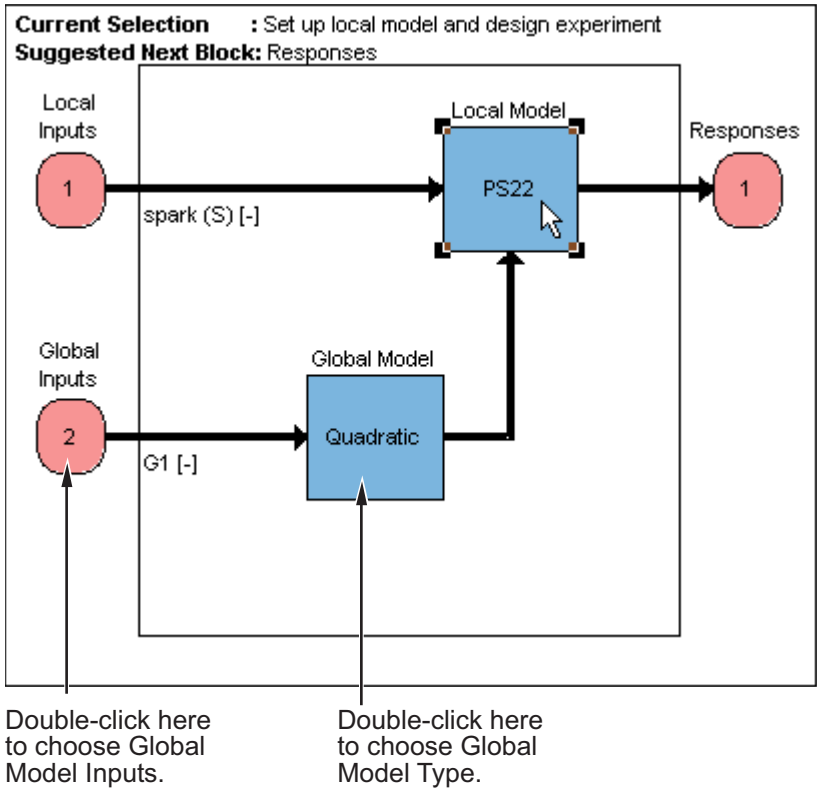

Specifying the Global Model Inputs

The inputs to the global model are the variables that determine the operating point of the system being modeled. In this example, the operating point of the engine is determined by the engine's speed in revolutions per minute (rpm - often called N), load (L), and air/fuel ratio (afr).

To specify these inputs,

1 Double-click the Global Inputs icon on the model diagram.

The Global Input Factor Setup dialog box appears.

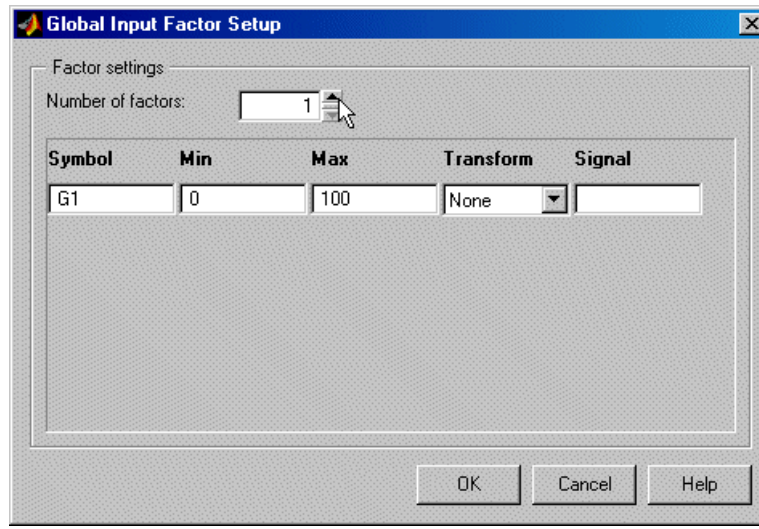

By default, there is one input to the global model. Because this engine model has three input factors, you need to increase the input factors as follows:

a Click the up arrow button indicated by the cursor above to increase the number of factors to three.

b Edit the three factors to create the engine model input. In each case, change the symbols and signals to the following:

Symbol	Signal
N	n
L	load
A	afr

c Leave the Min and Max boxes at the defaults (you fill them during the data selection process). You might want to set factor ranges at this stage if you were designing an experiment, but in this case there is already data available, so you use the actual range of the data to model instead.

2 Click **OK** to dismiss the dialog box.

Specifying the Global Model Type

Fitting the local model finds values for each model coefficient or response feature (for example, `knot`) for each test. These coefficients then become the data to which you fit the global model.

By default, quadratic polynomials are used to build the global model for each response feature. In this case you use the default.

To specify quadratic curves as the global model curve fitting method,

1. Double-click the icon representing the global model in the two-stage model diagram.

   The Global Model Setup dialog box appears.

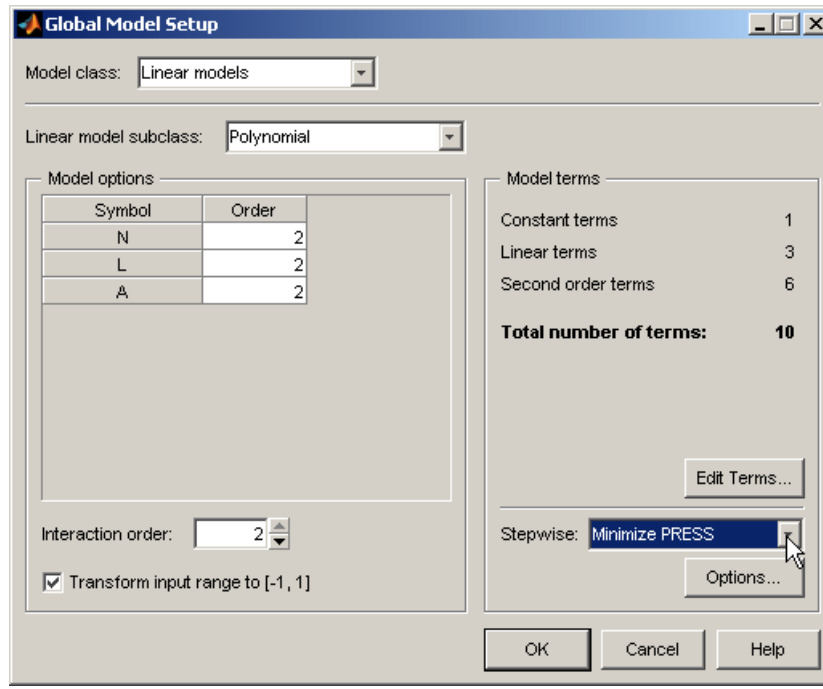

   a. `Polynomial` should already be selected from the **Linear Model Subclass** pop-up menu. Under **Model options**, the order for the three variables `N`, `L`, and `A` is set by default to `2`, which is required.

   b Set **Stepwise** to `Minimize PRESS` (PREdicted Sum Square error).

2. Click **OK** to accept the settings and dismiss the Model Settings dialog box.

You use the Stepwise feature to avoid overfitting the data; that is, you do not want to use unnecessarily complex models that "chase points" in an attempt to model random effects. Predicted error sum of squares (PRESS) is a measure of the predictive quality of a model. `Minimize PRESS` throws away terms in the model to improve its predictive quality, removing those terms that reduce the PRESS of the model.

This completes the setup of the global model.

Selecting Data

The model you have set up now needs data:

1. Double-click the Responses block in the Test Plan diagram. As no data has yet been selected for this test plan, this launches the Data Wizard.

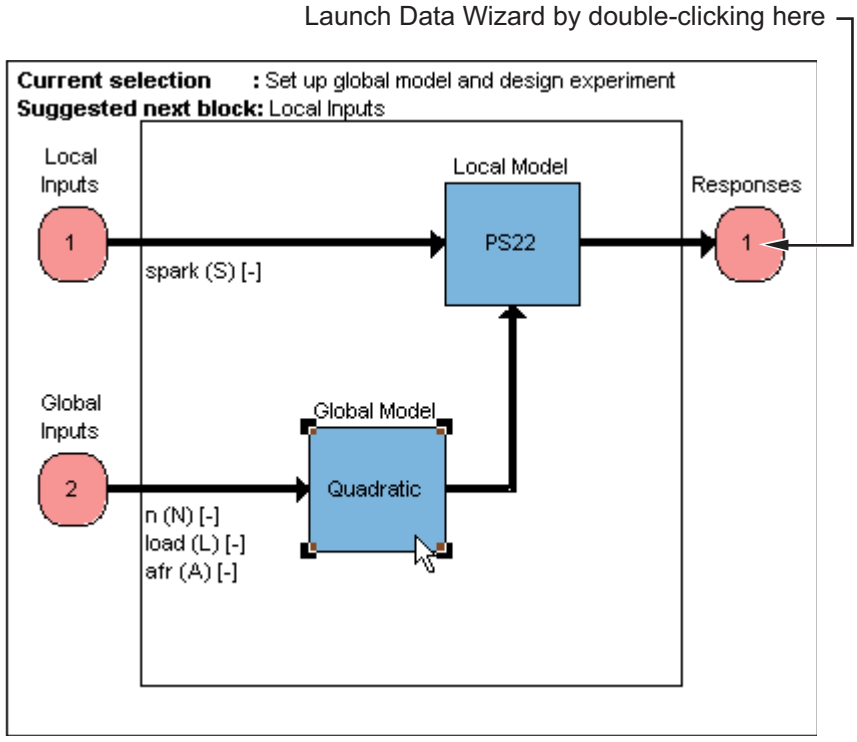

For the same result, you could also click the Select Data button in the toolbar of the Model Browser (or **TestPlan > Select Data** menu item). Also, if you did not already load a data set at the project node, you can do it at this point using **TestPlan > Load New Data**.

The Data Wizard dialog box appears.

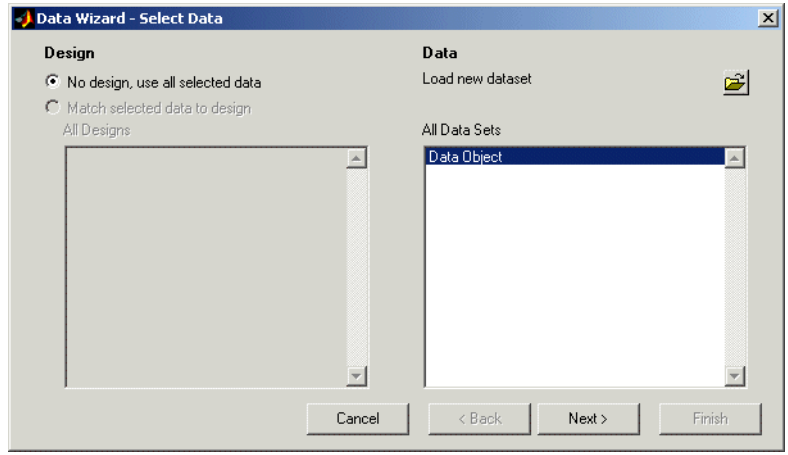

2 `Data Object` is already selected by default. Click **Next**.

3 Select `S` in the **Model Input Factors** box and `Spark` under **All Data Signals**.

If the signal name entered during the input factor setup matches a signal name in the data set, the wizard automatically selects the correct signal when the input factor is selected. If the name is not correct, you must select the correct signal manually by clicking. This autoselect facility can save time if the data set has a large number of signals.

4 Select the **Copy Range** check box, as shown in the following figure. This selection tells the model to use the range in the data for that factor.

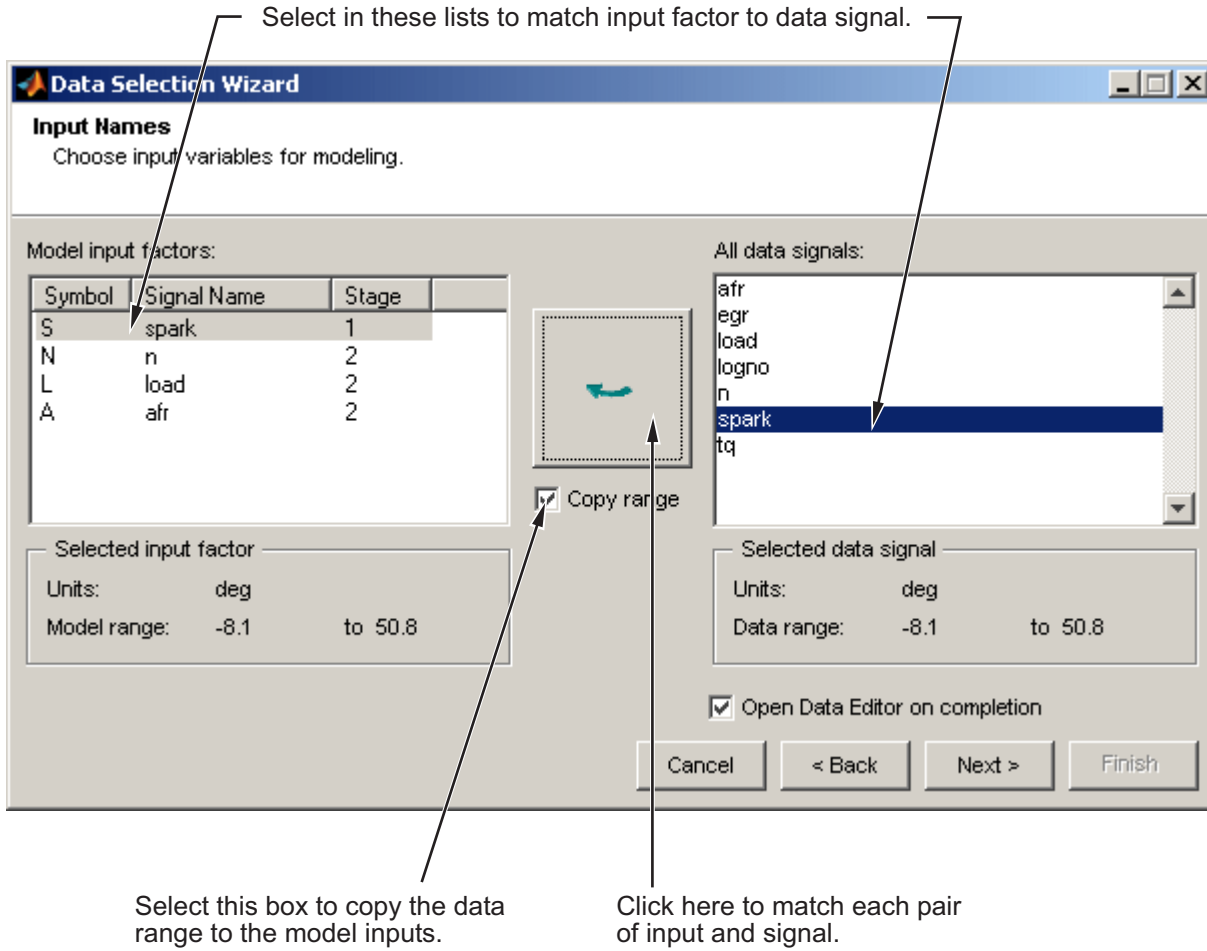

5 Click the large **Select Data Signal** button, as indicated above.

6 Repeat this process to match the correct data signals to the other three input factors, `N`, `L`, and `A` (from `n`, `load`, and `afr`).

7 When you have matched all four input factors to the correct data signals (for both stages of the two-stage model), click **Next**.

Specifying the Response Model

The model you just set up now needs a response specified (that is, the factor you want the model to predict, in this case, torque).

The next screen of the Data Wizard is for selecting response models.

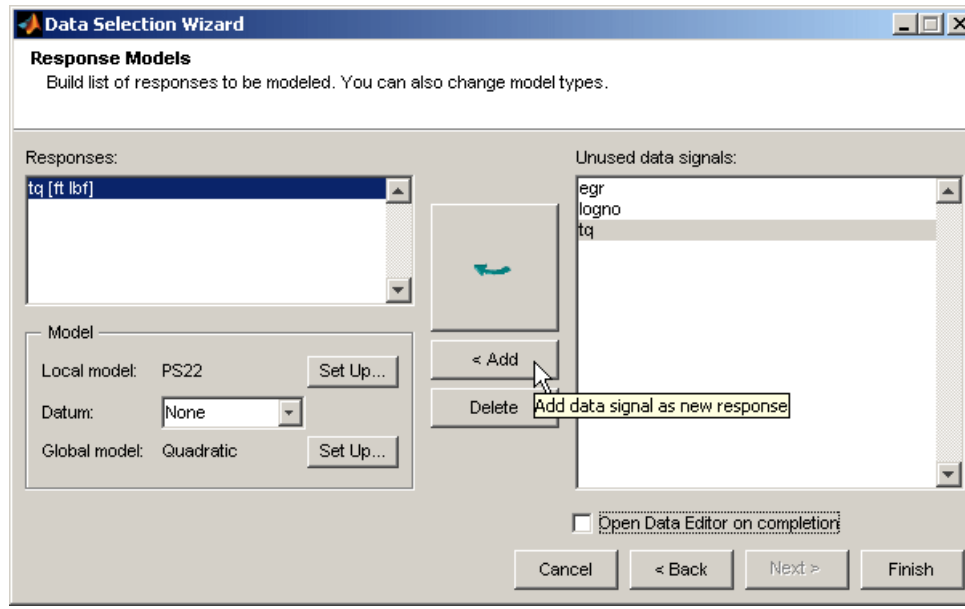

1. Select `tq` (torque) as the response.
2. Click **Add**. Torque appears in the **Responses** list.
3. Select `Maximum` under **Datum**.

   Only certain model types with a clearly defined maximum or minimum can support datum models. See “New Response Models and Datum Models”.
4. Clear the check box **Open Data Editor on completion**. You have already inspected the data.
5. Click **Finish**.

The Model Browser now calculates local and global models using the test plan models you just set up.

Notice that torque appears on the two-stage model diagram, and a new node appears on the tree in the **All Models** pane, called PS22.

Verifying the Model

In this section...

"Verifying the Local Model" on page 5-24
"Verifying the Global Model" on page 5-26
"Selecting the Two-Stage Model" on page 5-28
"Comparing the Local Model and the Two-Stage Model" on page 5-35
"Maximum Likelihood Estimation" on page 5-36
"Response Node" on page 5-39

Verifying the Local Model

The first step is to check that the local models agree well with the data:

1 Select `PS22` (the local node) on the Model Browser tree.

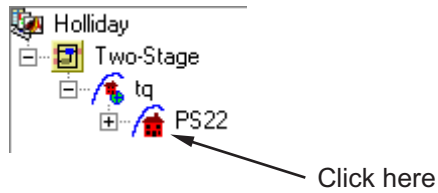

The **Local Model** pane appears, displaying the local model fitting the torque/spark data for the first test and diagnostic statistics that describe the fit. The display is flexible in that you can drag, open, and close the divider bars separating the regions of the screen to adjust the view.

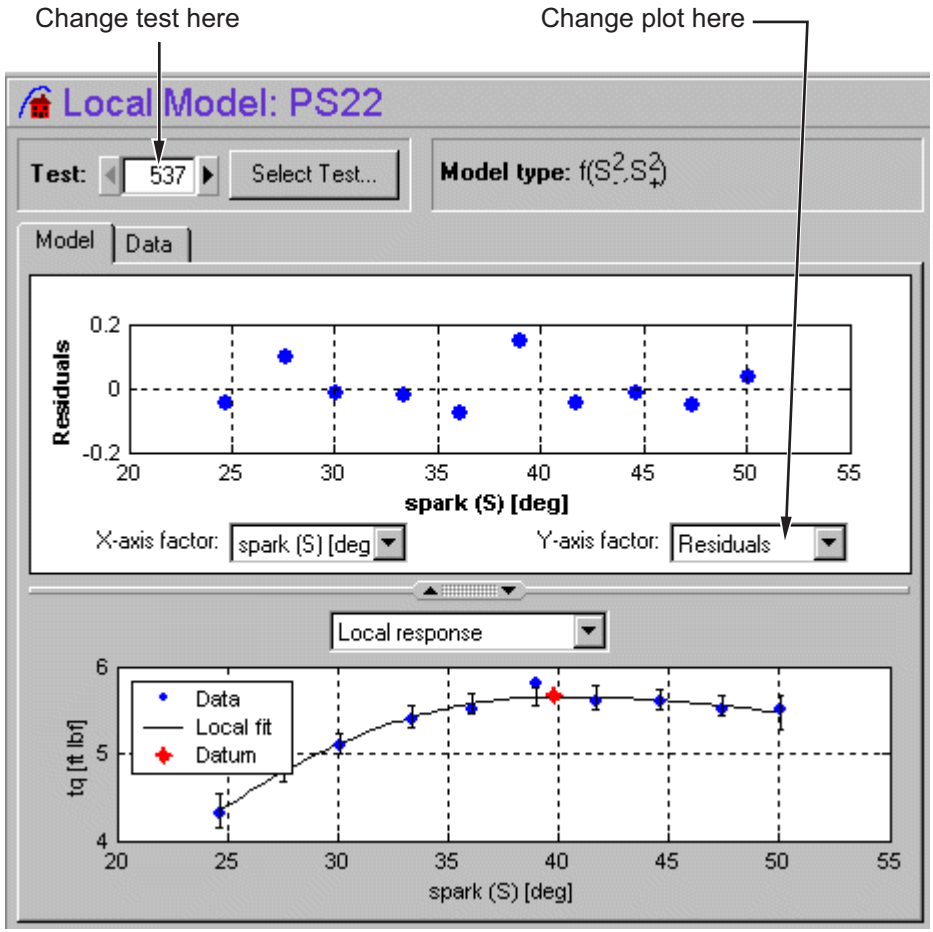

The lower plot shows the data being fitted by the model (blue dots) and the model itself (line). The red spot shows the position of the polyspline knot, at the datum (maximum) point.

2 In the upper scatter plot pane, click the **Y-axis factor** pop-up menu and select `Studentized residuals`.

3 To display plots and statistics for the other test data, scroll through the tests using the **Test** arrows at the top left, or by using the **Select Test** button.

4 Select Test 588. You see a data point outlined in red. This point has automatically been flagged as an outlier.

5 Right-click the scatter plot and select **Remove Outliers**. Observe that the model is refitted without the outlier.

Both plots have right-click pop-up menus offering various options such as removing and restoring outliers and confidence intervals. Clicking any data point marks it in red as an outlier.

You can use the **Test Notes** pane to record information on particular tests. Each test has its own notes pane. The test numbers of data points with notes recorded against them are colored in the global model plots, and you can choose the color using the **Test Number Color** button in the **Test Notes** pane. You can quickly locate tests with notes by clicking **Select Test**.

Verifying the Global Model

The next step is to check through the global models to see how well they fit the data:

1 Expand the `PS22` local node on the Model Browser tree by clicking the plus sign (+) to the left of the icon. Under this node are four response features of the local model. Each of these is a feature of the local model of the response, which is torque.

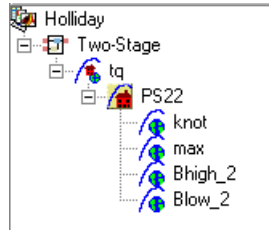

2 Select the first of the global models, `knot`.

The **Response Feature** pane appears, showing the fit of the global model to the data for `knot`. Fitting the local model is the process of finding values for these coefficients or *response features*. The local models produce a value of `knot` for each test. These values are the data for the global model for

`knot`. The data for each response feature come from the fit of the local model to each test.

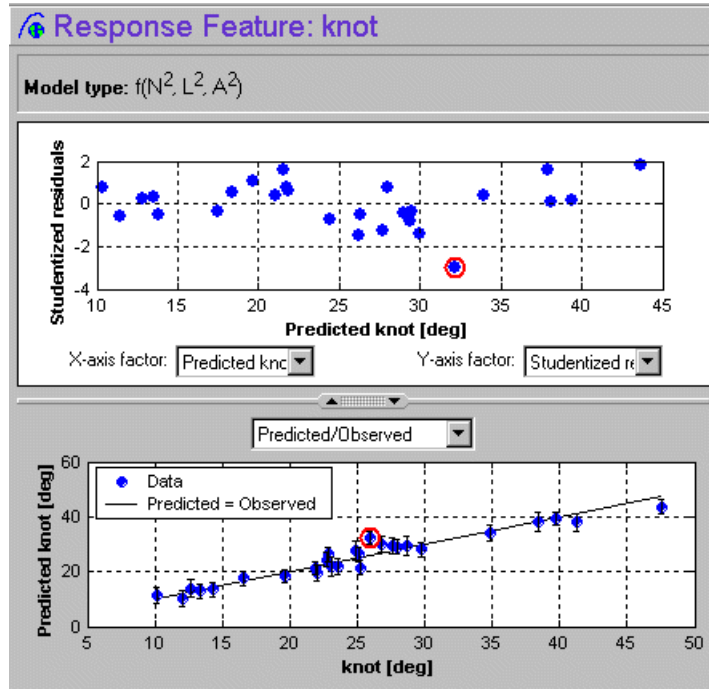

3 Select the response feature `Bhigh_2`. One outlier is marked. Points with an absolute studentized residual value of more than 3 are automatically suggested as outliers (but included in the model unless you take action). You can use the right-click menu to remove suggested outliers (or any others you select) in the same way as from the Local Model plots. Leave this one. If you zoom in on the plot (**Shift**-click-drag or middle-click-drag) you can see the value of the studentized residual of this point more clearly. Double-click to return to the previous view.

Note Never remove outliers as a matter of course. However, this tutorial is designed to show you how the toolbox helps you to do this when required. The default outlier selection criterion is a studentized residual greater than 3, to bring your attention to possible outliers, but you should never remove data without good reasons. Remove enough points and the model will simply interpolate the data and become useless for prediction. You can customize the criteria for outlier selection. Use the plot of Cook's Distance to see the influence of each point on the model fit to help you decide whether to remove an outlier.

4 Select the other response features in turn: `max` and `Blow 2`. You will see that `Blow 2` has a suggested outlier with a very large studentized residual; it is a good distance away from all the other data points for this response feature. All the other points are so clustered that removing this one could greatly improve the fit of the model to the remaining points, so remove it.

Return to the **Local Model** pane by clicking the local node `PS22` in the Model Browser tree.

Selecting the Two-Stage Model

Recall how two-stage models are constructed: two-stage modeling partitions the variation separately between tests and within tests, by fitting local and global models separately. A model is fitted to each test independently (local models). These local models are used to generate global models that are fitted across all tests.

For each sweep (test) of spark against torque, you fit a local model. The local model in this case is a spline curve, which has the fitted response features of `knot`, `max`, `Bhigh_2`, and `Blow_2`. The result of fitting a local model is a value for `knot` (and the other coefficients) for each test. The global model for `knot` is fitted to these values (that is, the `knot` global model fits `knot` as a function of the global variables). The values of knot from the global model (along with the other global models) are then used to construct the two-stage model

The global models are used to reconstruct a model for the local response (in this case, torque) that spans all input factors. This is the two-stage model across the whole global space, derived from the global models.

Now you can use the model selection features to view the fit of this two-stage model in various ways, to compare it with both the data and the local model fit.

Within this tutorial, you use the following:

- "Tests View" on page 5-32
- "Response Surface View" on page 5-33

For more detailed help on all the views available in the Model Selection window, see "Selecting Models".

Note To construct a two-stage model from the local and global models, you click the local node in the model tree (with the house icon) and click the **Select** button. This is the next step in the tutorial.

Once you are satisfied with the fit of the local and global models, it is time to construct a two-stage model from them. Return to the Local Model view by clicking the local node `PS22` in the Model Browser tree. The Model Browser should look like the following example.

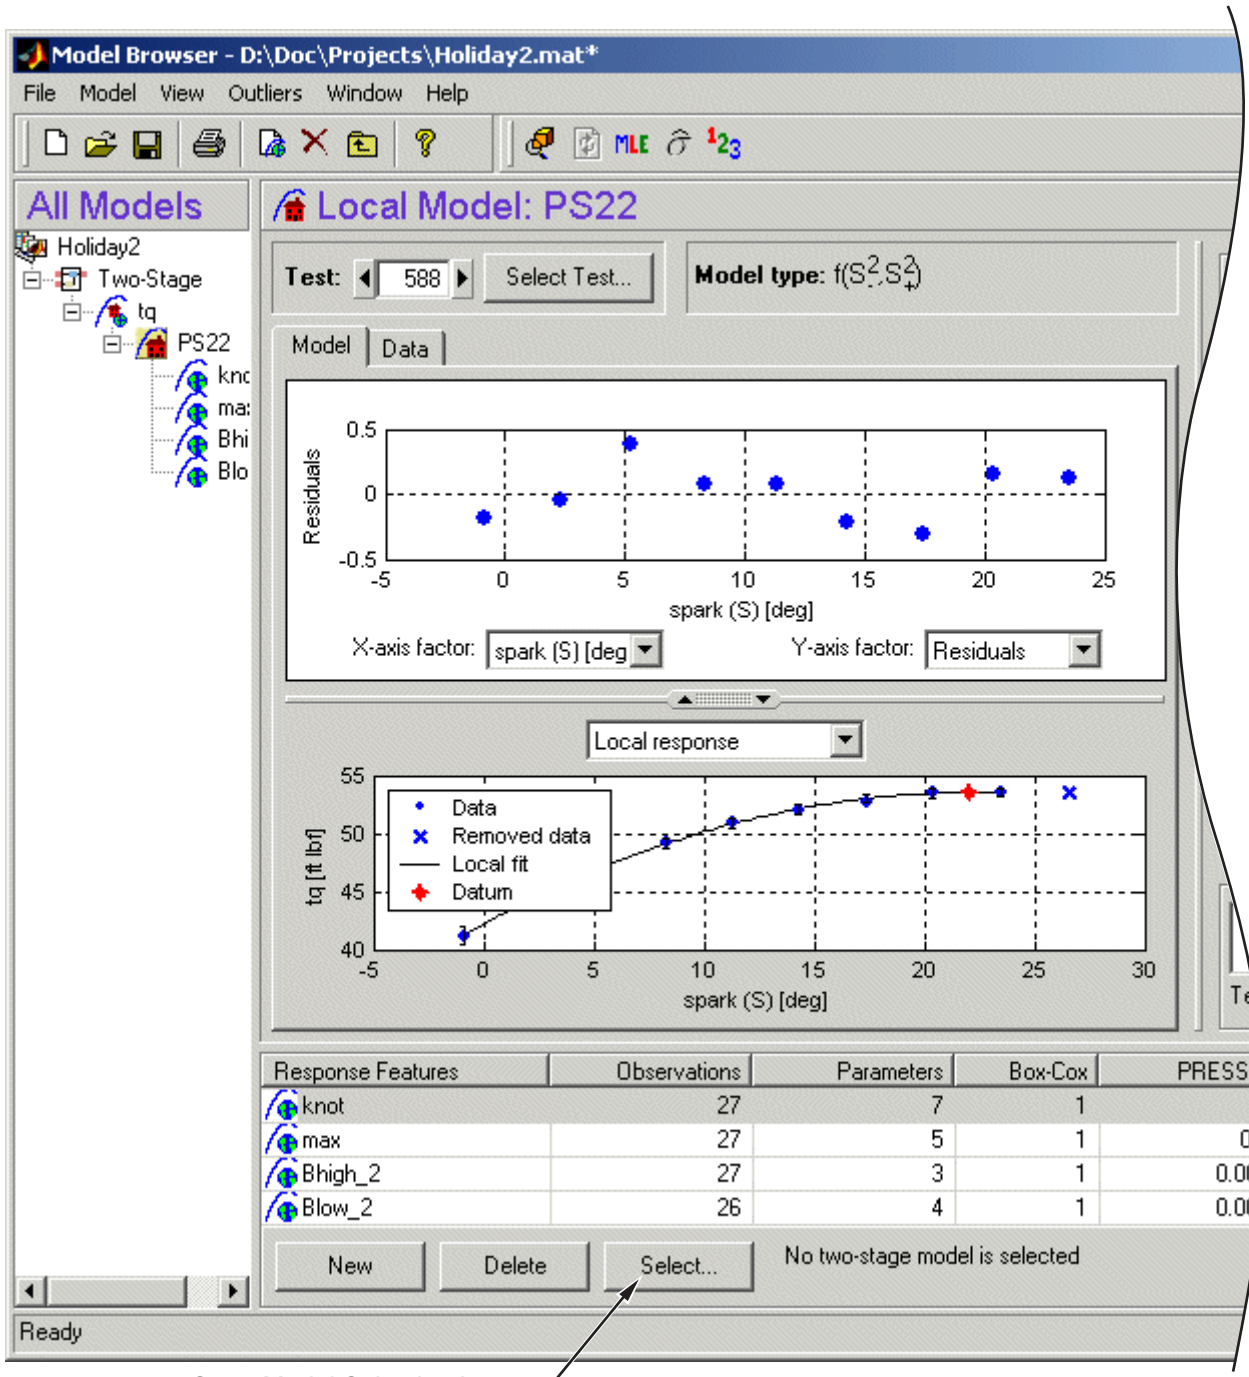

Open Model Selection here.

Click **Select** in the **Response Features** list pane, and the Model Selection window appears. This window is intended to help you select a *best model* by comparing several candidate models. A number of icons in the toolbar enable you to view the fit of the model in various ways. By default, the Tests view appears. These plots show how well the two-stage model agrees with the data.

Tests View

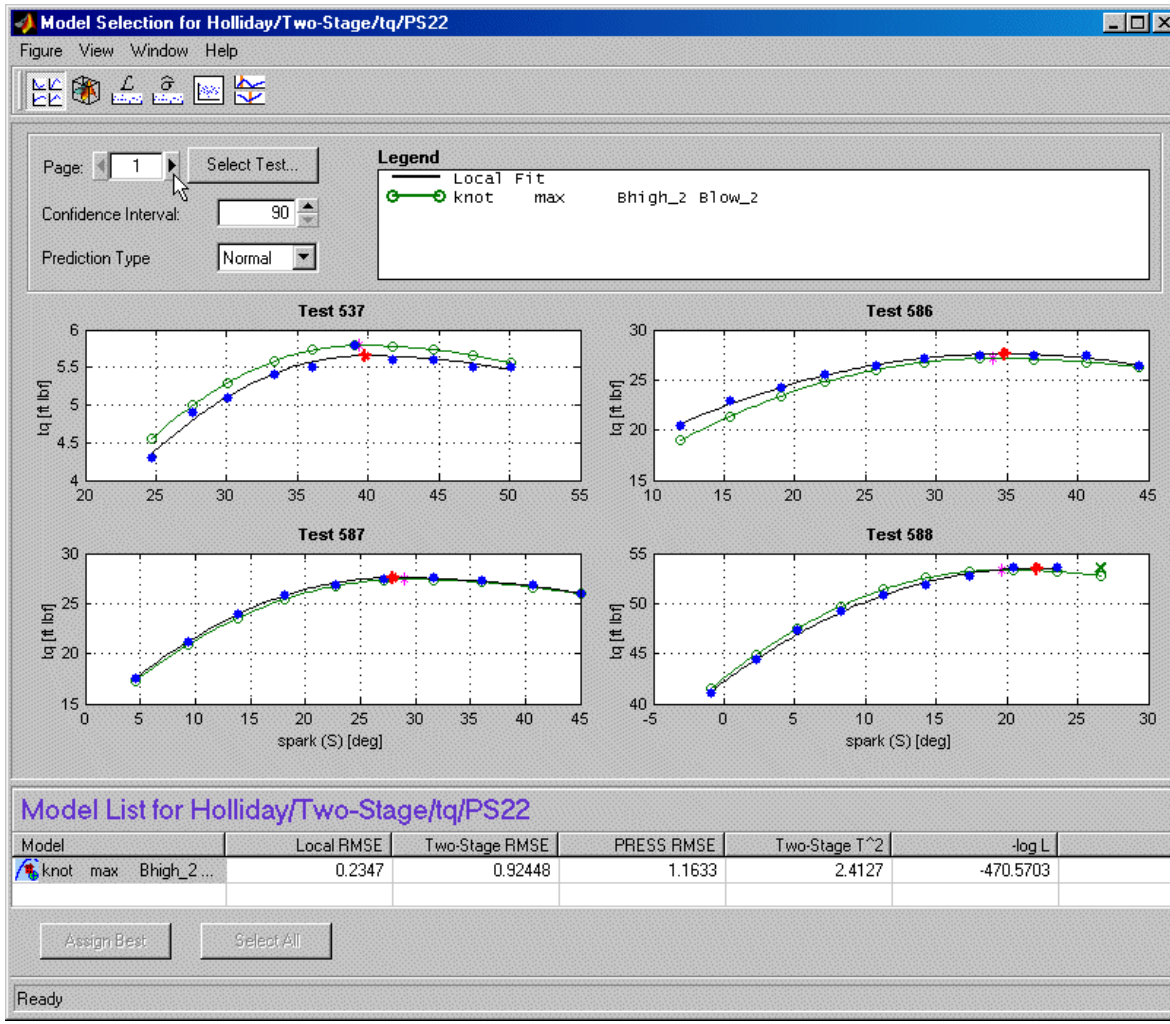

Scroll though the tests using the left/right arrows or the **Select Test** button at the top left. The plots show the fit of the two-stage model for each test (green open circles and line), compared with the fit of the local model (black line) and the data (blue dots). You can left-click (and hold) to see information on each test or zoom in on points of interest by **Shift**-click-dragging or middle-click-dragging. Double-click to return the plot to the original size.

Response Surface View

You view the model as a surface by clicking the **Response Surface** icon in the toolbar. You can rotate the plot by click-dragging it.

1. Click `Movie` in the **Display Type** list to see the surface (torque against spark and speed) vary through different values of `load`. Click **Replay** to see it again.

2. Take a look at some of the other display types.

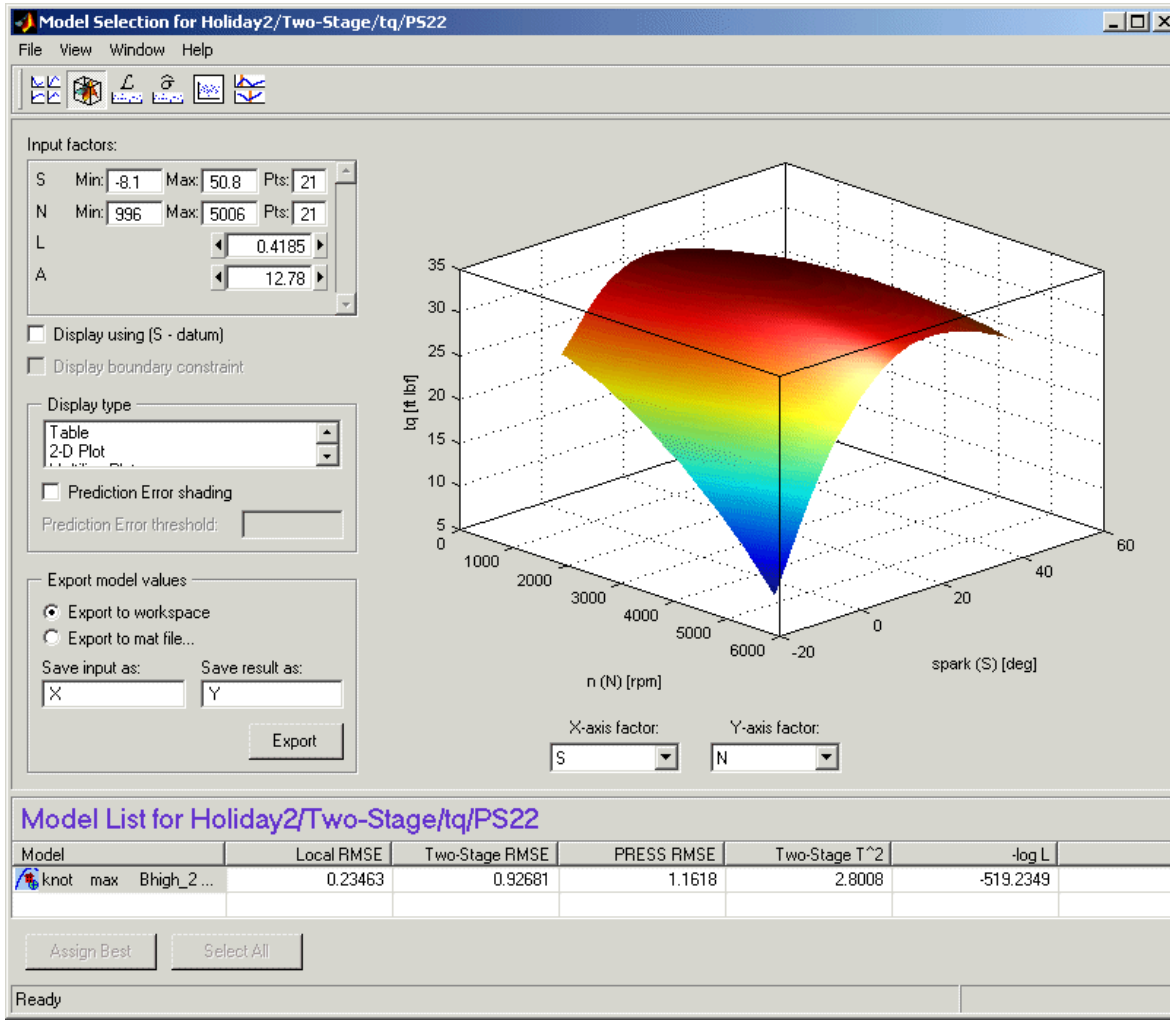

3 Dismiss the Model Selection window, and accept the best model by clicking **Yes** in the Model Selection dialog (it is the only two-stage model so far).

4 The MLE dialog appears, prompting you to calculate the maximum likelihood estimate (MLE) for the two-stage model. Click **Cancel**. You can calculate MLE later.

Comparing the Local Model and the Two-Stage Model

Now the lower plots in the **Local Model** pane show two lines fitted to the test data: the Local Model line (black), and the Two-Stage Model line (green). The plots also show the data (in blue), so you can compare how close the two-stage model fit is to both the data and the local fit for each test.

You can scroll through the various tests (using the arrows at the top left or the **Select Test** button) to compare the local and two-stage models for different tests.

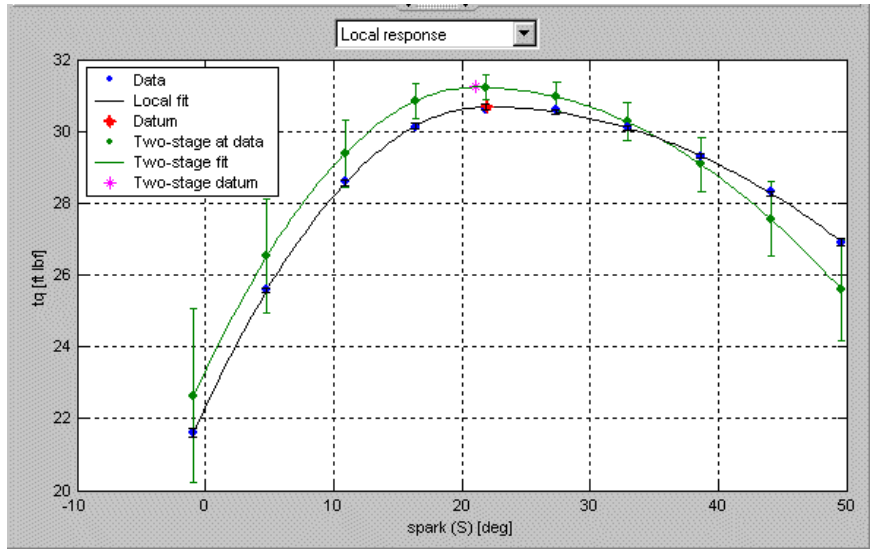

Notice that the local model icon has changed (from the local icon showing a house, to a two-stage icon showing a house and a globe) to indicate that a two-stage model has been calcluated.

Click the

button in the toolbar to calculate the maximum likelihood estimate.

Maximum Likelihood Estimation

The global models were created in isolation without accounting for any correlations between the response features. Using MLE (maximum likelihood estimation) to fit the two-stage model takes account of possible correlations between response features. In cases where such correlations occur, using MLE significantly improves the two-stage model.

1. You reach the MLE dialog from the local node (PS22 in this case) by
   - Clicking the button in the toolbar
   - Or by choosing **Model > Calculate MLE**
2. Leave the algorithm default settings and click **Start** to calculate MLE.
3. Watch the progress indicators until the process finishes and a two-stage RMSE (root mean square error) value appears.
4. Click **OK** to leave the MLE dialog.

   Now the plots on the **Local Model** pane all show the two-stage model in purple to indicate that it is an MLE model. This is also indicated in the legend. Notice that all the model icons in the tree (the response, the local model, and the response features) have also changed to purple to indicate that they are MLE models.

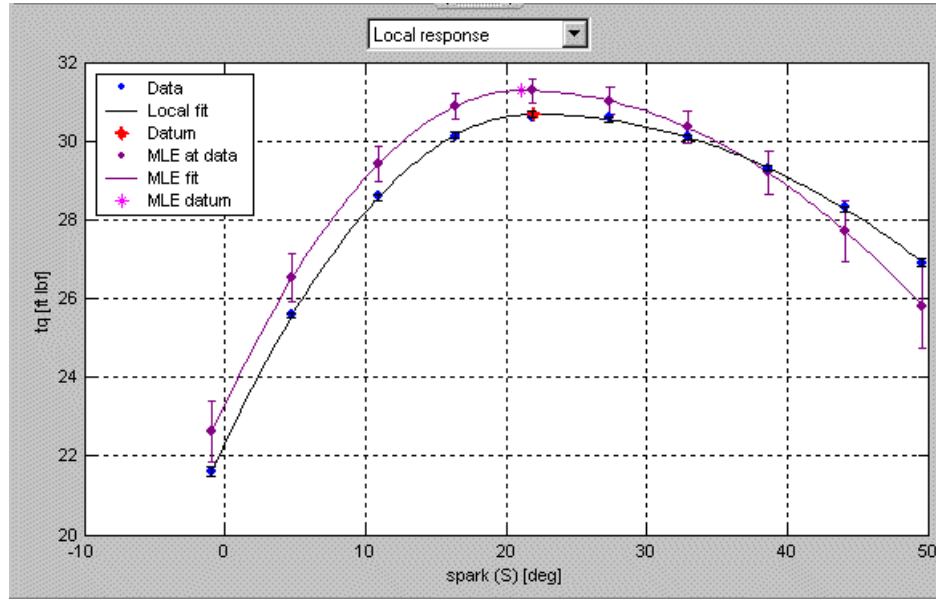

5 Click the **Select** button. This takes you to the Model Selection window.

Here you can compare MLE with the univariate model previously constructed (without correlations). By default, the local fit is plotted against the MLE model.

6 Select both MLE and the Univariate model for plotting by holding down **Shift** while you click the Univariate model in the **Model List** at the bottom of the view.

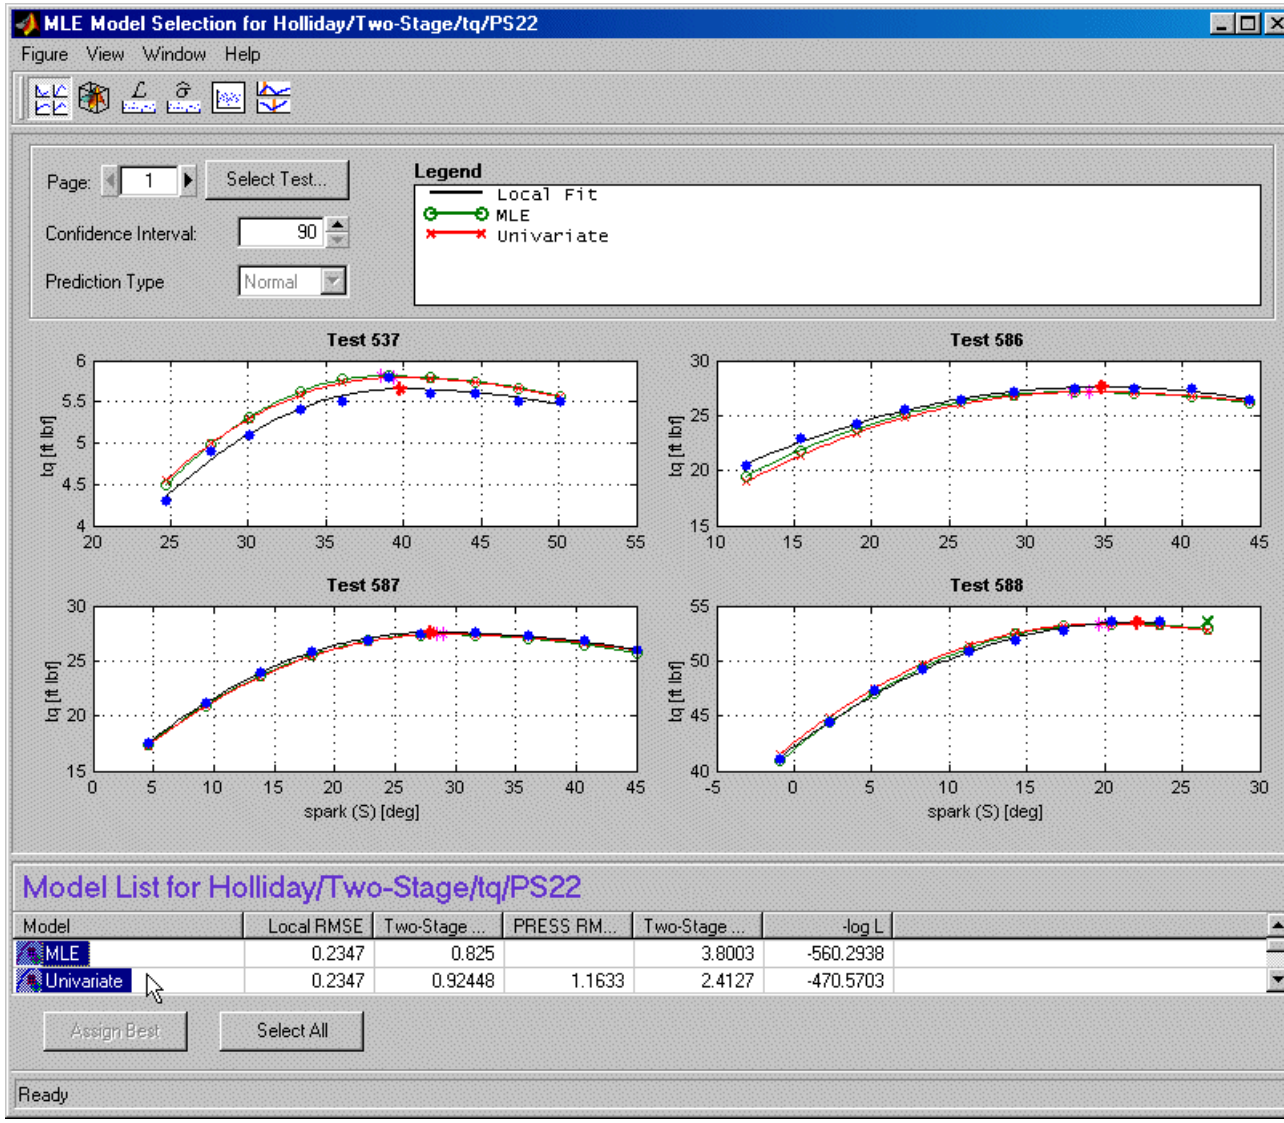

7 Close the Model Selection window. Click **Yes** to accept the MLE model as the best.

Response Node

Click the Response node (`tq`) in the Model Browser tree.

Now at the Response node in the Model Browser tree (`tq`), which was previously blank, you see this:

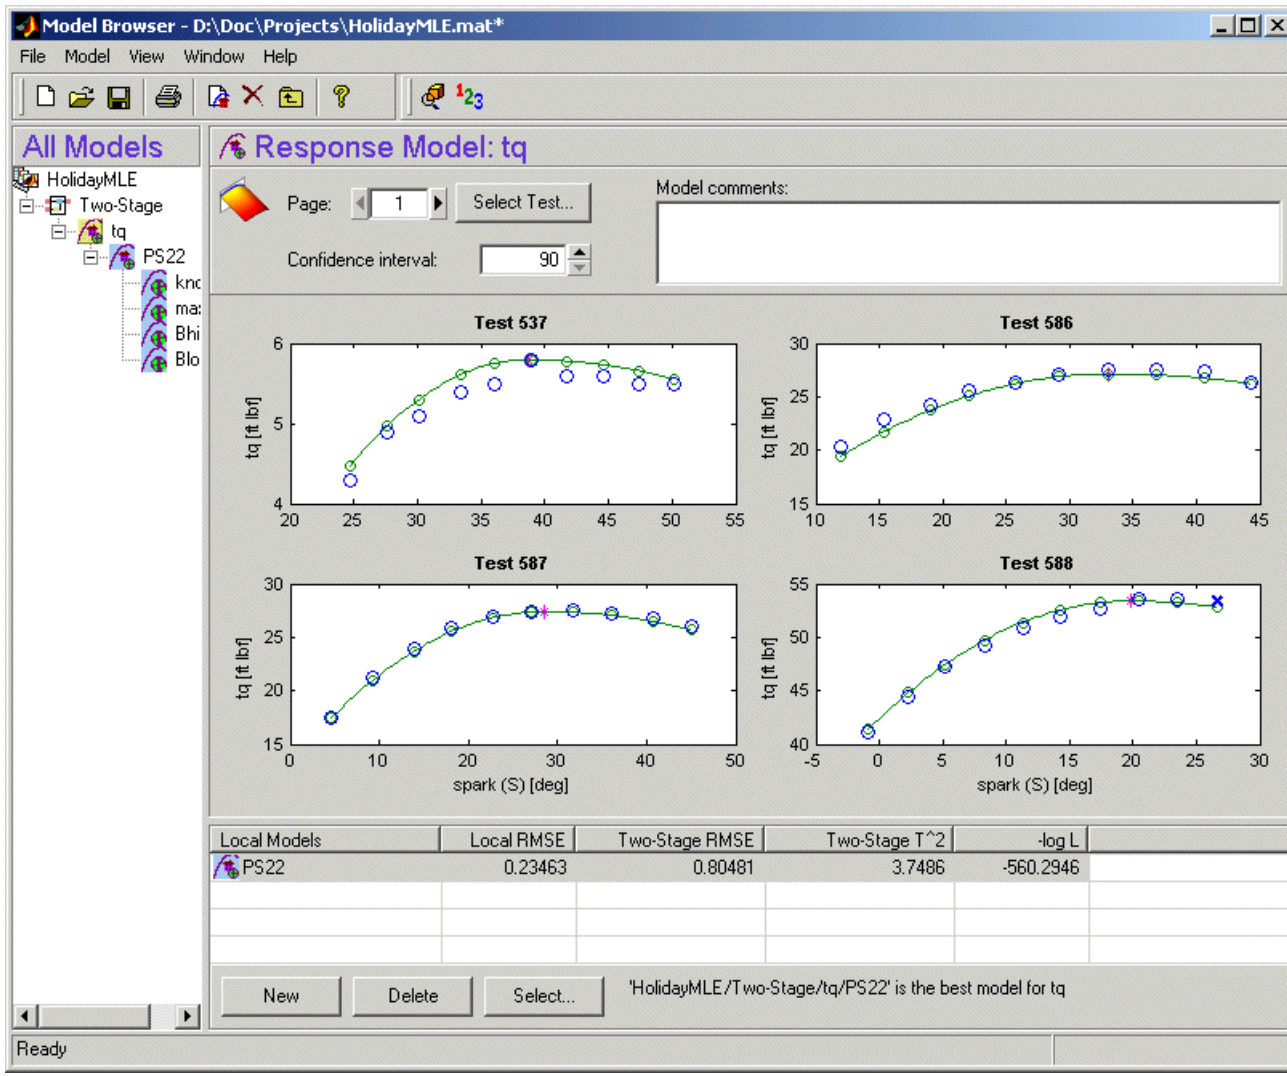

This shows you the fit of the two-stage model to the data. You can scroll through the tests, using the arrows at top left, to view the two-stage MLE model (in green) against the data (in blue) for each test.

You have now completed setting up and verifying a two-stage model.

Exporting the Model

All models created in the Model Browser are exported using the **File** menu. A model can be exported to the MATLAB workspace, to a file, or to a Simulink model.

1. Click the `tq` node in the model tree.
2. Choose **File > Export Models**. The Export Model dialog box appears.
3. Choose `File` from the **Export to** pop-up menu. This saves the work as a file for use within the Model-Based Calibration Toolbox product, for instance, to create calibrations in the CAGE Browser.
4. In the **Export Options** frame, select the destination of the file. You can do this by typing directly in the edit box, or using the Browse button if you want to locate a directory or use an existing file.
5. Ensure that **Export datum models** is selected, as this allows the datum global model to be exported. The datum model in this case is MBT (the spark angle at maximum brake torque).

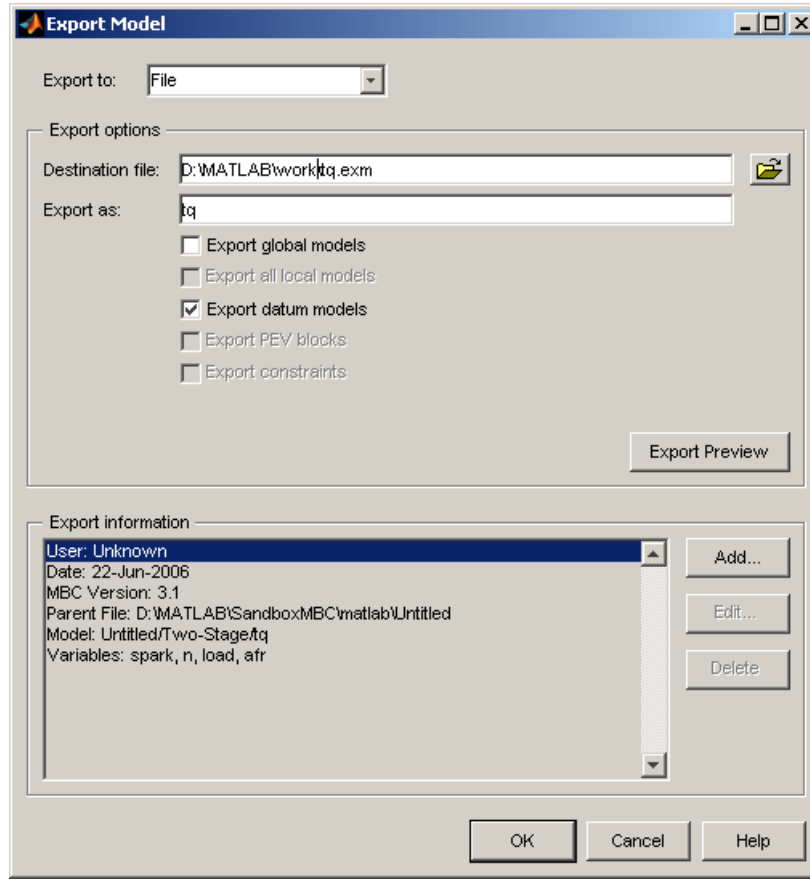

6 Click **OK** to export the models.

Creating Multiple Models to Compare

In this section...

"Methods For Creating More Models" on page 5-43
"Creating New Local Models" on page 5-43
"Adding New Response Features" on page 5-46
"Comparing Models" on page 5-47
"Creating New Global Models" on page 5-49
"Creating Multiple Models Using Build Models" on page 5-52

Methods For Creating More Models

Once you have fitted and examined a single model, you will normally want to create more models to search for the best fit. You can create individual new models, use the **Build Models** function to create a selection of models at once, or create a template to save a variety of model settings for reuse.

You can create new child nodes by clicking the **New** button from any modeling node. The Model Setup dialog appears where you can change the type and settings, and when you close the dialog the view switches to the new child node on the tree. You can do this for multiple child nodes to create a selection of different model types fitted to the same data. You can also use the Build Models dialog to quickly create a selection of different child nodes to compare. The following exercises show you examples of these processes. Note that you need to complete the previous tutorial sections to have a complete two-stage model as a starting point.

Creating New Local Models

1. As an example, select the `tq` response node and click **New** in the **Local Models** list pane.

   The Local Model Setup dialog appears.

2. Select a `Polynomial Spline` with a spline order of `3` below the knot and `2` above. Click **OK**.

A new set of local models (and associated response feature models) is calculated.

3 Return to the parent `tq` response node , and click **New** again, in the **Local Models** list pane.

4 Select a `Polynomial` with an order of `2` in the Local Model Setup dialog. Click **OK**.

A new set of local models and response feature models is calculated.

Now you have three alternative local models to compare: two polynomial splines (order 3,2 and order 2,2) and a polynomial (order 2), as shown.

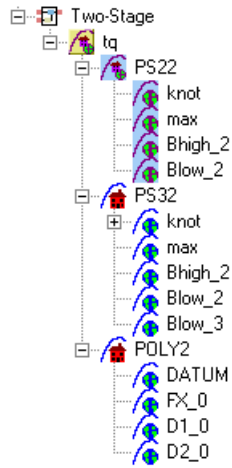

You can select the alternative local models in turn and compare their statistics. For an example, follow these steps:

1 Select the new local model node `PS32`.

2 Select test `587` in the **Test** edit box.

3 In the **Diagnostic statistics** pane, select `Local diagnostics` from the drop-down menu. Observe the value of `s_i` in this pane. This is the value of RMSE (root mean squared error) for the current (i^{th}) test.

The RMSE value is our basic measure of how closely a model fits some data, which measures the average mismatch between each data point and the model. This is why you should look at the RMSE values as your first tool to inspect the quality of the fit — high RMSE values can indicate problems.

4 Now select the local model node `POLY2` and see how the value of `s_i` changes.

Observe that the shape of the torque/spark sweep for this test is better suited to a polynomial spline model than a polynomial model. The curve is not symmetrical because curvature differs above and below the maximum (marked by the red cross at the datum). This explains why the value of `s_i` is much lower for `PS32` (the polynomial spline) than for the `POLY2` (polynomial) for this test. The polynomial spline is a better fit for the current test.

5 Look through some other tests and compare the values of `s_i` for the different local models. To choose the most suitable local model you must decide which fits the majority of tests better, as there are likely to be differences among best fit for different tests.

6 To help you quickly identify which local models have the highest RMSE, indicating problems with the model fit, click RMSE Plots ( ) in the toolbar (or select **View > RMSE Plots**) to open the RMSE Explorer dialog.

a Use the plot to help you identify problem tests. Use the drop-down menus to change the display. For example, select `s_knot` to investigate the error values for knot (MBT), or `s_e` to look at overall error.

b You can navigate to a test of interest from the RMSE Explorer by double-clicking a point in the plot to select the test in the Model Browser local model view.

7 Look at the value of `Local RMSE` reported in the **Pooled Statistics** pane on the right (this is pooled between all tests). Now switch between the `POLY2` and the `PS32` local models again and observe how this value changes.

8 You can compare these values directly by selecting the parent `tq` response node, when the Local RMSE is reported for each child local model in the list at the bottom.

When all child models have a two-stage model calculated, you can also compare two-stage values of RMSE here. Remember, you can always see statistics for the list of child models of the currently selected node in this bottom list pane.

When comparing models, look for lower RMSE values to indicate better fits. However, remember that a model that interpolates between all the points can have an RMSE of zero but be useless for predicting between points. Always use the graphical displays to visually examine model fits and beware of "overfitting" — chasing points at the expense of prediction quality. You will return to the problem of overfitting in a later section when you have two-stage models to compare.

Adding New Response Features

Recall that two-stage models are made up of local models and global models. The global models are fitted to the response features of the local models. The response features available are specific to the type of local model. You can add different response features to see which combination of response features makes the best two-stage model as follows:

1. Select the local model node `PS32`.
2. Click the **New** button under the list of response features.

   A dialog appears with a list of available response features.
3. Select `f(x+datum)` from the list and enter `-10` in the **Value** edit box. Click **OK**.

   A new response feature called `FX_less10` is added under the `PS32` local model. Recall that the datum marks the maximum, in this case maximum torque. The spark angle at maximum torque is referred to as maximum brake torque (MBT). You have defined this response feature (`f(x+datum)`) to measure the value of the model (torque) at (`-10 + MBT`) for each test. It can be useful to use a response feature like this to track a value such as maximum brake torque (MBT) minus 10 degrees of spark angle. This response feature is not an abstract property of a curve, so engineering knowledge can then be applied to increase confidence in the models.

4 Select the local node PS32, and click **Select**. Notice that there are four possible two-stage models this time. This is because you added a sixth response feature. Only five (which must include knot) are required for the two-stage model, so you can see the combinations available and compare them. Note that not all combinations of five response features can completely describe the shape of the curve for the two-stage model, so only the possible alternatives are shown.

5 Close the Model Selection window and click **OK** to accept one of the models as best. Click **Cancel** to avoid calculating MLE.

Notice that the four response features chosen to calculate the two-stage model are highlighted in blue, and the unused response feature is not highlighted, as shown.

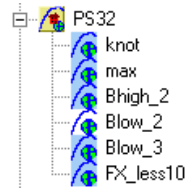

6 Select the tq response node to see a comparison of the statistics of both two-stage models (your original PS22 and the new PS32).

Remember that the POLY2 local model has no two-stage model yet; no two-stage statistics are reported for POLY2 in the bottom list pane. You also cannot use the Model Selection window to fully compare the two-stage models until every local model in the test plan has a two-stage model calculated.

7 To calculate the two-stage model for POLY2, click **Select** at the POLY2 node. Either double-click to assign a model as best or close the Model Selection window and click **OK** to accept the best model. Click **Cancel** to avoid calculating MLE, then the two-stage model is calculated.

Comparing Models

1 Now you have three two-stage models. Select the tq response node and look at the statistics, particularly Local RMSE, Two-Stage RMSE, and PRESS RMSE, reported in the list of child models at the bottom.

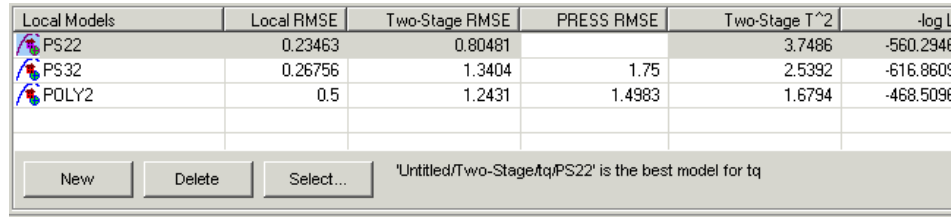

Local Models	Local RMSE	Two-Stage RMSE	PRESS RMSE	Two-Stage T^2	-log L
PS22	0.23463	0.80481		3.7486	-560.2946
PS32	0.26756	1.3404	1.75	2.5392	-616.8609
POLY2	0.5	1.2431	1.4983	1.6794	-468.5096

New Delete Select... 'Untitled/Two-Stage/tq/PS22' is the best model for tq

- Look for lower RMSE values to indicate better fits.
- Look for lower PRESS RMSE values to indicate better fits without overfitting. PRESS RMSE is a measure of the predictive power of your models.

  It is useful to compare PRESS RMSE with RMSE as this may indicate problems with overfitting. RMSE is minimized when the model gets close to each data point; "chasing" the data will therefore improve RMSE. However, chasing the data can sometimes lead to strong oscillations in the model between the data points; this behavior can give good values of RMSE but is not representative of the data and will not give reliable prediction values where you do not already have data. The PRESS RMSE statistic guards against this by testing how well the current model would predict each of the points in the data set (in turn) if they were not included in the regression. To get a small PRESS RMSE usually indicates that the model is not overly sensitive to any single data point.

  If the value of PRESS RMSE is much bigger than the RMSE, then you are overfitting - the model is unnecessarily complex. For a fuller description of the meaning of overfitting, and how RMSE and PRESS can help you select good models, see "Model Selection Guide". As a rule of thumb, if you have about 100 data points, you should aim for a PRESS RMSE no more than 5% larger than the RMSE (remember here you have only 27 tests).

  Notice that your first two-stage model (`PS22`) does not have a PRESS RMSE value. This is because it cannot be displayed for MLE models. You need non-MLE models to be able to use PRESS for direct comparison.
- Look for lower T^2 values. A large T^2 value indicates that there is a problem with the response feature models.
- Look for large negative log likelihood values to indicate better fits.

  See "Pooled Statistics" for more on T^2 and log likelihood.

2 Now click **Select** to open Model Selection to compare all three two-stage models simultaneously. Here you can see the same statistics to compare the models in the bottom list, but you can also make use of a variety of views to look for the best fit:

   - You can plot the models simultaneously on the Tests, Residuals and Cross Section views (**Shift**- or **Ctrl**-click to select models in the list)
   - You can view each model in the Response Surface view as a surface; movie, contour or multiline plot, and as a table

3 You can select a model and click **Assign Best** in the Model Selection window, or double-click a model to assign it as best.

4 When you close the Model Selection window and return to the Model Browser, the model you selected as best is copied to the parent response node, `tq`.

Creating New Global Models

In this example, you have not yet searched for the best global model types. You would normally do this before creating and comparing two-stage models. For the purpose of this tutorial, you have already created two-stage models and used the Model Selection tool to introduce the use of RMSE and PRESS to help you identify better models. The principle is the same at each level in the model tree: add new child models and use the Model Selection window to choose the best.

1 Select one of the response feature nodes under the `PS32` node, such as `knot`.

2 Click **New**. Click **OK** in the Model Setup dialog without changing any settings, to create a copy of the parent model. Return to the parent model and repeat.

   Two new global model child nodes now appear underneath `knot`, as shown. Both are labeled `Quadratic`, as they are currently copies of the parent model. You can create any number of child nodes to search for the best global model fit for each response feature in your tree. When you choose the best, it is copied to the parent node, in this case `knot`, including any outliers you decide to exclude.

A good technique for creating multiple models can be to leave the first child node unchanged, then you always have a copy of the original model for comparison.

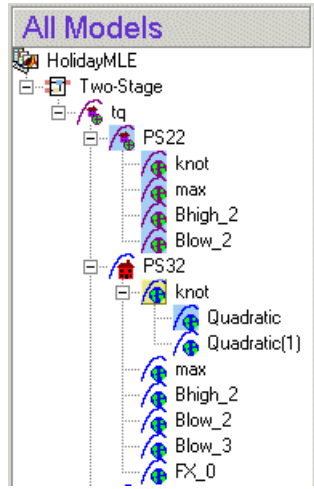

3 Select one of the new `Quadratic` nodes, then select the menu item **Model > Set Up**.

The Global Model Setup dialog appears. Here you can change the type and settings of the model to see if you can find a better fit to the data with a different model type.

4 Use the drop-down menu to change the **Model class** to `Hybrid RBF` and click **OK**.

The new model fit is calculated, and the `Quadratic` node's name changes to `Linear-RBF`.

5 Select the remaining `Quadratic` node, then select **Model > Set Up**.

6 Use the drop-down menu to change the **Model class** to `Radial Basis Function`, and click **OK**. There are many other settings you can alter for both these model types, but for a quick exploration of the trends in the data it is worth trying the default model settings.

The new model fit is calculated and the `Quadratic` node's name changes to `RBF-multiquadric`.

7 To compare the two child node models, select the parent node `knot` and click **Select**. Whichever model you assign as best is copied to the `knot` node when you close the Model Selection window and click **OK**.

Notice that the child node model assigned as best is highlighted in blue, and the local node has changed from the two-stage icon back to the local model icon (a red house) as shown. This is because you have changed one of the response feature models, and so you need to recalculate the two-stage model using the new global model for this response feature. First you need to select best global models for every response feature.

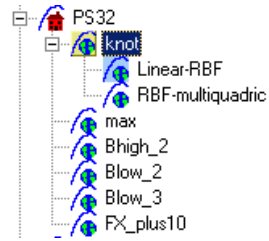

8 Add two more child nodes to the `knot` global model (select `knot`, then click **New**, and repeat).

Notice that now the new nodes are copies of Linear-`RBF`, because that model was selected as best.

9 Select the two new nodes in turn and change their model types. Try a cubic and quadratic polynomial model type.

a Select the menu item **Model > Set Up**.

b Choose `Linear model` from the **Model class** drop-down menu and set the polynomial order for each factor to `3` for one model, then `2` for the other. Click **OK**.

10 To compare all four child node models, select the parent node `knot` and click **Select**. `Linear-RBF` still performs the best for `PRESS RMSE`. Whichever model you assign as best is copied to the `knot` node.

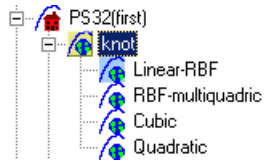

11 Select the knot model node, then select **Model > Make Template**. Browse to a suitable work directory and enter the name `Mytemplate`. Click **OK**.

Creating Multiple Models Using Build Models

The quickest way to create multiple different models to compare is to use the Build Models function. You can use this to select a template and build a selection of models as child nodes of the current node. The best model of this selection of child nodes is automatically selected (it will have a blue icon), based on the selection criteria you choose (such as PRESS RMSE, RMSE, Box-Cox, and so on).

1 Before calculating MLE, select a global model such as `max`.

   You cannot reach the Build Models dialog from an MLE global model. Note that calculating MLE is not irreversible — to go back you can always go to Model Selection (from the local node) and select the `Univariate` model as best.

2 Click Build Models in the toolbar.

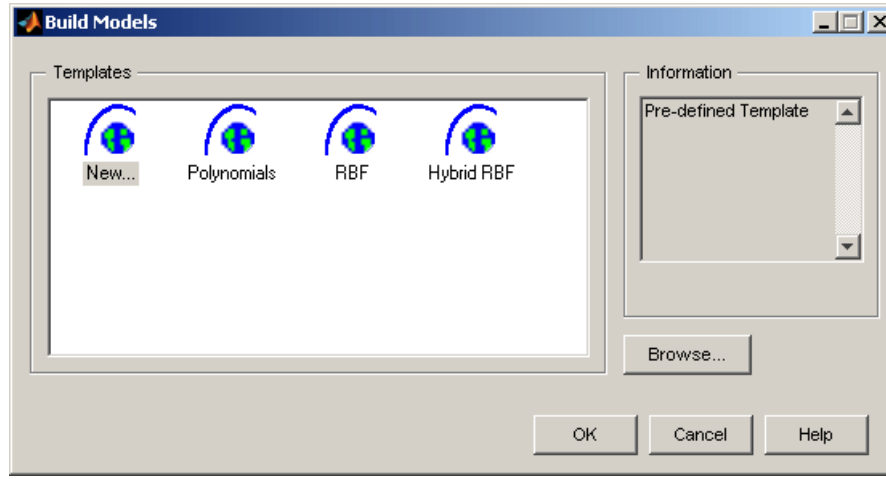

The Build Models dialog appears. Here you can choose a template for the type of models you want to build. There are predefined templates for polynomials, RBFs, hybrid RBFs, free knot splines (for single input models), or you can select a suitable parent node in the current project to use as a template.

You can also create templates of whatever models you choose by selecting the `New` template in the Build Models dialog box, or using the **Model > Make Template** menu item at a model node, as you did in the previous section. Your user-defined templates can then be found via the Build Models dialog. You can use the Browse button to find stored templates that are not in the default directory.

3 Click Browse and select the directory containing the template you created earlier, named `Mytemplate`. Click **OK**.

Your new template (called `Mytemplate`) now appears in the Build Models dialog along with the defaults. Note that you can set the default directory where the toolbox looks for templates (and models, data, and projects) using **File > Preferences**.

4 Select `Mytemplate`. Notice the four model types you saved in this template appear in the **Information** pane. Click **OK**.

The Model Selection dialog appears, where you can choose a criterion for automatically selecting the best model out of the child nodes.

5 Use the drop-down menu to choose `PRESS RMSE` as the selection criteria for the best model, and click **OK**.

Four child nodes appear: `Linear-RBF`, `RBF-multiquadric`, `Cubic`, and `Quadratic`. These are the model types you selected when you built the template and are now fitted to the data for `max`.

The most favorable child node model, based on `PRESS RMSE`, is selected as best (highlighted in blue) as shown in the following figure. This model is also copied to the parent node `max`, in the same way as if you had used the Model Selection window to assign a best model.

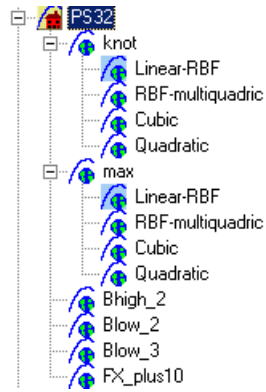

Try one of the default templates in the Build Models dialog box as follows:

1 Select another global model such as `Blow_2`.

2 Click Build Models in the toolbar.

3 Select `RBF` and click **OK**. Click **Build** in the following Model Building Options dialog to build a selection of child nodes.

Similarly, you can use the Build Models dialog to automatically build a selection of polynomial or hybrid RBF models, or your own selection of model types, to search for the best fit.

4. Click **OK** in the Model Selection dialog to accept the default, `PRESS RMSE`, as the selection criteria for the best model.

5. Look at the statistics in the lower list pane to quickly compare all the different RBF kernel child models. If one model performs significantly better in terms of PRESS RMSE and RMSE you might choose not to click **Select** to compare all the child node models. However, it is usually useful to visually inspect the models to see how they compare.

6. When you have chosen a best model, it can be useful to select some (or all) of the rejected models in the bottom list pane and press **Delete**. You can also select **File > Clean Up Tree**. This deletes all rejected child models where best models have been chosen; only the child nodes selected as best remain.

Creating a template containing a list of all the models you want is a very efficient way to quickly build a selection of alternative model child nodes for many global models. Use these techniques to find models well suited to the data for each of your global models.

When you have chosen best global models for all your response features, you need to recalculate the two-stage model. Click **Select** at the local model (`PS32`) node to calculate the two-stage model.

6

Tutorial: Design of Experiment

This section discusses the following topics:

- "What Is Design of Experiment?" on page 6-2
- "Getting Started with the Design Editor" on page 6-6
- "Creating Optimal Designs" on page 6-10
- "Viewing Design Displays" on page 6-17
- "Using the Prediction Error Variance Viewer" on page 6-21
- "Creating Classical Designs" on page 6-27
- "Using the Design Evaluation Tool" on page 6-34
- "Creating Space-Filling Designs" on page 6-37
- "Applying Constraints" on page 6-39
- "Saving Designs" on page 6-45

What Is Design of Experiment?

In this section...

"Why Use Design of Experiment?" on page 6-2

"Design Styles" on page 6-3

"Structure of This Design Editor Tutorial" on page 6-4

Why Use Design of Experiment?

With today's ever-increasing complexity of models, design of experiment has become an essential part of the modeling process. The Design Editor within the Model-Based Calibration Toolbox product is crucial for the efficient collection of engine data. Dyno-cell time is expensive, and the savings in time and money can be considerable when a careful experimental design takes only the most useful data. Dramatically reducing test time is growing more and more important as the number of controllable variables in more complex engines is growing. With increasing engine complexity, the test time increases exponentially.

The traditional method of collecting large quantities of data by holding each factor constant in turn until all possibilities have been tested is an approach that quickly becomes impossible as the number of factors increases. A full factorial design (that is, testing for torque at every combination of speed, load, air/fuel ratio, and exhaust gas recirculation on a direct injection gasoline engine with stratified combustion capability) is not feasible for newer engines. Simple calculation estimates that, for recently developed engines, to calibrate in the traditional way would take 99 years!

With a five-factor experiment including a multiknot spline dimension and 20 levels in each factor, the number of points in a full factorial design quickly becomes thousands, making the experiment prohibitively expensive to run. The Design Editor solves this problem by choosing a set of experimental points that allow estimation of the model with the maximum confidence using just a fraction of the number of experimental runs; for the preceding example just 100 optimally chosen runs is more than enough to fit the model. Obviously, this approach can be advantageous for any complex experimental design, not just engine research.

The Design Editor offers a systematic, rigorous approach to the data collection stage. When you plan a sequence of tests to be run on an example engine, you can base your design on engineering expertise and existing physical and analytical models. During testing, you can compare your design with the latest data and optimize the remaining tests to get maximum benefit.

The Design Editor provides prebuilt standard designs to allow a user with a minimal knowledge of the subject to quickly create experiments. You can apply engineering knowledge to define variable ranges and apply constraints to exclude impractical points. You can increase modeling sophistication by altering optimality criteria, forcing or removing specific design points, and optimally augmenting existing designs with additional points.

Design Styles

The Design Editor provides the interface for building experimental designs. You can make three different styles of design: classical, space-filling, and optimal.

Optimal designs are best for cases with high system knowledge, where previous studies have given confidence on the best type of model to be fitted, and the constraints of the system are well understood. See “Creating Optimal Designs” on page 6-10.

Space-filling designs are better when there is low system knowledge. In cases where you are not sure what type of model is appropriate, and the constraints are uncertain, space-filling designs collect data in such as a way as to maximize coverage of the factors’ ranges as quickly as possible. See “Creating Space-Filling Designs” on page 6-37.

Classical designs (including full factorial) are very well researched and are suitable for simple regions (hypercube or sphere). Engines have complex constraints and models (high-order polynomials and splines). See “Creating Classical Designs” on page 6-27.

You can augment any design by optimally adding points. Working in this way allows new experiments to enhance the original, rather than simply being a second attempt to gain the necessary knowledge.

Structure of This Design Editor Tutorial

The following sections guide you through constructing optimal, classical, and space-filling designs; how to compare designs using the prediction error variance (PEV) viewer and Design Evaluation tool; and how to apply constraints to designs.

1. To start the tutorial, you pick a model to design an experiment for, enter the Design Editor, and construct an optimal design. Once you create a design, you can use the displays and tools to examine the properties of the design, save the design, and make changes.

   See

   - "Getting Started with the Design Editor" on page 6-6
   - "Creating Optimal Designs" on page 6-10
   - "Viewing Design Displays" on page 6-17
   - "Using the Prediction Error Variance Viewer" on page 6-21
   - "Saving Designs" on page 6-45
   - "Improving the Design" on page 6-24

2. Next you create a classical design, and use the Prediction Error Variance Viewer to compare it with the previous design. You can also use the Design Evaluation tool to view all details of any design; it is introduced in this example.

   See

   - "Creating Classical Designs" on page 6-27
   - "Using the Design Evaluation Tool" on page 6-34

3. Lastly you construct a space-filling design and compare it with the others using the Prediction Error Variance Viewer. Then you construct and apply two different constraints to this design and view the results. Normally you would design constraints before constructing a design, but for the purposes of this tutorial you make constraints last so you can view the effects on your design.

   See

- "Creating Space-Filling Designs" on page 6-37
- "Applying Constraints" on page 6-39

For more details on functionality in the Design Editor, see the reference section "Designs".

Getting Started with the Design Editor

In this section...

"Setting Up a Model" on page 6-6

"Starting the Design Editor" on page 6-7

"Creating a New Design" on page 6-8

Setting Up a Model

You must first have a model for which to design an experiment.

1. From the Model Browser at startup, click the button in the toolbar, or click **New** in the **Test Plans** pane, or choose **File > New Test Plan**.
2. Select `Two-Stage Model` and click **OK**.
3. Click the new `Two-Stage` node that appears in the model tree (in the **All Models** pane), or double-click `Two Stage` in the **Test Plans** list at the bottom. The Two-Stage Model diagram appears.

If you already have a project open, you can select any existing model within the test plans in the Model Browser tree. For the purposes of this tutorial, you design experiments for the default Two-Stage global model, which is a quadratic.

There is only one input to the global model by default. To increase the number of input factors:

1. Double-click the Global Model Inputs block in the diagram. The Input Factors Setup dialog appears.
2. Increase the number of factors to three by clicking the **Number of Factors** up/down buttons or entering 3 in the edit box.
3. Change the symbols of the three input factors to `N`, `L`, and `A`. This matches the global factors modeled in the Quick Start tutorial: speed (n), load (L), and air/fuel ratio (A).

4 Click **OK** to leave the Input Factor Setup dialog.

Starting the Design Editor

To access the Design Editor, use either of the following methods:

- Right-click the global model in the diagram and choose **Design Experiment**, as shown.
- You can also access the Design Editor by selecting the menu item **TestPlan > Design Experiment**.

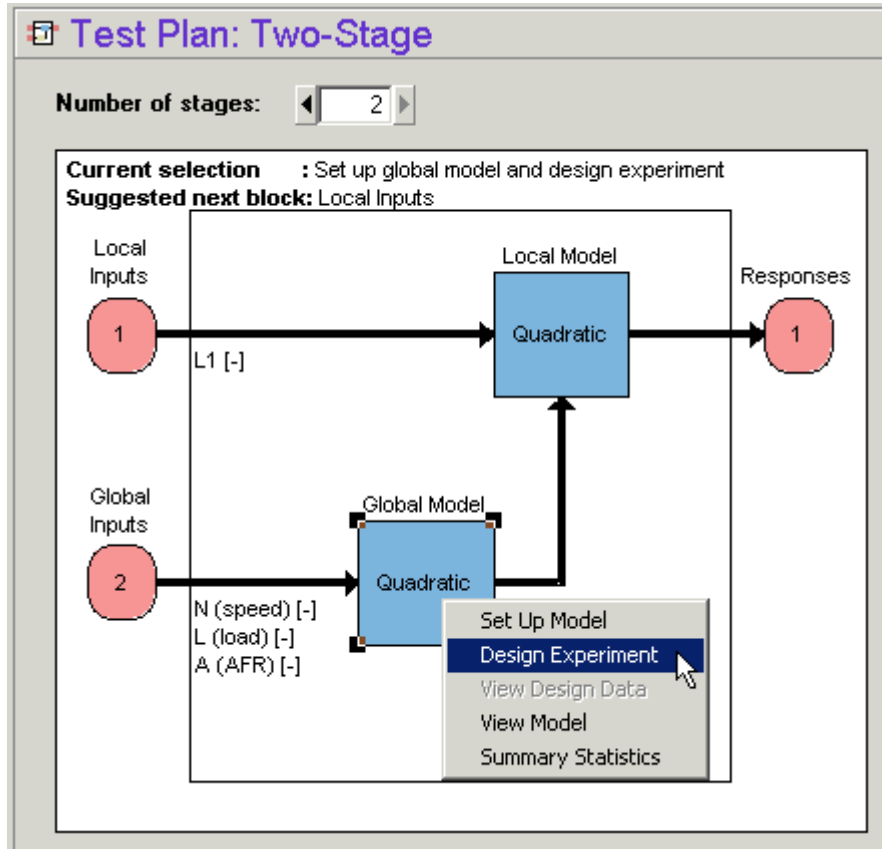

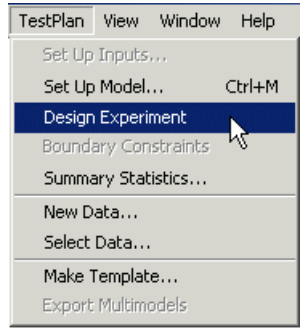

The Design Editor window appears.

Creating a New Design

1 Click the button in the toolbar or select **File > New**. A new node called `Linear Model Design` appears.

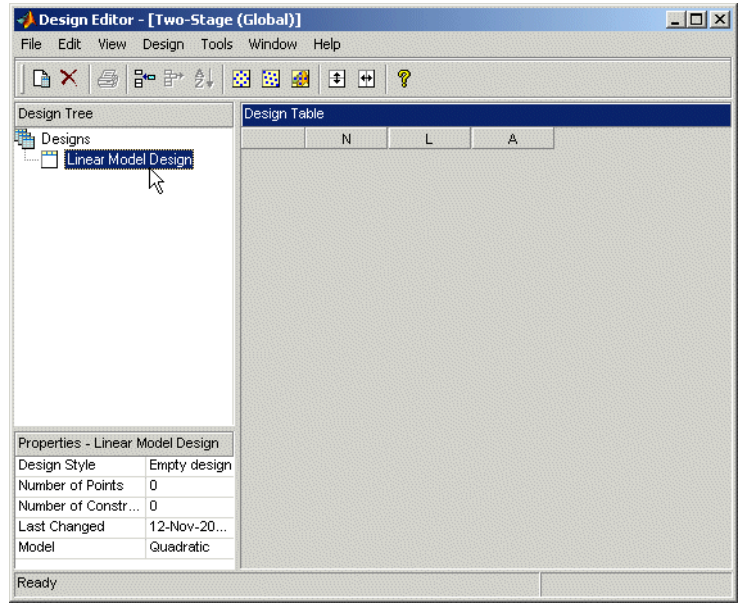

2. The new `Linear Model Design` node is automatically selected. An empty Design Table appears (see above) because you have not yet chosen a design. For this example you create an optimal design for the default global model, which is a quadratic.

   You can change the model for which you are designing an experiment from within the Design Editor window by selecting **Edit > Model**.

3. Rename the new node `Optimal` (you can edit the names by clicking again on a node when it is already selected, or by pressing **F2**, as when selecting to rename in Windows Explorer).

Creating Optimal Designs

In this section...
“Introducing Optimal Designs” on page 6-10
“Start Point Tab” on page 6-11
“Candidate Set Tab” on page 6-12
“Algorithm Tab” on page 6-15

Introducing Optimal Designs

Choose an optimal design by clicking the button in the toolbar, or choose **Design > Optimal**.

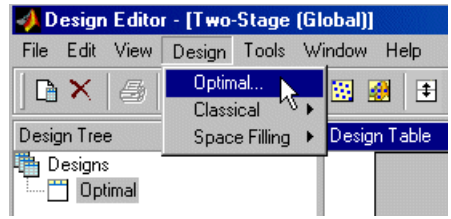

Optimal designs are best for cases with high system knowledge, where previous studies have given confidence on the best type of model to be fitted, and the constraints of the system are well understood.

The optimal designs in the Design Editor are formed using the following process:

- An initial starting design is chosen at random from a set of defined candidate points.
- `m` additional points are added to the design, either optimally or at random. These points are chosen from the candidate set.

- `m` points are deleted from the design, either optimally or at random.
- If the resulting design is better than the original, it is kept.

This process is repeated until either (a) the maximum number of iterations is exceeded or (b) a certain number of iterations has occurred without an appreciable change in the optimality value for the design.

The Optimal Design dialog consists of several tabs that contain the settings for three main aspects of the design:

- Starting point and number of points in the design
- Candidate set of points from which the design points are chosen
- Options for the algorithm that is used to generate the points

Start Point Tab

The **Start Point** tab allows you to define the composition of the initial design: how many points to keep from the current design and how many extra to choose from the candidate set.

1. Leave the optimality criteria at the default to create a `V-Optimal` design.
2. Increase the total number of points to 30 by clicking the **Optional additional points** up/down buttons or by typing directly into the edit box. You can edit the additional points and/or the total number of points.

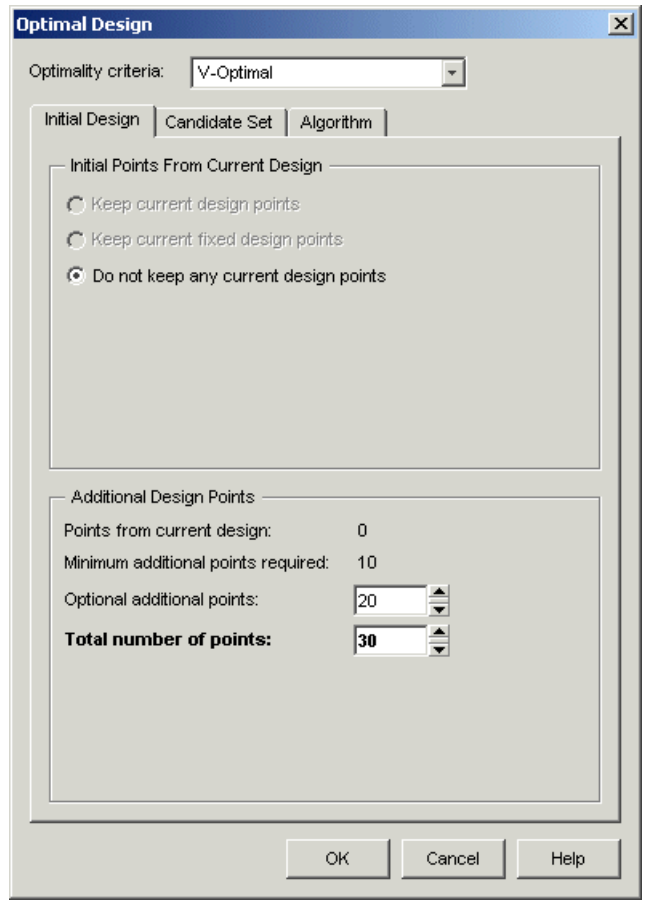

Candidate Set Tab

The **Candidate Set** tab allows you to set up a candidate set of potential test points. This typically ranges from a few hundred points to several hundred thousand.

1. Choose `Grid` for this example. Note that you could choose different schemes for different factors.

2. This tab also has buttons for creating plots of the candidate sets. Try them to preview the grid.

3 Notice that you can see 1-D, 2-D, 3-D, and 4-D displays (the fourth factor is color, but this example only uses three factors) at the same time as they appear in separate windows (see example following). Look at a display window while changing the number of levels for the different factors. See the effects of changing the number of levels on different factors, then return them all to the default of 21 levels.

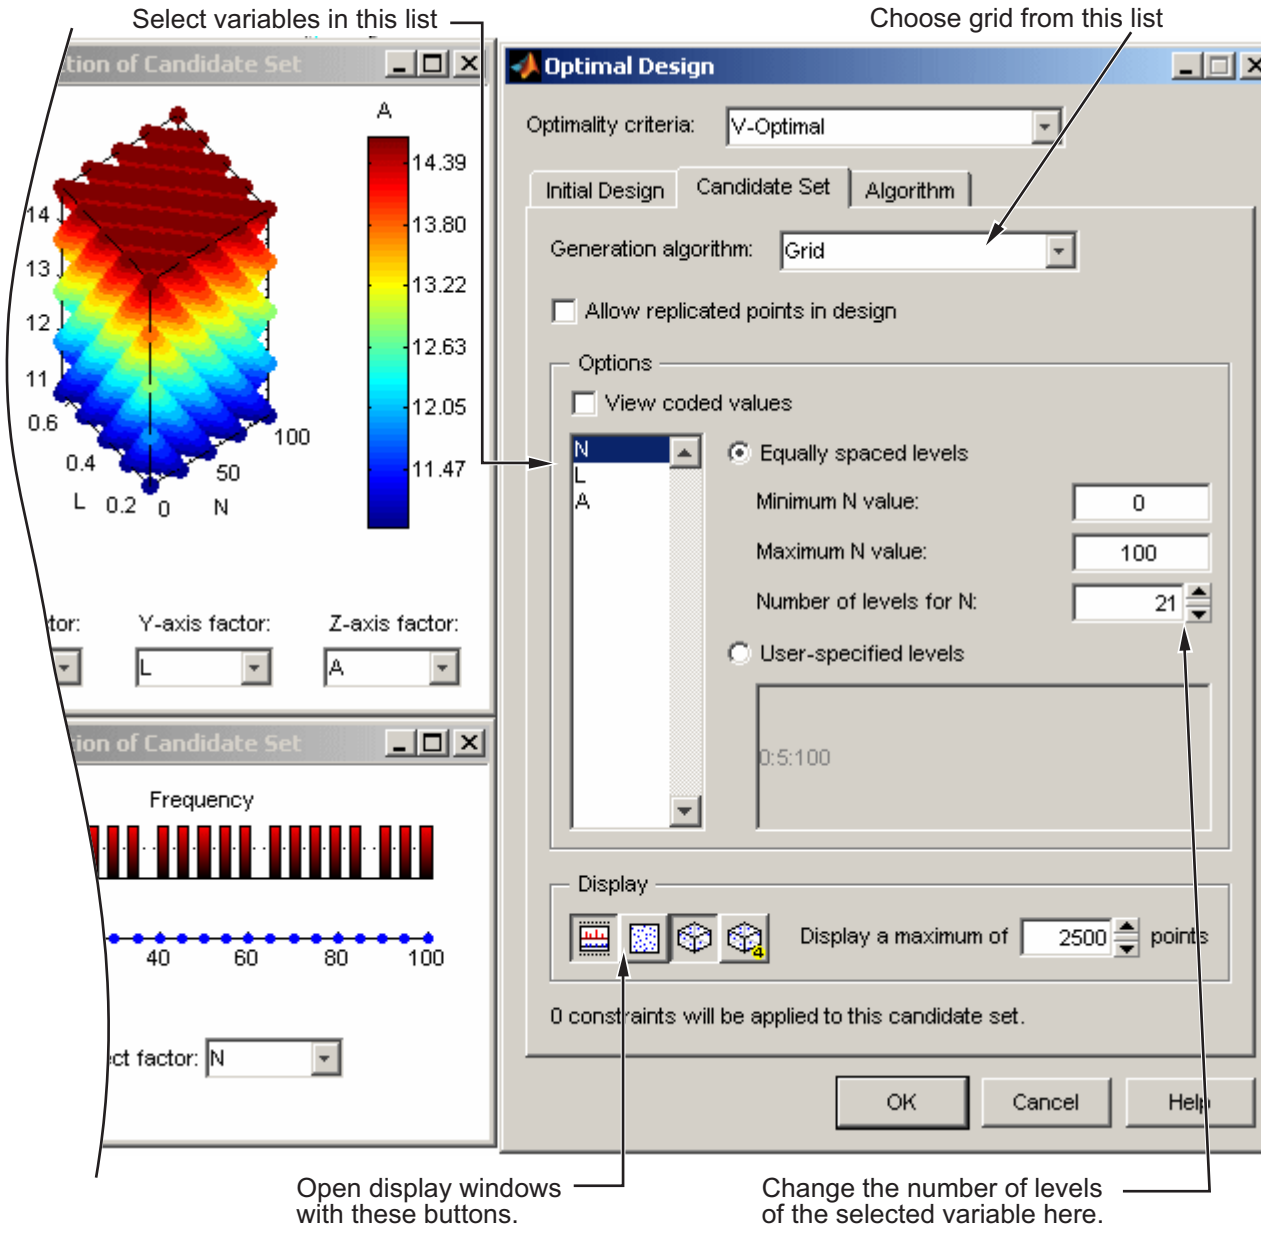
Select variables in this list
Choose grid from this list
Optimal Design
Optimality criteria:
V-Optimal
Initial Design
Candidate Set
Algorithm
Generation algorithm:
Grid
Allow replicated points in design
Options
View coded values
N
L
A
Equally spaced levels
Minimum N value:
0
Maximum N value:
100
Number of levels for N:
21
User-specified levels
0:5:100
Display
Display a maximum of
2500
points
0 constraints will be applied to this candidate set.
OK
Cancel
Help
Y-axis factor:
L
Z-axis factor:
A
Frequency
40
60
80
100
N
Open display windows with these buttons.
Change the number of levels of the selected variable here.

Algorithm Tab

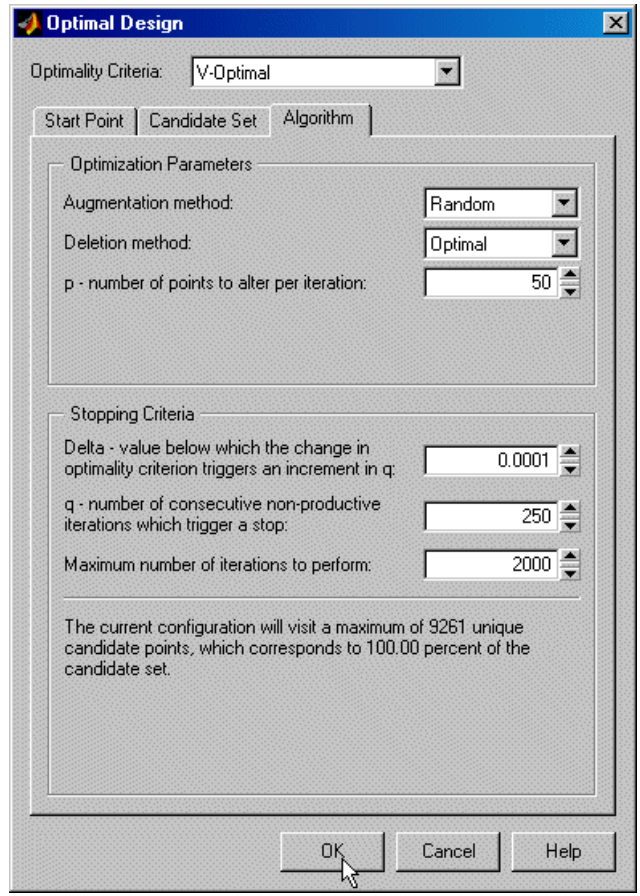

1 Leave the algorithm settings at the defaults and click **OK** to start optimizing the design.

   When you click the **OK** button on the Optimal Design dialog, the Optimizing Design dialog appears, containing a graph. This dialog shows the progress of the optimization and has two buttons: **Accept** and **Cancel**. **Accept** stops the optimization early and takes the current design from it. **Cancel** stops the optimization and reverts to the original design.

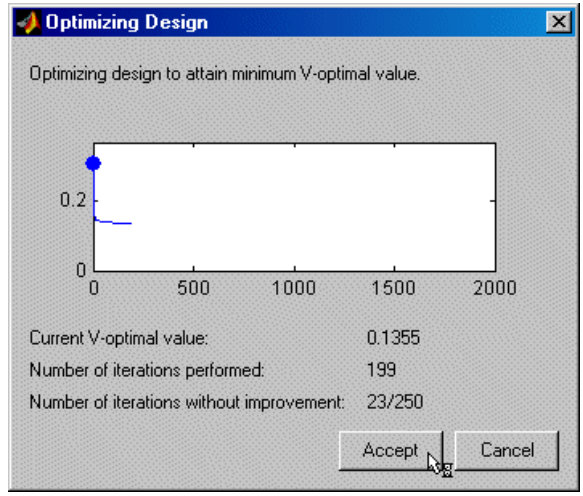

2 Click **Accept** when iterations are not producing noticeable improvements; that is, the graph becomes very flat.

Viewing Design Displays

When you press the **Accept** button, you return to the Design Editor.

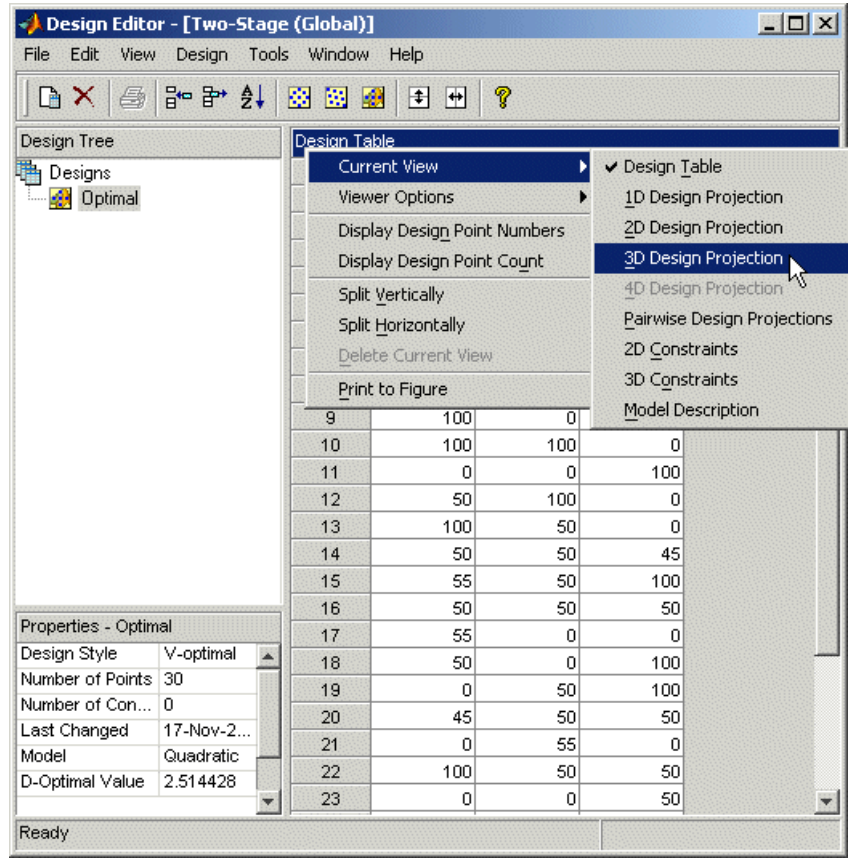

When you first see the main display area, it shows the default **Design Table** view of the design (see preceding example). There is a context menu, available by right-clicking on the title bar, in which you can change the view of the design to 1-D, 2-D, 3-D, 4-D, and pairwise design projections, 2-D and 3-D constraint views, and the table view (also under the **View** menu). This menu also allows you to split the display either horizontally or vertically so that you simultaneously have two different views on the current design. You can also use the toolbar buttons to do this. The split can be merged again. After

splitting, each view has the same functionality; that is, you can continue to split views until you have as many as you want. When you click a view, its title bar becomes blue to show it is the active view.

The currently available designs are displayed on the left in a tree structure. For details, see “The Design Tree”.

Display Options

The Design Editor can display multiple design views at once, so while working on a design you can keep a table of design points open in one corner of the window, a 3-D projection of the constraints below it and a 2-D or 3-D plot of the current design points as the main plot. The following example shows several views in use at once.

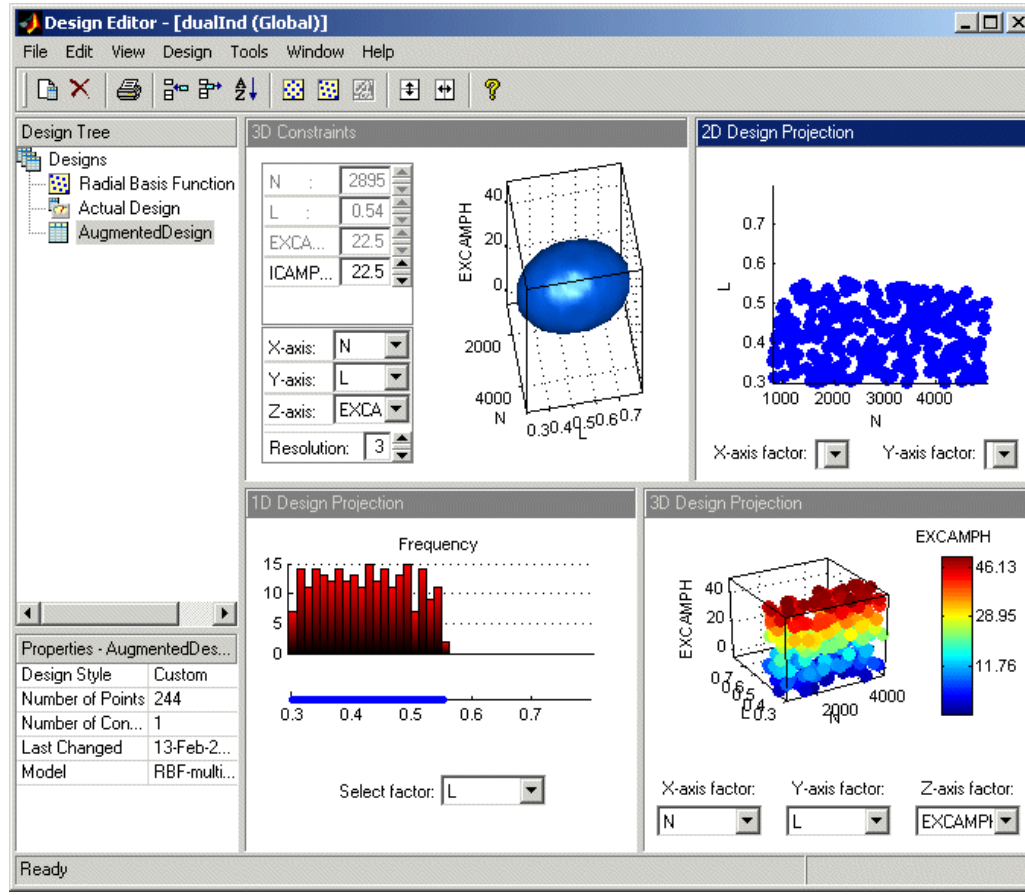

The current view and options for the current view are available either through the context menu or the **View** menu on the Design Editor window.

1 Change the main display to 3-D Projection view.

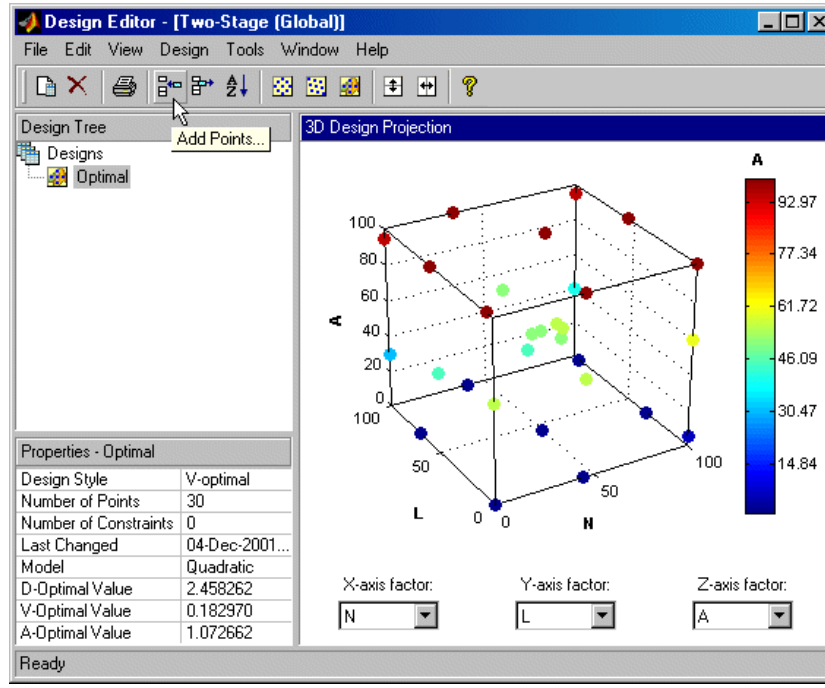

2 You can rotate the projection with click-drag mouse movement. View your design in several projections (singly, or simultaneously by dividing the pane) by using the right-click context menu in the display pane.

Using the Prediction Error Variance Viewer

In this section...

"Introducing the Prediction Error Variance Viewer" on page 6-21

"Improving the Design" on page 6-24

Introducing the Prediction Error Variance Viewer

A useful measure of the quality of a design is its prediction error variance (PEV). The PEV hypersurface is an indicator of how capable the design is in estimating the response in the underlying model. A bad design is either not able to fit the chosen model or is very poor at predicting the response. The Prediction Error Variance Viewer is only available for linear models. The Prediction Error Variance Viewer is not available when designs are rank deficient; that is, they do not contain enough points to fit the model. Optimal designs attempt to minimize the average PEV over the design region.

Select **Tools > Prediction Error Variance Viewer**.

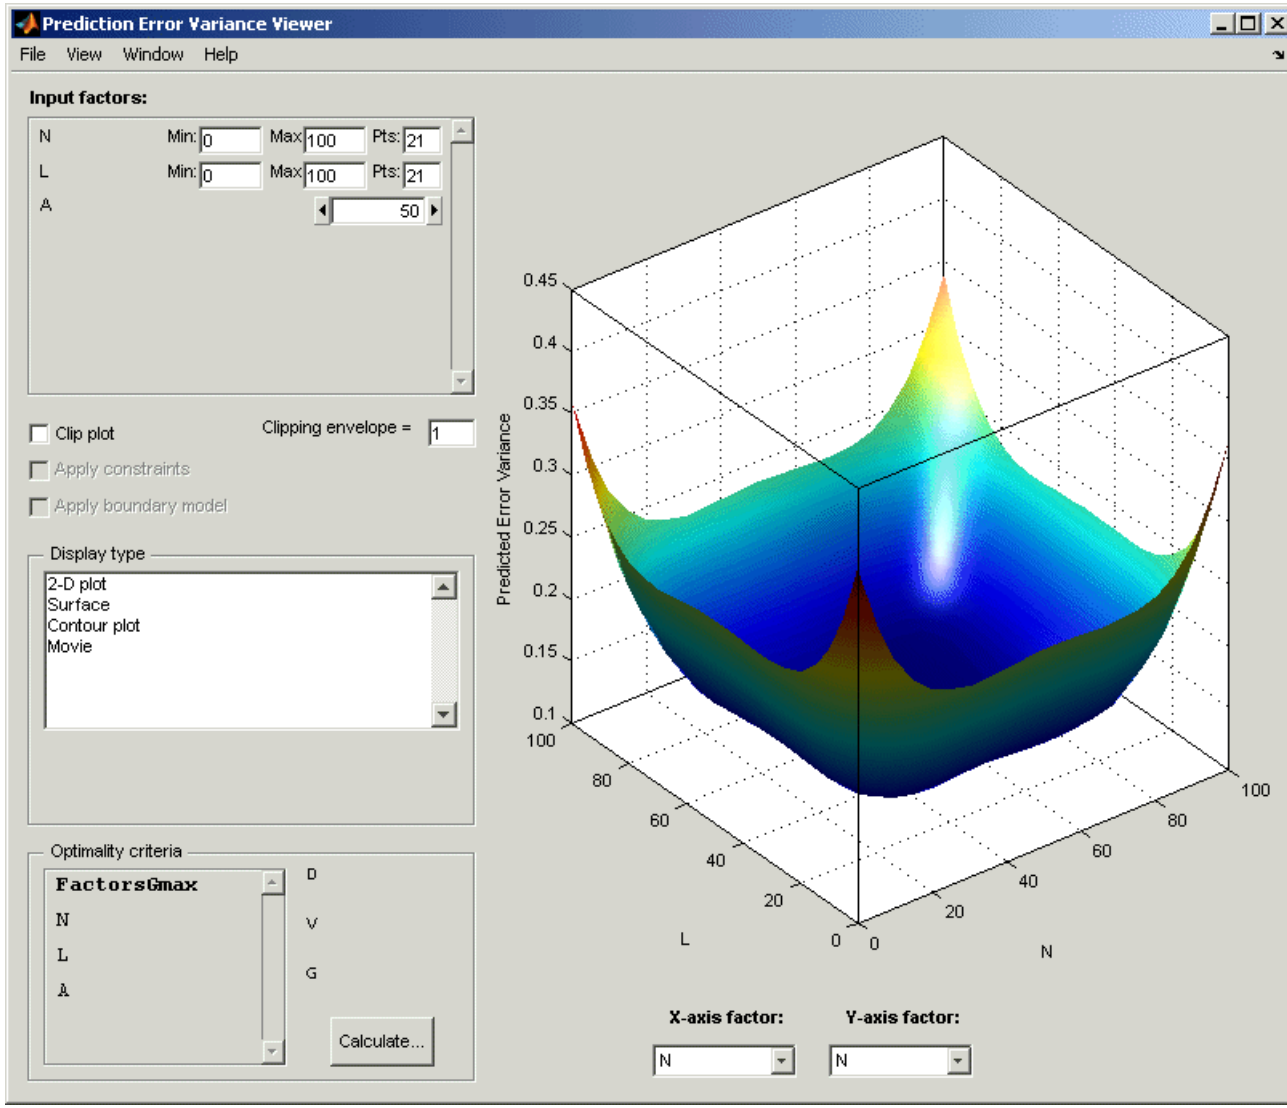

The default view is a 3-D plot of the PEV surface.

This shows where the response predictions are best. This example optimal design predicts well in the center and the middle of the faces (one factor high

and the other midrange), but in the corners the design has the highest error. Look at the scale to see how much difference there is between the areas of higher and lower error. For the best predictive power, you want low PEV (close to zero).

You can examine PEV for designs and models. The two are related in this way:

Accuracy of model predictions (model PEV)=Design PEV * MSE (Mean Square Error in measurements).

You can think of the design PEV as multiplying the errors in the data. The smaller the PEV, the greater the accuracy of your final model. You can read more about the calculation of PEV in “Prediction Error Variance”.

Try the other display options.

- The **View** menu has many options to change the look of the plots.
- You can change the factors displayed in the 2-D and 3-D plots. The pop-up menus below the plot select the factors, while the unselected factors are held constant. You can change the values of the unselected factors using the buttons and edit boxes in the **Input factors** list, top left.
- The `Movie` option shows a sequence of surface plots as a third input factor's value is changed. You can change the factors, replay, and change the frame rate.
- You can change the number, position, and color of the contours on the contour plot with the **Contours** button, as shown.

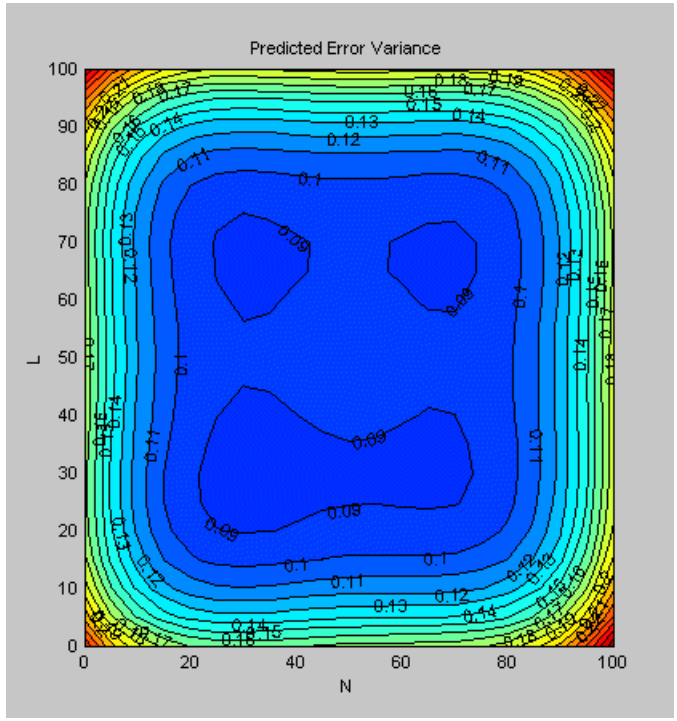

Improving the Design

You can further optimize the design by returning to the Optimal Design dialog, where you can delete or add points optimally or at random. The most efficient way is to delete points *optimally* and add new points *randomly* — these are the default algorithm settings. Only the existing points need to be searched for the most optimal ones to delete (the least useful), but the entire candidate set has to be searched for points to add optimally.

To strengthen the current optimal design:

1. Return to the Design Editor window.
2. Click the Optimal Design button in the toolbar again to reenter the dialog, and add 60 more points. Keep the existing points (which is the default).

3. Click **OK** and watch the optimization progress, then click **Accept** when the number of iterations without improvement starts increasing.

4. View the improvements to the design in the main displays.

5. Once again select **Tools > Prediction Error Variance Viewer** and review the plots of prediction error variance and the new values of optimality criteria in the optimality frame (bottom left). The shape of the PEV projection might not change dramatically, but note the changes in the scales as the design improves. The values of D, V, and G optimality criteria will also change (you have to click **Calculate** to see the values).

To see more dramatic changes to the design, return to the Design Editor window (no need to close the Prediction Error Variance Viewer).

1. Split the display so you can see a 3-D projection at the same time as a Table view.

2. You can sort the points to make it easier to select points in one corner. For example, to pick points where N is 100 and L is 0,

   a. Select **Edit > Sort Points**.

   b. Choose to sort by `N` only (reduce the number of sort variables to one) and click **OK**.

3. Choose **Edit > Delete Point**.

4. Using the Table and 3-D views as a guide, in the Delete Points dialog, pick six points to remove along one corner. Add the relevant point numbers to the delete list by clicking the add (>) button.

5. Click **OK** to remove the points. See the changes in the main design displays and look at the new Surface plot in the Prediction Error Variance Viewer (see the example following).

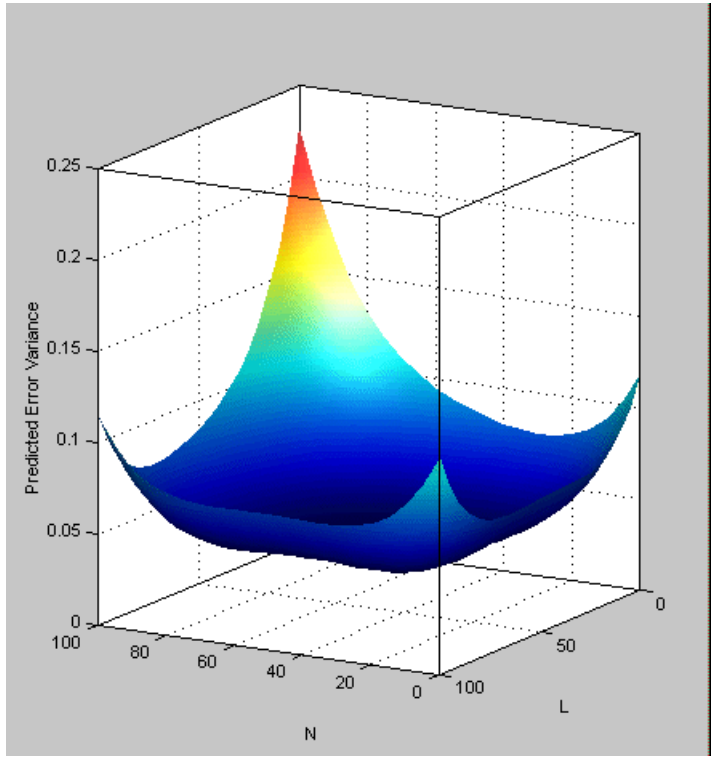
Predicted Error Variance
0.25
0.2
0.15
0.1
0.05
0
100
80
60
40
20
0
N
0
50
100
L

Creating Classical Designs

In this section...
“Adding a Classical Design” on page 6-27
“Classical Design Browser” on page 6-29
“Setting Up and Viewing a Classical Design” on page 6-30

Adding a Classical Design

1. In the Design Editor window, select the Optimal design in the design tree by clicking.

2. Add a new design. Use the first toolbar button, or select **File > New**.

   A new child node appears in the tree, called `Optimal_1`. Notice that the parent node now has a padlock on the icon. This indicates it is locked. This maintains the relationship between designs and their child nodes. The tree arrangement lets you try different operations starting from a basic design, then select the most appropriate one to use. The hierarchy allows clear viewing of the effects of changes on designs. The locking of parent designs also gives you the ability to easily reverse out of changes by retreating back up the tree.

3. Select the new design node in the tree. Notice that the display remains the same — all the points from the previous design remain, to be deleted or added to as necessary. The new design inherits all its initial settings from the currently selected design and becomes a child node of that design.

4. Rename the new node `Classical` by clicking again or by pressing **F2**.

5. Click the button in the toolbar or select **Design > Classical > Design Browser**.

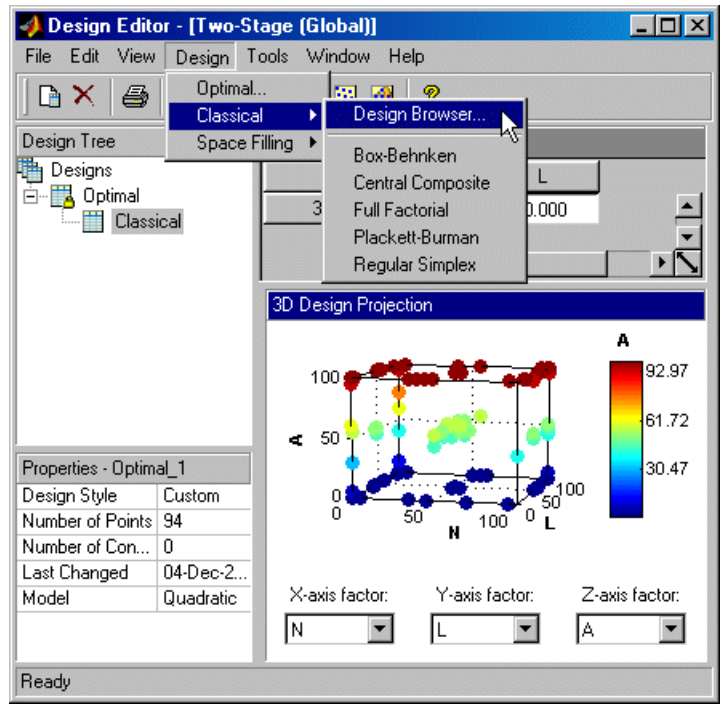

Note In cases where the preferred type of classical design is known, you can go straight to one of the five options under **Design > Classical**. Choosing the **Design Browser** option allows you to see graphical previews of these same five options before making a choice.

A dialog appears because there are already points from the previous design. You must choose between replacing and adding to those points or keeping only fixed points from the design.

6 Choose the default, replace current points with a new design, and click **OK**.

The Classical Design Browser appears.

Classical Design Browser

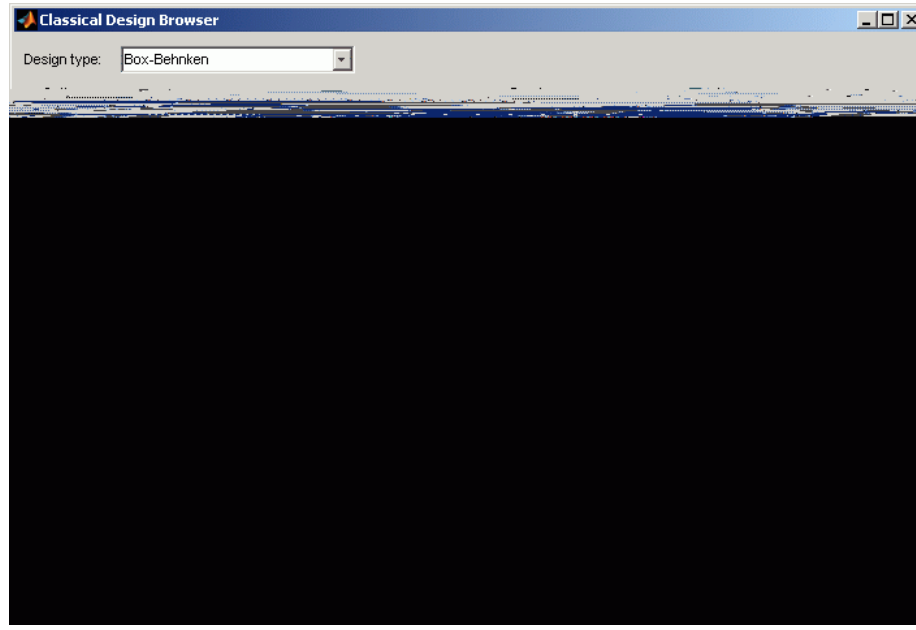

In the **Design Style** drop-down menu, there are five classical design options:

- `Central Composite`

  Generates a design that has a center point, a point at each of the design volume corners, and a point at the center of each of the design volume faces. You can choose a ratio value between the corner points and the face points for each factor and the number of center points to add. You can also specify a spherical design. Five levels are used for each factor.

- `Box-Behnken`

  Similar to Central Composite designs, but only three levels per factor are required, and the design is always spherical in shape. All the design points (except the center point) lie on the same sphere, so there should be at least three to five runs at the center point. There are no face points. These designs are particularly suited to spherical regions, when prediction at the corners is not required.

- `Full Factorial`

b Set the number of rows to be `10` and press Enter.

c Leave the value for each cell set to `0`.

d Click **Apply**.

6 Repeat step 5 for `F_SPK`.

Note The icons change as you initialize each table or function.

---

7 Click **Close** to leave the Calibration Manager.

After completing these steps, you can calibrate the lookup tables.

Calibrating a Feature

In this section...

"Process For Feature Calibration" on page 8-16
"Calibrating the Normalizers" on page 8-17
"Calibrating the Tables" on page 8-21
"Calibrating the Feature" on page 8-28

Process For Feature Calibration

The feature contains both a strategy (which is a collection of tables) and a model. You can use CAGE to fill the lookup tables using the model as a reference.

These are the three steps to calibrate a feature, described in these sections:

1 Calibrate the normalizers.

2 Calibrate the tables.

3 Calibrate the feature as a whole.

Click the expand icon, ⊞, to expand the nodes and display all the tables and normalizers in the feature.

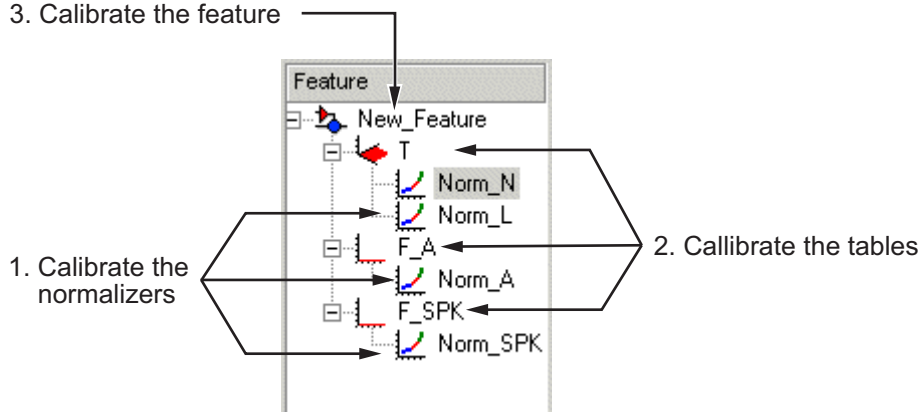

Each node in the display has a different view and different operations.

Calibrating the Normalizers

Normalizers are the axes for the lookup tables. Currently, `Norm_N` has 12 breakpoints; the other normalizers have 10 breakpoints each. This section describes how to set values for the normalizers `Norm_N` and `Norm_L`, based on the torque model, `tq`.

To display the Normalizer view, select the normalizer `Norm_N` in the branch display.

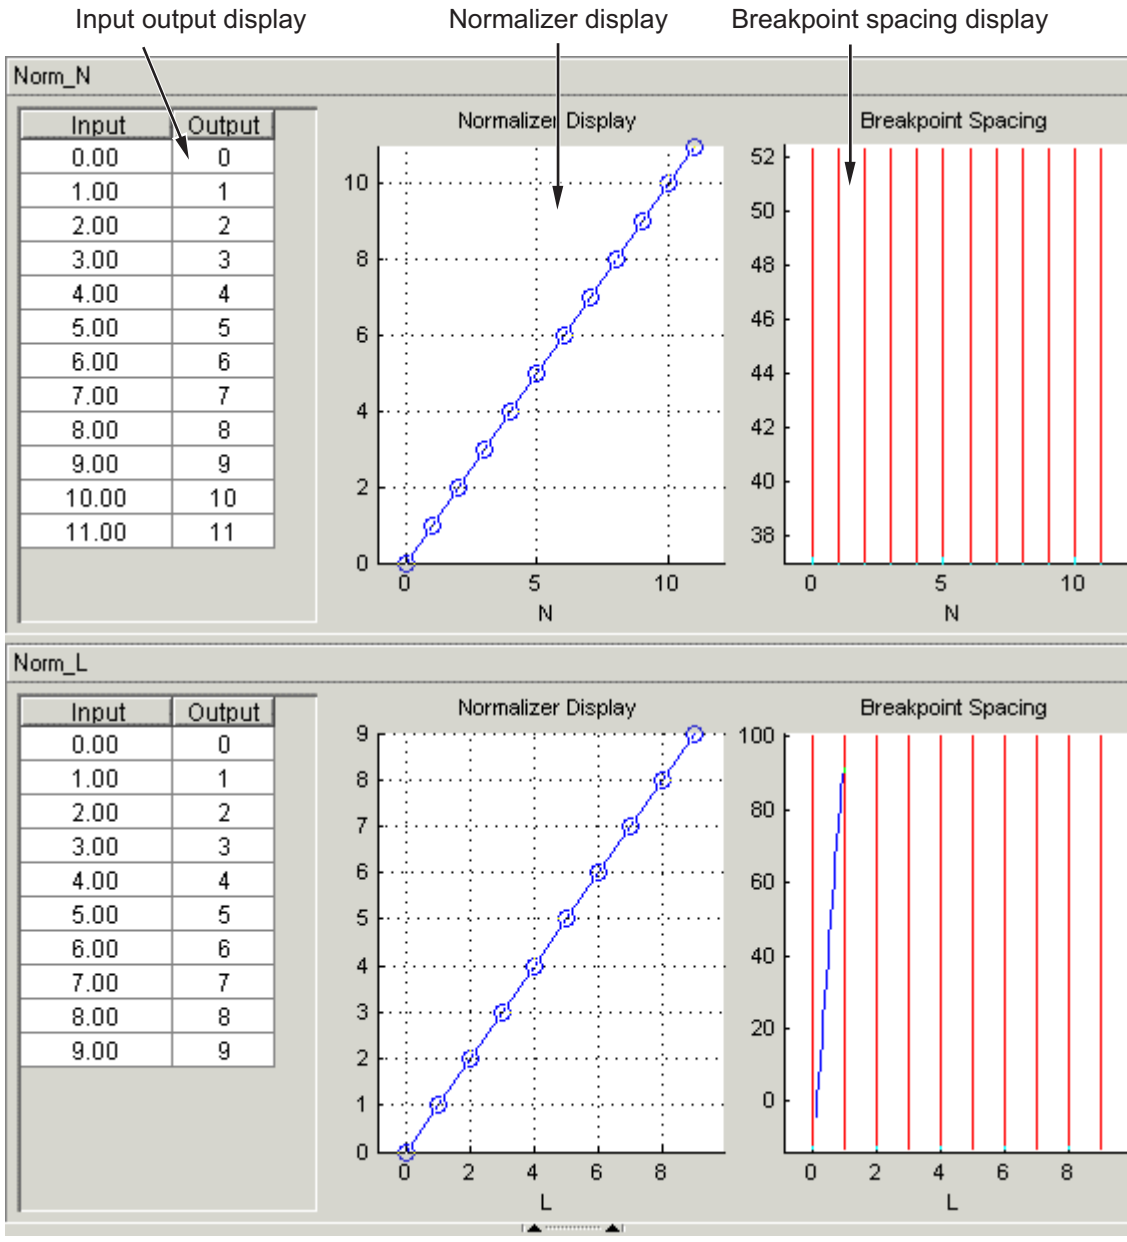
Input output display
Normalizer display
Breakpoint spacing display
Norm_N
Input
Output
0.00 0
1.00 1
2.00 2
3.00 3
4.00 4
5.00 5
6.00 6
7.00 7
8.00 8
9.00 9
10.00 10
11.00 11
Normalizer Display
Breakpoint Spacing
N
Norm_L
Input
Output
0.00 0
1.00 1
2.00 2
3.00 3
4.00 4
5.00 5
6.00 6
7.00 7
8.00 8
9.00 9
Normalizer Display
Breakpoint Spacing
L

The Normalizer view has two panes, **Norm_N** and **Norm_L**.

In each pane, you see

- An input/output table
- A normalizer display
- A breakpoint spacing display

In both Normalizer panes, the Input Output table and the **Normalizer Display** show the position of the breakpoints.

The **Breakpoint Spacing** display shows a blue slice through the model with the break points overlaid as red lines.

For a more detailed description of the Normalizer view, see "Normalizer View" in the CAGE documentation.

Placing the Breakpoints Automatically

You now must space the breakpoints across the range of each variable. For example, `Norm_N` takes values from 500 to 6500, the range of the engine speed.

To space the breakpoints evenly throughout the data values,

1. Click Initialize in the toolbar. Alternatively, select **Normalizer > Initialize**.

   This opens a dialog box that suggests ranges for `Norm_N` and `Norm_L`.

2. To accept the default ranges of values of the data, click **OK**.

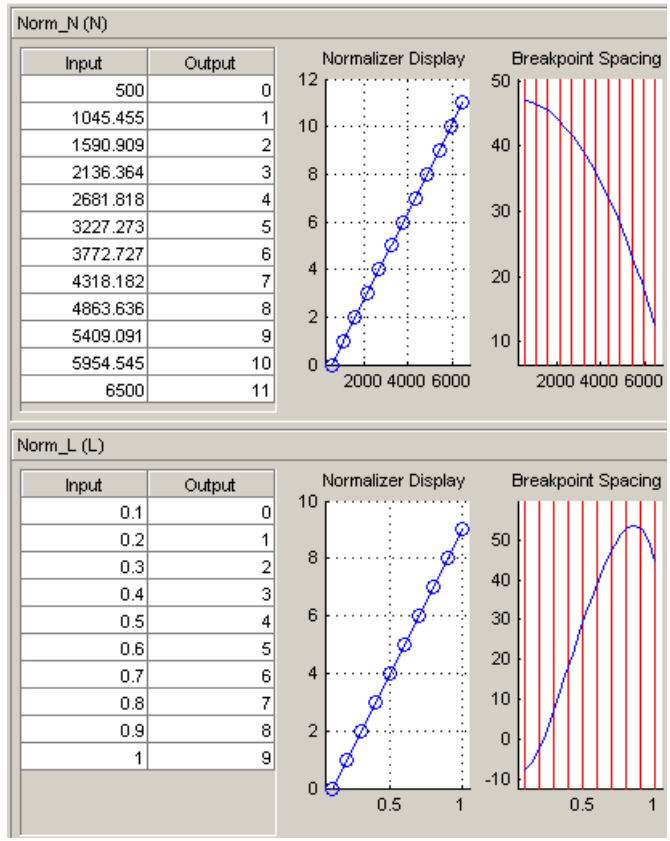

A better fit between model and table can often be achieved by spacing the breakpoints nonlinearly.

1 Click Fill in the toolbar. Alternatively, select **Normalizer > Fill**.

This opens a dialog box that suggests ranges for `Norm_N` and `Norm_L`. It also suggests values for `AFR` and `SPK`; these values are the set points for `AFR` and `SPK`.

2 To accept the values in the dialog box, click **OK**.

This ensures that the majority of the breakpoints are where the model is most curved. The table now has most values where the model changes

most. So, with the same number of breakpoints, the table is a better match to the model.

For more information about calibrating the normalizers, see “About Normalizers” in the CAGE documentation.

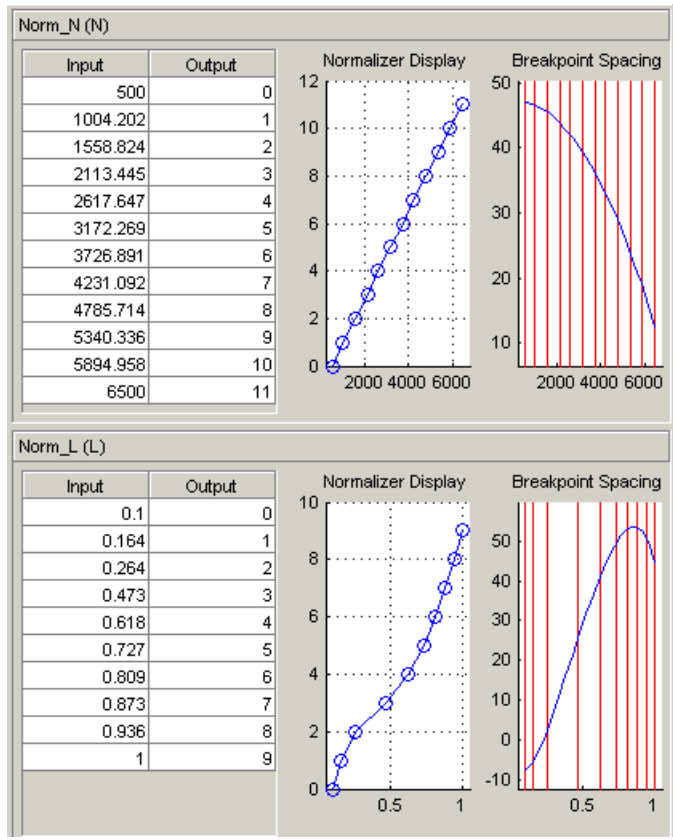

You can now calibrate the lookup tables; this is described in the next section.

Calibrating the Tables

The lookup tables currently have zero as the entry for each cell. This section demonstrates how to fill the table `T` with values of torque using the torque model, `tq`. To view the Table display, click the `T` node.

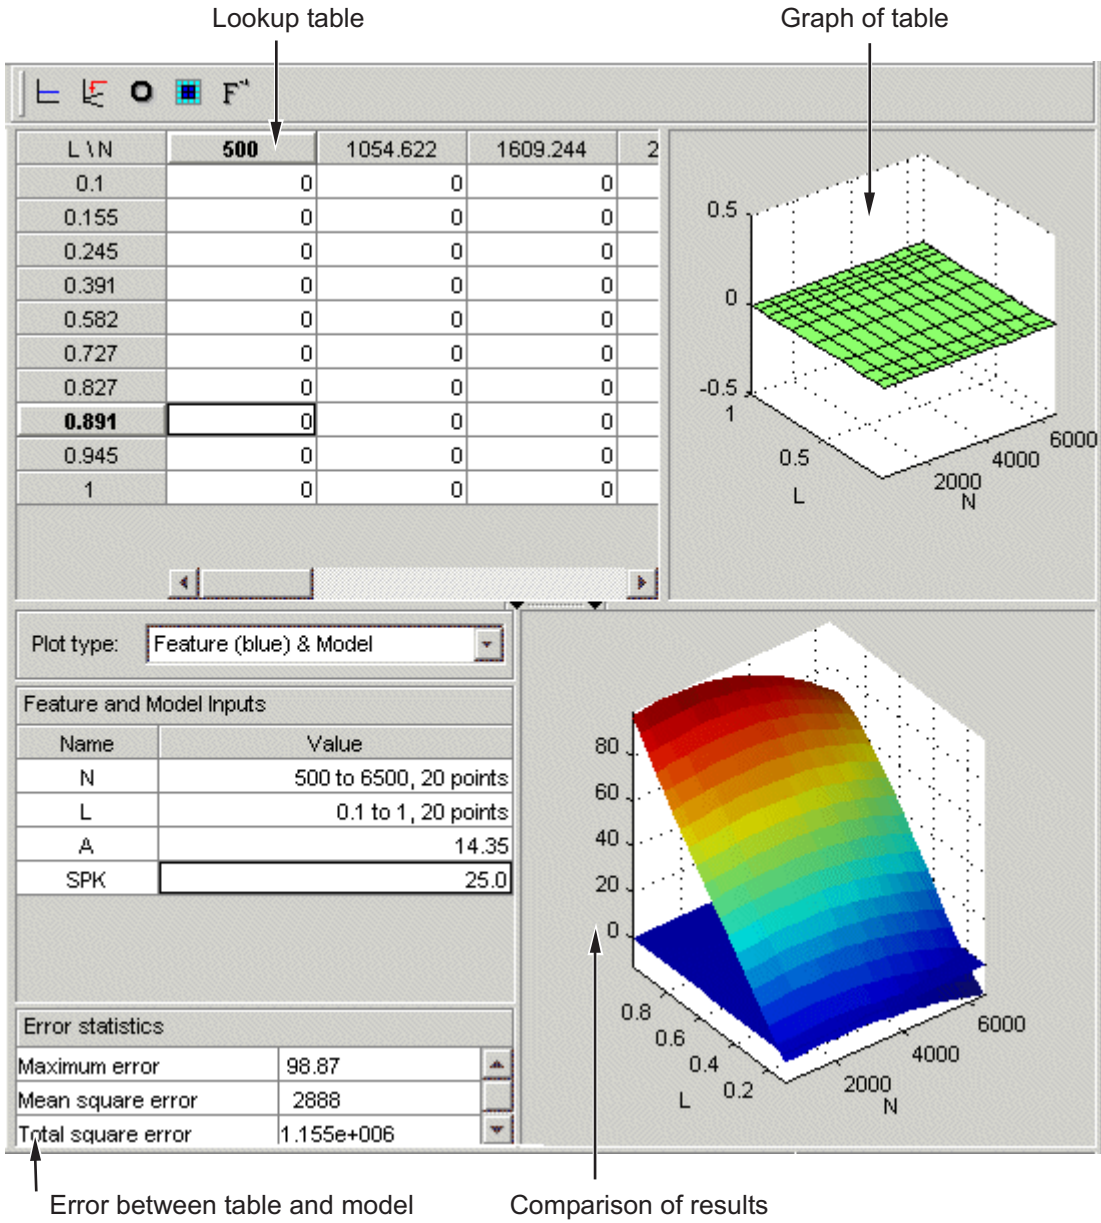
Lookup table
Graph of table
L \ N
500
1054.622
1609.244
0.1
0.155
0.245
0.391
0.582
0.727
0.827
0.891
0.945
1
0.5
0
-0.5
1
0.5
L
2000
4000
6000
N
Plot type:
Feature (blue) & Model
Feature and Model Inputs
Name
Value
N
500 to 6500, 20 points
L
0.1 to 1, 20 points
A
14.35
SPK
25.0
Error statistics
Maximum error
98.87
Mean square error
2888
Total square error
1.155e+006
80
60
40
20
0
0.8
0.6
0.4
0.2
L
2000
4000
6000
N
Error between table and model
Comparison of results

This view has three panes: the table, the graph, and the comparison-of-results pane.

To fill the table with values of the model at the appropriate operating points,

1 Click Fill on the toolbar.

This opens the Feature Fill Wizard.

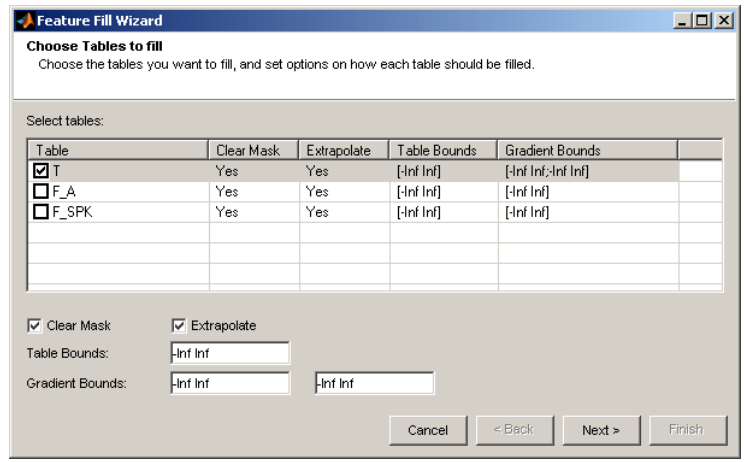

Observe the `T` table check box is selected, but you can also fill multiple tables at once using the wizard. You can **Fill** from the top feature node or from any table node in a feature. On the first screen you can set table bounds to avoid extrapolating to infeasible values, choose whether to only fill cells in the extrapolation mask, and choose whether to extrapolate automatically.

Leave the settings at the defaults and click **Next**.

2 On this screen you can see the `tq` model is selected to fill the table. Here you could also set up constraints, for example using a boundary model to constrain filling to table areas where data was collected, and you can link other models or features to inputs.

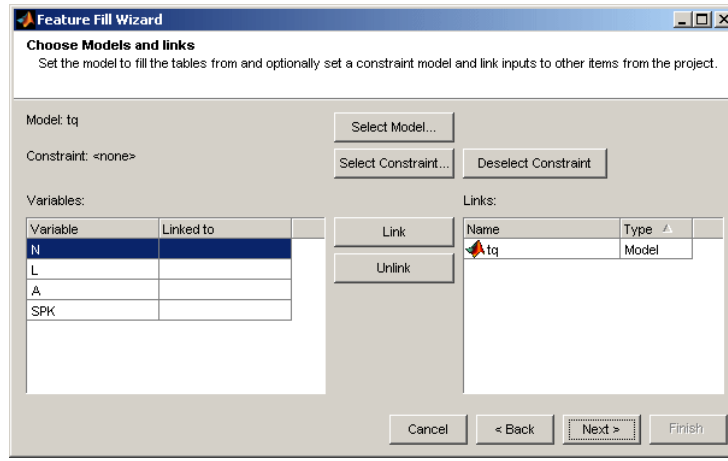

Leave the settings at the defaults and click **Next**.

3 Here you can set variable values for optimizing over.

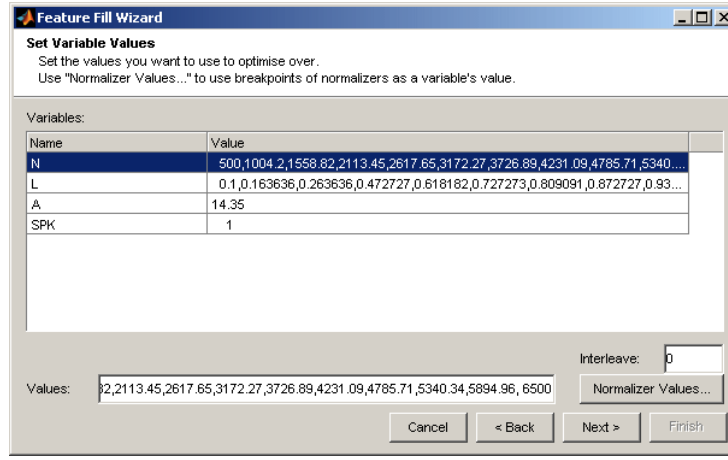

By default the table's normalizer breakpoints (here N and L) and the set points of the other variables (A and SPK) are selected. You can select

different normalizers, and edit values in the **Values** edit box to optimize over a range rather than at a single point. If you choose a range of values the table will be filled using the average model value at each cell. You can use the Interleave setting to add values between normalizer breakpoints to optimize over a finer grid than the number of table cells.

Leave the settings at the defaults and click **Next**.

4 Click **Fill Tables**.

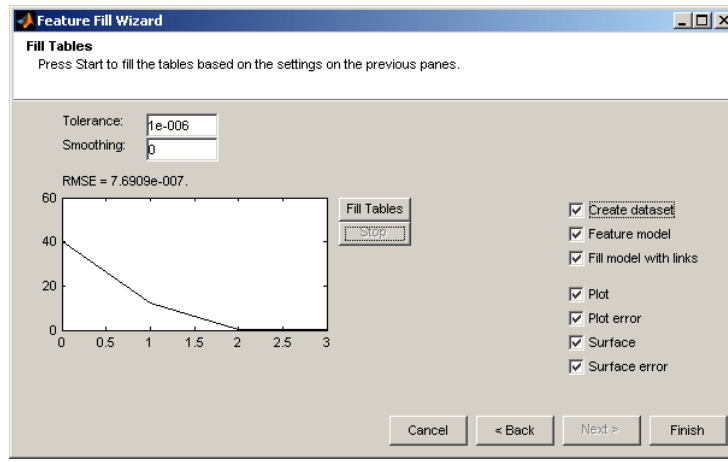

The graph shows the progress of the optimization. Select all the check boxes and click **Finish**.

Plots are created of the filled table surface and the error between the table values and the model.

5 Click **OK**.

The following view shows the table filled with values of the model.

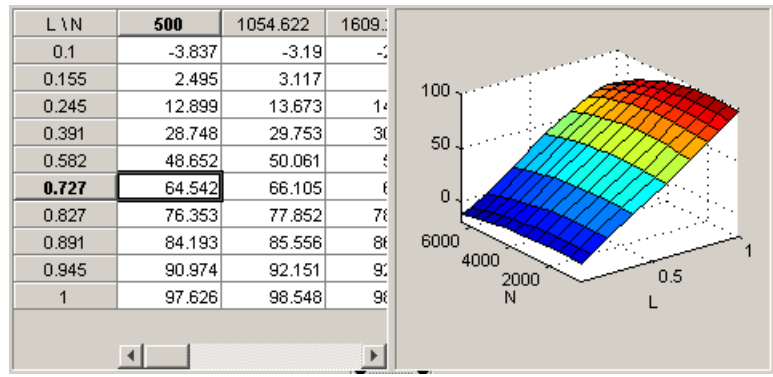

L \ N	500	1054.622	1609.
0.1	-3.837	-3.19	-2
0.155	2.495	3.117	
0.245	12.899	13.673	14
0.391	28.748	29.753	30
0.582	48.652	50.061	5
0.727	64.542	66.105	6
0.827	76.353	77.852	78
0.891	84.193	85.556	86
0.945	90.974	92.151	92
1	97.626	98.548	96

The following comparison-of-results pane shows how good a fit the strategy is to the model.

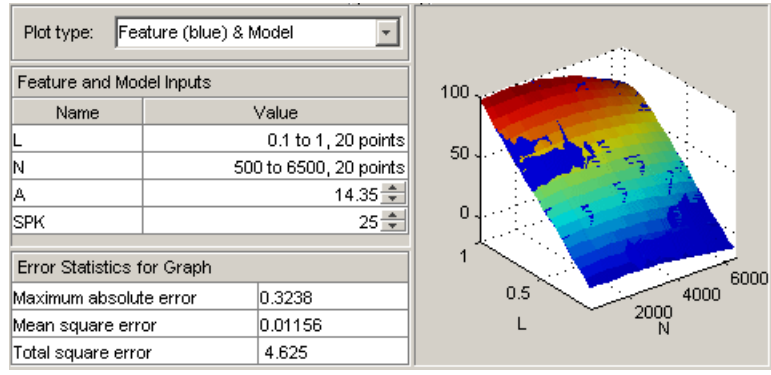

Plot type: Feature (blue) & Model

Feature and Model Inputs

Name	Value
L	0.1 to 1, 20 points
N	500 to 6500, 20 points
A	14.35
SPK	25

Error Statistics for Graph

Maximum absolute error	0.3238
Mean square error	0.01156
Total square error	4.625

The model is represented by the multicolored surface and the strategy is the blue surface.

The table `T` is now filled with optimized values compared to the model at these operating points.

For more information about the process of filling tables, see “Calibrating the Tables” in the CAGE documentation

Now you must fill the tables `F_A` and `F_SPK` and their normalizers. The tables are modifiers for AFR and the spark angle respectively. These steps are described in the next section.

Calibrating the Feature

A feature is a strategy (which is a collection of tables) and a model. Currently the torque table, `T`, is filled with optimized values compared to the torque model, `tq`. You must now calibrate the normalizers and tables for `F_A` and `F_SPK`.

You could calibrate the normalizers and then the tables for F_A and F_SPK in turn. However, CAGE enables you to calibrate the entire feature in one procedure.

To view the Feature view following, click the `New_Feature` node.

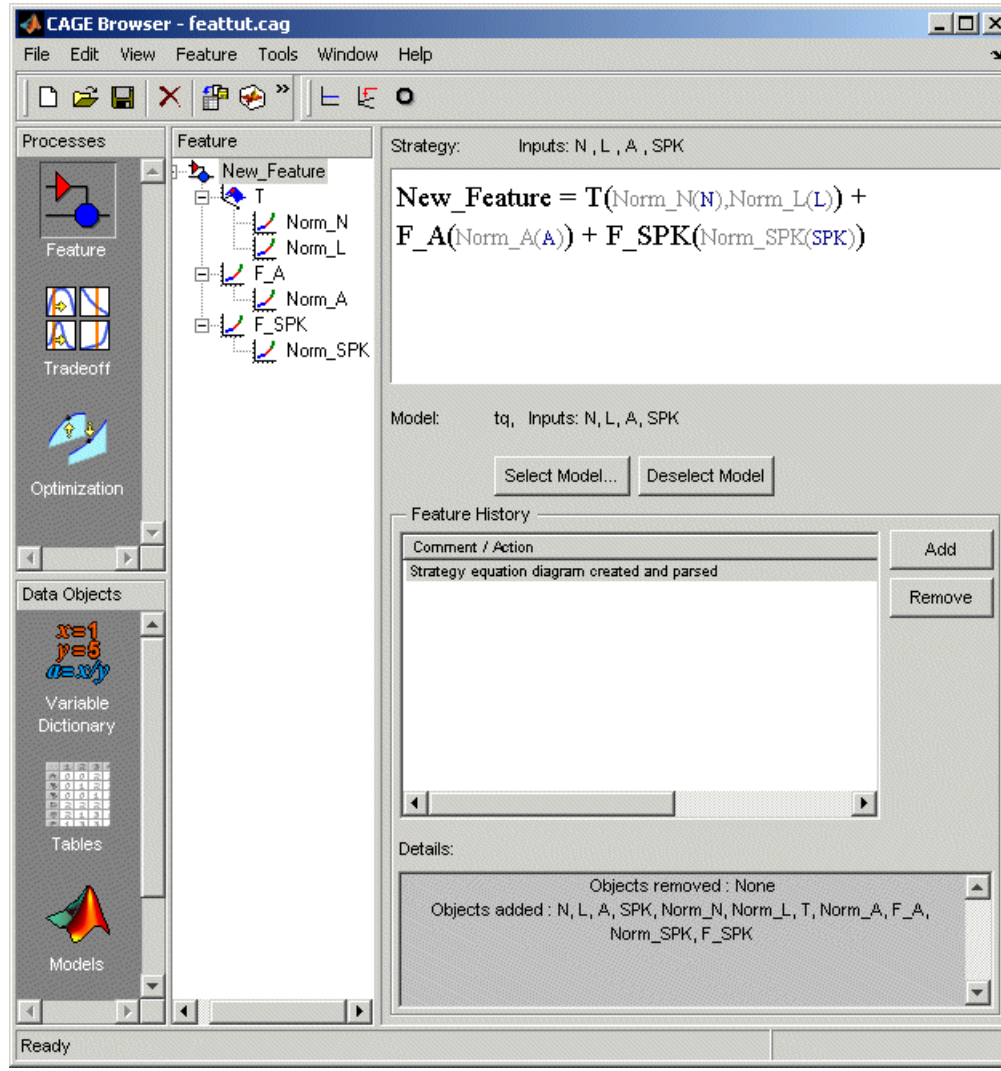

To calibrate all the tables and their normalizers,

1 Select **Feature > Initialize** (or use the Initialize toolbar button). The Feature Initialization Options dialog appears.

2 Clear the Enable check boxes for Breakpoints of T, and Values of T, as shown.

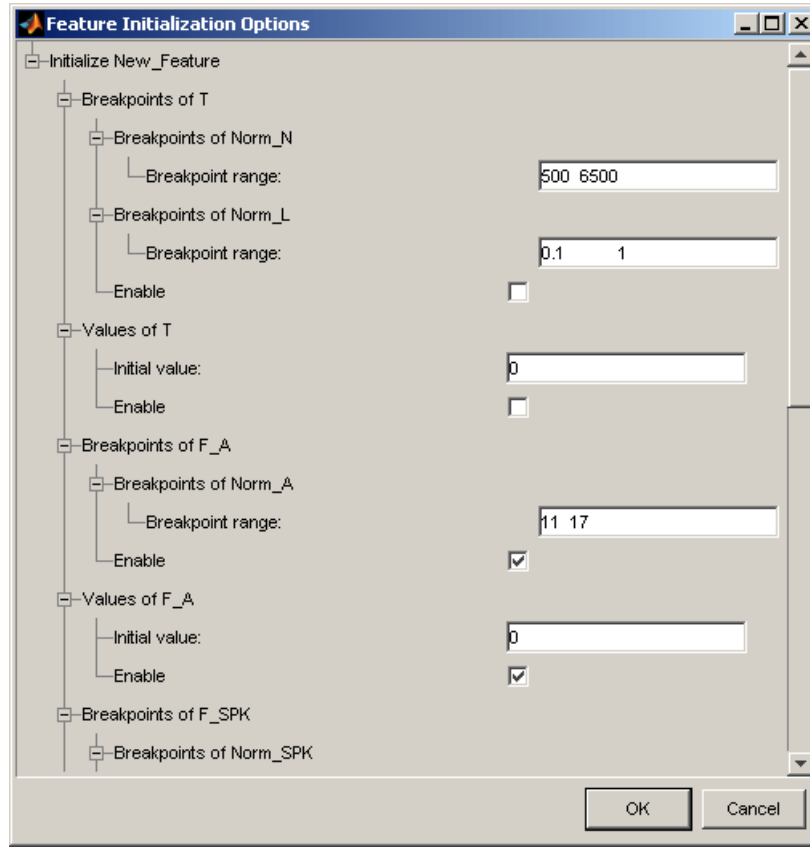

You have already optimized the breakpoints and table values for table T, so you only want to initialize the other tables `F_A` and `F_SPK`.

Click **OK**.

3 Select **Feature > Fill** (or use the Fill toolbar button) to open the Feature Fill Wizard. This time select only the F_A table to fill, and follow the steps in the wizard to fill this table: click Next 3 times then click **Fill Tables**. Select the F_A table node to view the results.

4 Select F_SPK table node and click Fill . Repeat the wizard steps to fill the table.

All three tables and normalizers are filled.

As the model and the feature are four-dimensional objects, it is difficult to fully view a comparison between the feature and the model. A meaningful comparison is shown in the lower half of the following figure (select the F_A node in the branch display). The equation *model* = *strategy* is rearranged so that the table is compared to the model and the remainder of the strategy. CAGE runs an optimization routine over the feature to minimize the total square error between the model and the feature.

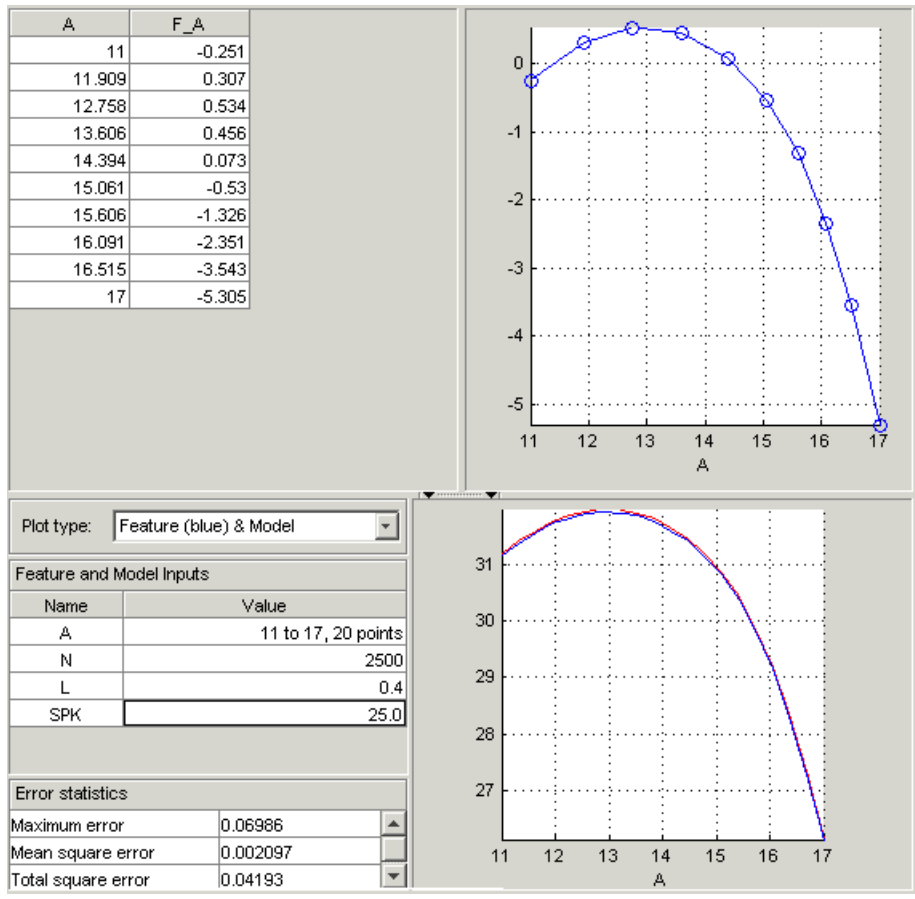

This display shows that the range of the normalizer for F_A is 11 to 17, the range of AFR. The lower pane shows a comparison between the blue line of the strategy and a red slice through the model, over the range of AFR.

This completes the calibration of the torque feature.

For more information about calibrating features, see “Calibrating the Feature Node” in the CAGE documentation.

You can now export the calibration.

Exporting Calibrations

To export your feature,

1. Select the `New_Feature` node in the branch display.
2. Select **File > Export > Calibration > Selected Item**.
3. Choose the type of file you want to save your calibrations as. For the purposes of this tutorial, select `Comma Separated Value (.csv)`.
4. Enter `tutorial.csv` as the file name and click **Save**.

This exports the calibration.

Note that when you choose to export**Selected Item** rather than **All Items**, what you export depends on which node is highlighted:

- Selecting a normalizer node outputs the values of the normalizer.
- Selecting a table node outputs the values of the table and its normalizers.
- Selecting a feature outputs the whole feature (all tables and normalizers).
- Selecting a branch node outputs all the features under the branch.

You have now completed the feature calibration tutorial.

9

Tutorial: Tradeoff Calibration

This section includes the following topics:

- "What Is a Tradeoff Calibration?" on page 9-2
- "Setting Up a Tradeoff Calibration" on page 9-3
- "Performing the Tradeoff Calibration" on page 9-8

What Is a Tradeoff Calibration?

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. Thus, a calibrator chooses the values of the input variables that produce this slight loss in torque over the values that produce the maximum value of torque.

This tutorial takes you through the various steps required for you to set up this tradeoff, and then to calibrate the lookup table for it.

Setting Up a Tradeoff Calibration

In this section...

"Creating a Tradeoff" on page 9-3

"Adding Tables to a Tradeoff Calibration" on page 9-5

"Displaying the Models" on page 9-6

Creating a Tradeoff

Start CAGE by typing

```
cage
```

at the MATLAB prompt.

Before you can calibrate the lookup tables, you must set up the calibration.

1. Select **File > Open Project** (or the toolbar button) to choose the `tradeoffInit.cag` file, found in the `matlab\toolbox\mbc\mbctraining` directory, then click **OK**.

   The `tradeoffInit.cag` project contains two models and all the variables necessary for this tutorial. For information about how to set up models and variables, see "Variables and Models" in the CAGE documentation.

2. To create a tradeoff calibration, select **File > New > Tradeoff**.

   This takes you to the **Tradeoff** view. You need to add tables and display models to the tradeoff, which are described step by step in the following sections:

   - "Adding Tables to a Tradeoff Calibration" on page 9-5.
   - "Displaying the Models" on page 9-6 describes how you display the models of torque and NOX emissions.

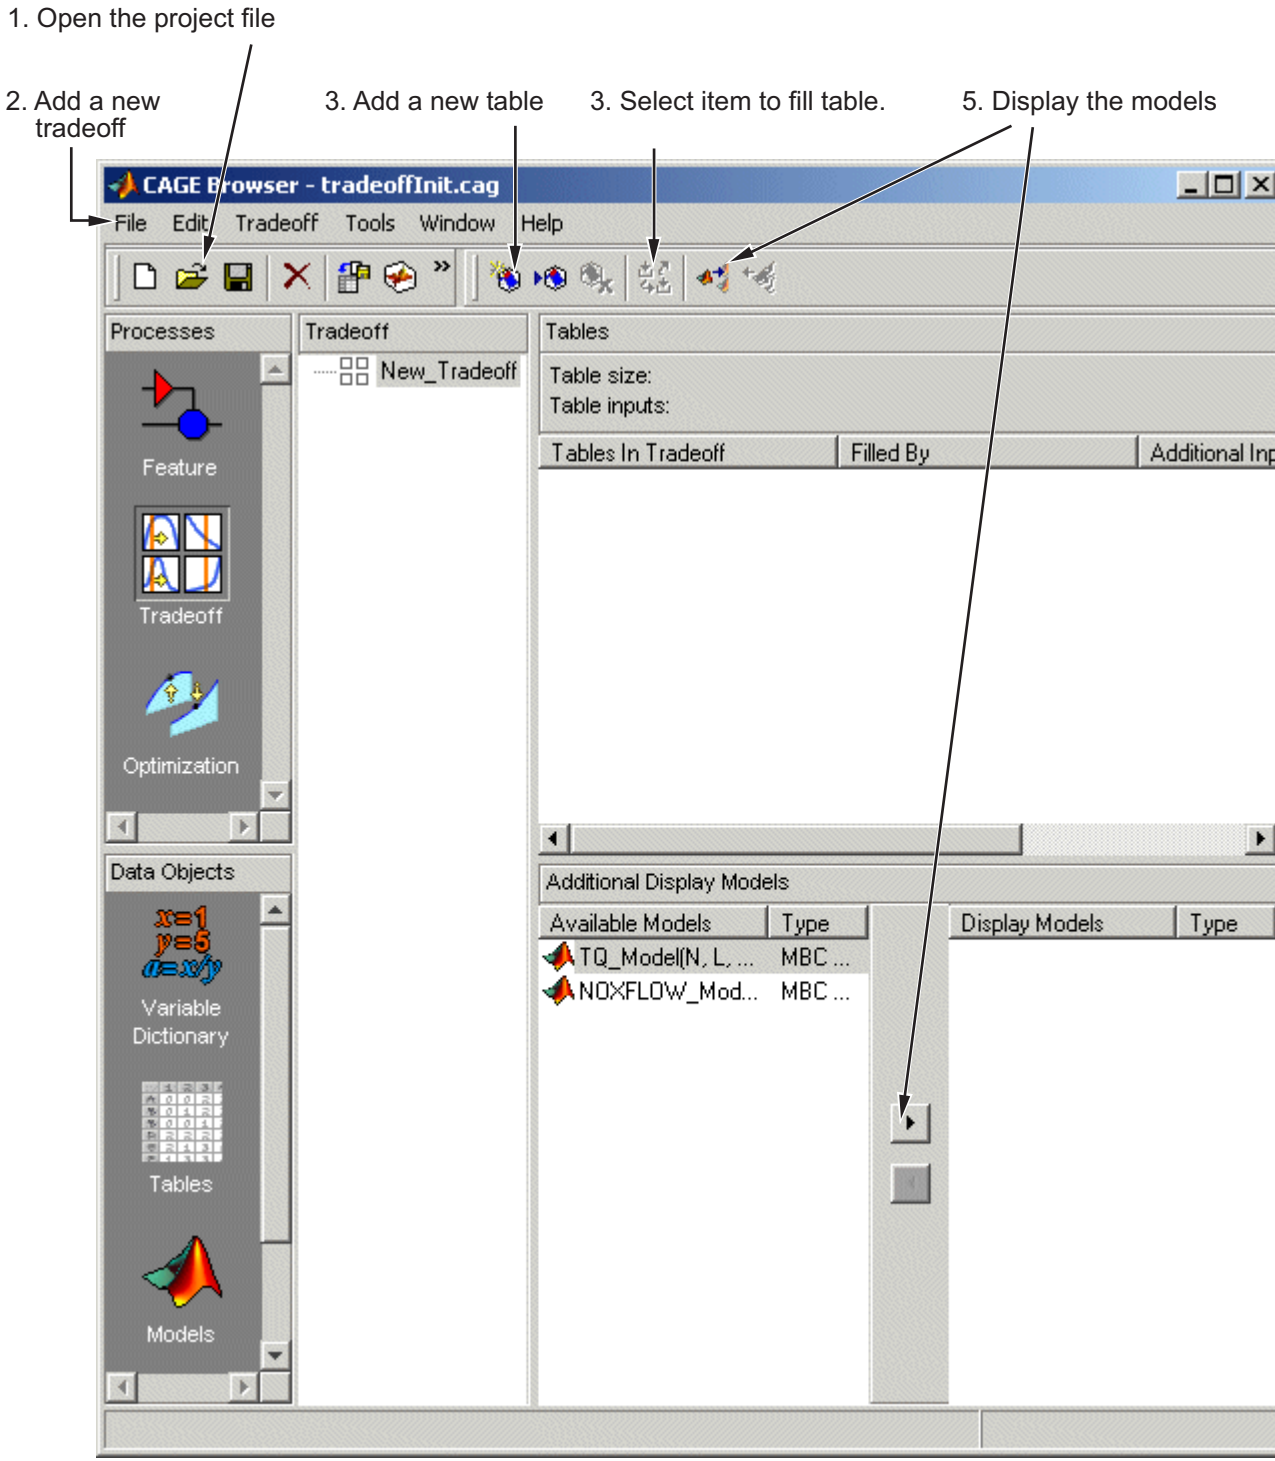
1. Open the project file
2. Add a new tradeoff
3. Add a new table
3. Select item to fill table.
5. Display the models
CAGE Browser - tradeoffInit.cag
File
Edit
Tradeoff
Tools
Window
Help
Processes
Feature
Tradeoff
Optimization
Data Objects
Variable Dictionary
Tables
Models
Tradeoff
New_Tradeoff
Tables
Table size:
Table inputs:
Tables In Tradeoff
Filled By
Additional Inp
Additional Display Models
Available Models
Type
TQ_Model(N, L, ...
MBC ...
NOXFLOW_Mod...
MBC ...
Display Models
Type

Adding Tables to a Tradeoff Calibration

The models of torque and NOX are in the current session. You must add the lookup table to calibrate.

Both models have five inputs. The inputs for the torque and NOX models are

- Exhaust gas recycling (EGR)
- Air/fuel ratio (AFR)
- Spark angle
- Speed
- Load

For this tutorial, you are interested in the spark angle over the range of speed and load.

To generate a lookup table for the spark angle,

1 Click (Add New Table) in the toolbar. This opens the Table Setup dialog.

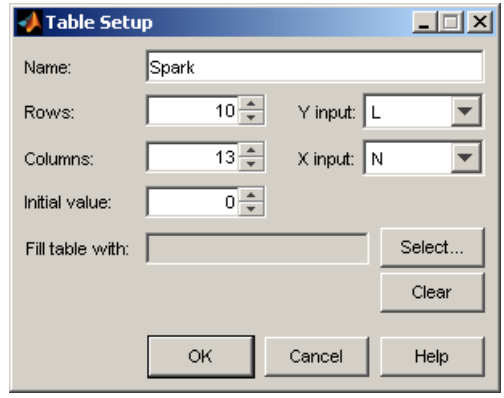

2 Enter `Spark` as the table **Name**.

3 Check that `N` is the **X input** and `L` is the **Y input** (these are selected automatically as the first two variables in the current Variable Dictionary).

4 Enter 10 as the size of the load axis (**Rows**).

5 Enter 13 as the size of the speed axis (**Columns**).

6 Click **Select** to open the dialog Select Filling Item.

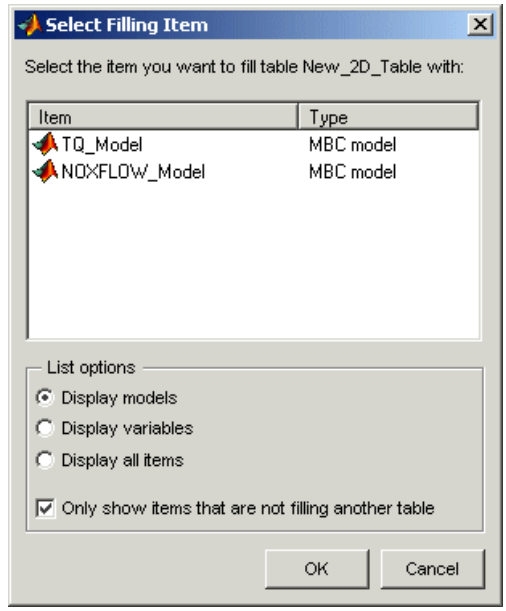

Select the radio button **Display variables**, then select SPK to fill the table and click **OK**.

7 Click **OK** to close the Table Setup dialog.

Before you can perform the calibration, you must display the models.

Displaying the Models

For this tutorial, you are comparing values of the torque and NOX models. Thus, you need to display these models.

To display both models,

- Click Add Model to Display List in the toolbar twice. This will move both available models into the Display list.
- Alternatively, **Shift**-click to select both models in the **Available Models** list and click to include both models in the current display. In this case you want to include all available models. You can click to select particular models in the list to display.

The **Display Models** pane following shows both models selected for display.

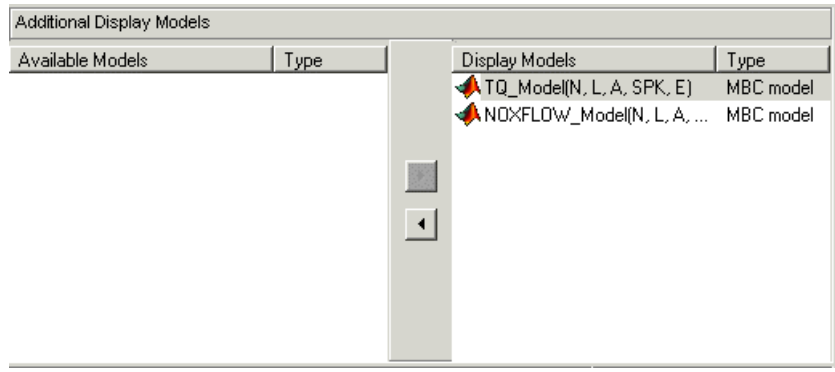

You can now calibrate the tradeoff.

Performing the Tradeoff Calibration

In this section...

"Process Overview" on page 9-8
"Checking the Normalizers" on page 9-10
"Setting Values for Other Variables" on page 9-11
"Filling Key Operating Points" on page 9-13
"Filling the Table by Extrapolation" on page 9-17
"Exporting Calibrations" on page 9-19

Process Overview

You now fill the lookup table for spark angle by trading off gain in torque for reduction in NOX emissions.

The method that you use to fill the lookup table is

- Obtain the maximum possible torque.
- Restrict NOX to below 250 g/hr at any operating point.

To perform the tradeoff calibration, follow the instructions in the next four sections:

1. Check the normalizers.
2. Set values for the other variables, AFR and EGR.
3. Fill key operating points with values for spark angle.
4. Fill the table by extrapolation.

Once you have completed the calibration, you can export the calibration for use in the electronic control unit.

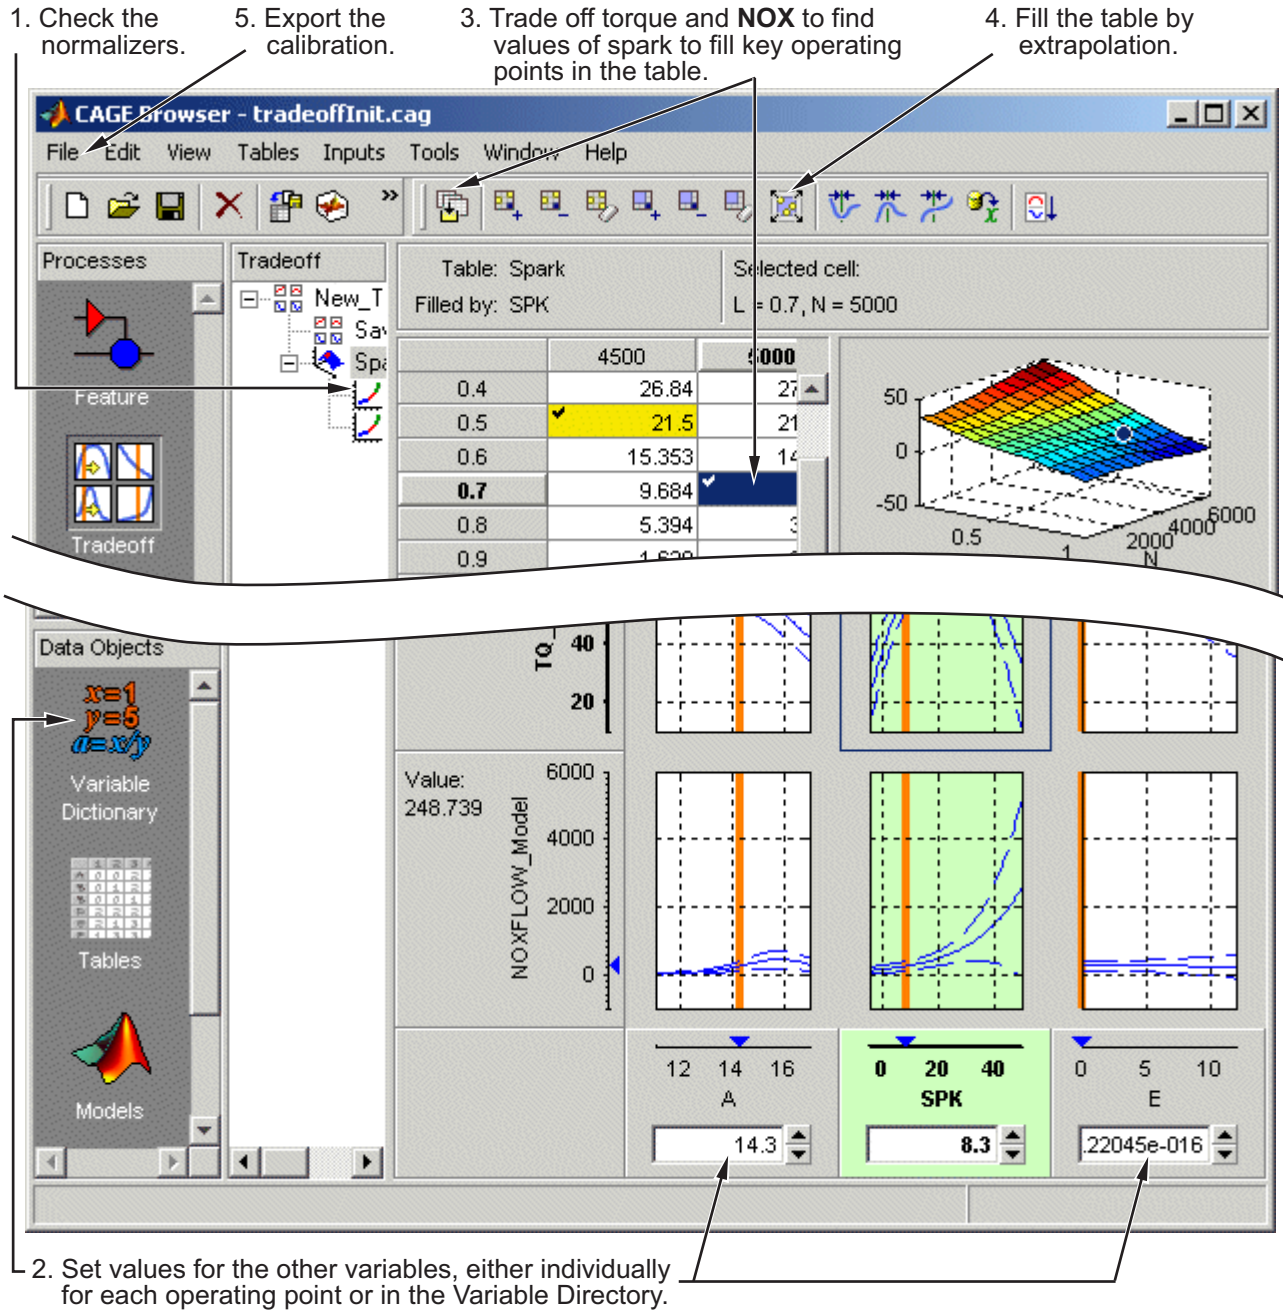
1. Check the normalizers.
5. Export the calibration.
3. Trade off torque and **NOX** to find values of spark to fill key operating points in the table.
4. Fill the table by extrapolation.
CAGE Browser - tradeoffInit.cag
File Edit View Tables Inputs Tools Window Help
Processes
Feature
Tradeoff
Tradeoff
New_T
Table: Spark
Filled by: SPK
Selected cell:
L = 0.7, N = 5000
4500
0.4 26.84
0.5 21.5
0.6 15.353
0.7 9.684
0.8 5.394
0.9
Data Objects
Variable Dictionary
Tables
Models
TQ
Value: 248.739
NOXFLOW_Model
12 14 16
A
14.3
0 20 40
SPK
8.3
0 5 10
E
.22045e-016
2. Set values for the other variables, either individually for each operating point or in the Variable Directory.

Checking the Normalizers

A normalizer is the axis of the lookup table (which is the collection of breakpoints). The breakpoints of the normalizers are automatically spaced over the ranges of speed and load. These define the operating points that form the cells of the tradeoff table.

Expand the Tradeoff tree by clicking the plus sign in the display, so you can see the `Spark` table and its normalizers `Speed` and `Load`. Click to highlight either normalizer to see the normalizer view. A tradeoff calibration does not compare the model and the table directly, so you cannot space the breakpoints by reference to the model.

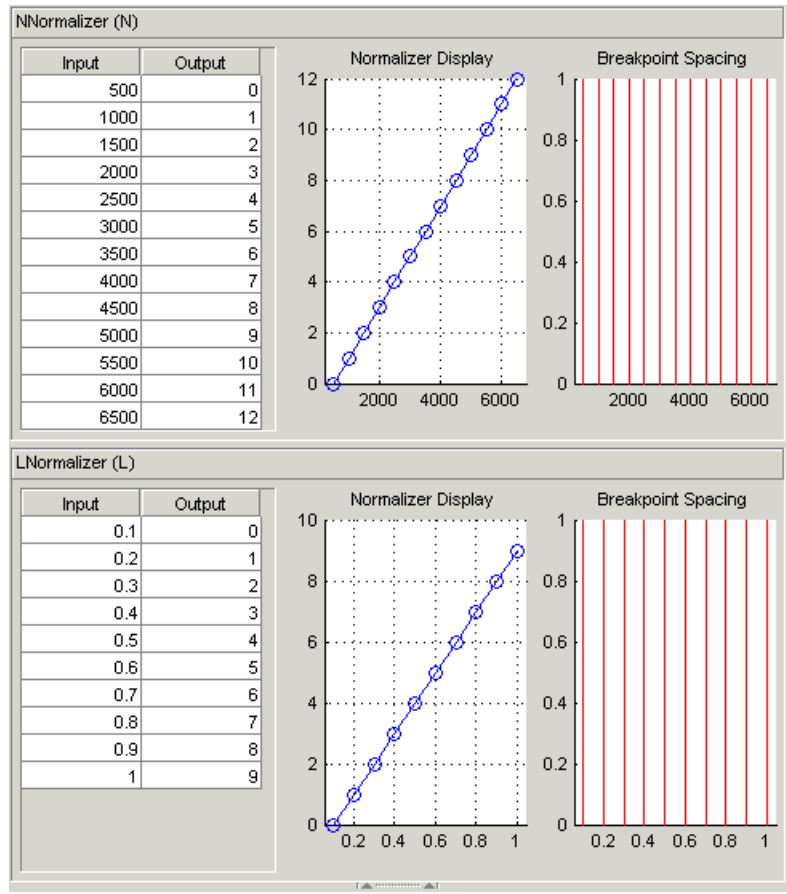

Setting Values for Other Variables

At each operating point, you must fill the values of the spark table. Both of the models depend on spark, AFR (labeled `A`, in the session), and EGR (labeled `E` in the session). You could set the values for AFR and EGR individually for each operating point in the table, but for simplicity you will set constant values for these model inputs.

To set constant values of AFR and EGR for all operating points,

1. Click **Variable Dictionary** in the **Data Objects** pane.
2. Click `A` and edit the **Set Point** to `14.3`, the stoichiometric constant, and press **Enter**.
3. Click `E` and change the **Set Point** to `0` and press **Enter**.

   You have set these values for every operating point in your tradeoff table. You can now fill the spark angle lookup table. The process is described next.
4. Click **Tradeoff** in the **Processes** pane to return to the tradeoff view.
5. Highlight the `Spark` table node in the Tradeoff tree display.
6. In the lower pane, check that the value for `A` is `14.3`, and the value for `E` is `0`, as shown in the following example. You leave these values unchanged for each operating point.

For each operating point you change the values of spark to trade off the torque and NOX objectives; that is, you search for the best value of spark that gives acceptable torque within the emissions constraint. The following example illustrates the controls you use, and there are step-by-step instructions in the following section.

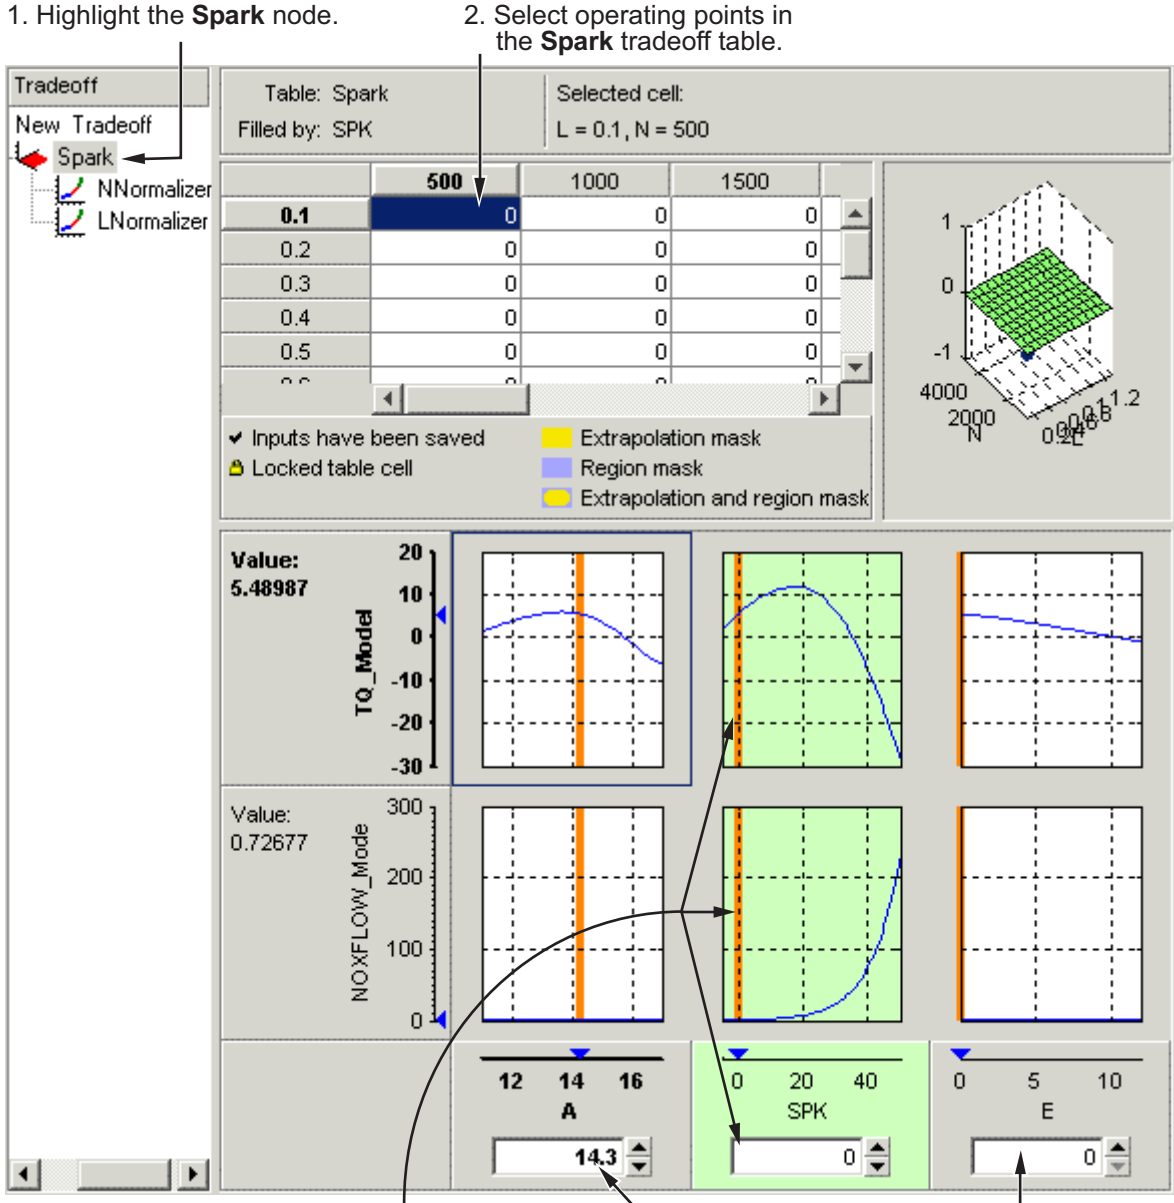
1. Highlight the **Spark** node.
2. Select operating points in the **Spark** tradeoff table.
Tradeoff
New Tradeoff
Spark
NNormalizer
LNormalizer
Table: Spark
Filled by: SPK
Selected cell:
L = 0.1, N = 500
500
1000
1500
0.1
0.2
0.3
0.4
0.5
Inputs have been saved
Locked table cell
Extrapolation mask
Region mask
Extrapolation and region mask
Value:
5.48987
TQ_Model
Value:
0.72677
NOXFLOW_Mode
12 14 16
A
14.3
0 20 40
SPK
0
0 5 10
E
0
3. Change values of **SPK** to trade off torque and NOX emissions at each operating point.
4. Leave A and E at the set point values.

Filling Key Operating Points

You now fill the key operating points in the lookup table for spark angle.

The upper pane displays the lookup table, and the lower pane displays the behavior of the torque and NOX emissions models with each variable.

The object is to maximize the torque and restrict NOX emissions to below 250 g/hr.

Determining the Value of Spark

At each operating point, the behavior of the model alters. The following display shows the behavior of the models over the range of the input variables at the operating point selected in the table, where speed (N) is 4500 and load (L) is 0.5. You can show confidence intervals by selecting **View > Display Confidence Intervals**.

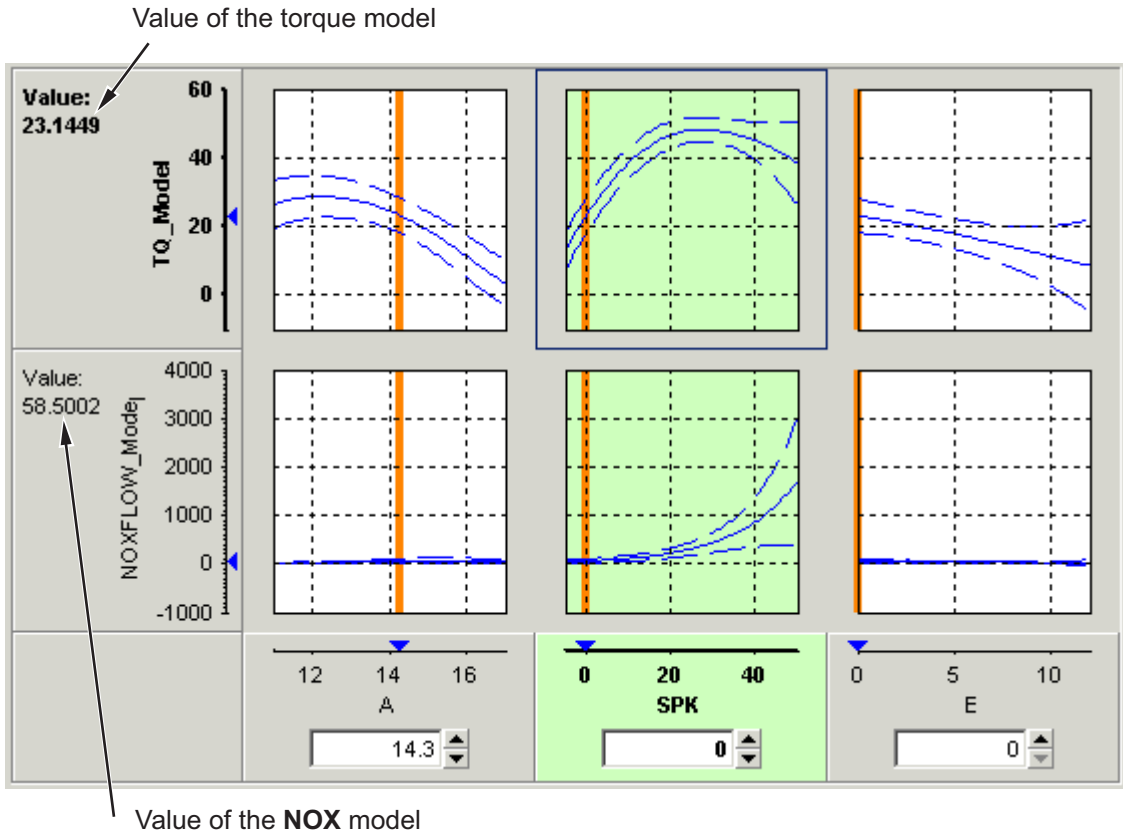

The top three graphs show how the torque model varies with the AFR (labeled A), the spark angle (SPK), and the EGR (E), respectively. The lower panes show how the NOX emissions model varies with these variables.

You are calibrating the Spark table, so the two spark (SPK) graphs are green, indicating that these graphs are directly linked to the currently selected lookup table.

1. Select the operating point N = 4500 and L = 0.5 in the lookup table.

2. Now try to find the spark angle that gives the maximum torque and restricts NOX emissions to below 250 g/hr. You can change the value of spark by clicking and dragging the orange line on the SPK graphs, or by

typing values into the `SPK` edit box. You can change the values of any of the other tradeoff variables in the same way, but as you have already set constant values for A and E you should not change these. Try different values of spark and look at the resulting values of the torque and NOX models.

3 Click to select the top SPK - TQ_Model graph (`TQ_Model` row, `SPK` column). When selected the graph is outlined as shown in the following example.

4 Now click 'Find maximum of output' ( ) in the toolbar. This calculates the value of spark that gives the maximum value of torque. The following display shows the behavior of the two models when the spark angle is `26.4458`, which gives maximum torque output.

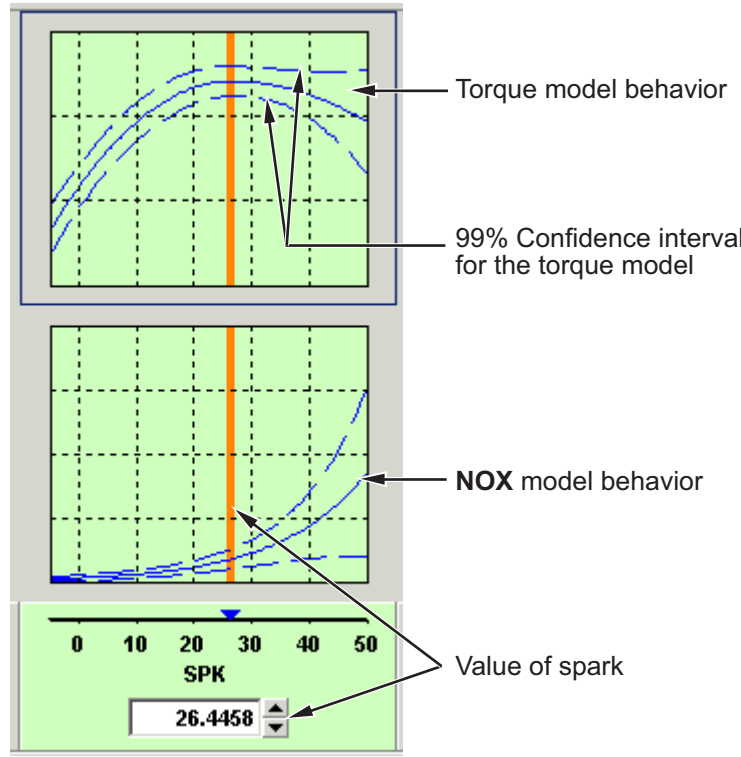

At this operating point, the maximum torque that is generated is `48.136` when the spark angle is `26.4989`. However, the value of NOX is `348.968`, which is greater than the restriction of 250 g/hr. Clearly you have to look at another value of spark angle.

5. Click and drag the orange bar to change to a lower value of spark. Notice the change in the resulting values of the torque and NOX models.

6. Enter `21.5` as the value of **SPK** in the edit box at the bottom of the SPK column.

   The value of the NOX emissions model is now `249.154`. This is within the restriction, and the value of torque is `47.2478`.

   At this operating point, this value of `21.5` degrees is acceptable for the spark angle lookup table, so you want to apply this point to your table.

7. Press **Ctrl**+T or click (Apply table filling values) in the toolbar to apply that value to the spark table.

   This automatically adds the selected value of spark to the table and turns this cell yellow. It is blue when selected, yellow if you click elsewhere. Look at the table legend to see what this means: yellow cells have been added to the extrapolation mask, and the tick mark indicates you saved this input value by applying it from the tradeoff. You can use the **View** menu to choose whether to display the legend.

8. Now repeat this process of finding acceptable values of spark at four more operating points listed in the table following. In each case,

   - Select the cell in the spark table at the specified values of speed and load.
   - Enter the value of spark given in the table (the spark angles listed all satisfy the requirements).
   - Press **Ctrl**+T or click (Apply table filling values) in the toolbar to apply that value to the spark table.

Speed, N	Load, L	Spark Angle, SPK
2500	0.3	25.75
3000	0.8	10.7

Speed, N	Load, L	Spark Angle, SPK
5000	0.7	8.2
6000	0.2	41.3

After you enter these key operating points, you can fill the table by extrapolation. This is described in the next section.

Filling the Table by Extrapolation

When you have calibrated several key operating points, you can produce a smooth extrapolation of these values across the whole table.

When you apply the value of the spark angle to the lookup table, the selected cell is automatically added to the extrapolation mask. This is why the cell is colored yellow. The extrapolation mask is the set of cells that are used as the basis for filling the table by extrapolation.

Click in the toolbar to fill the table by extrapolation.

The lookup table is filled with values of spark angle.

The following figure displays the view after extrapolation over the table.

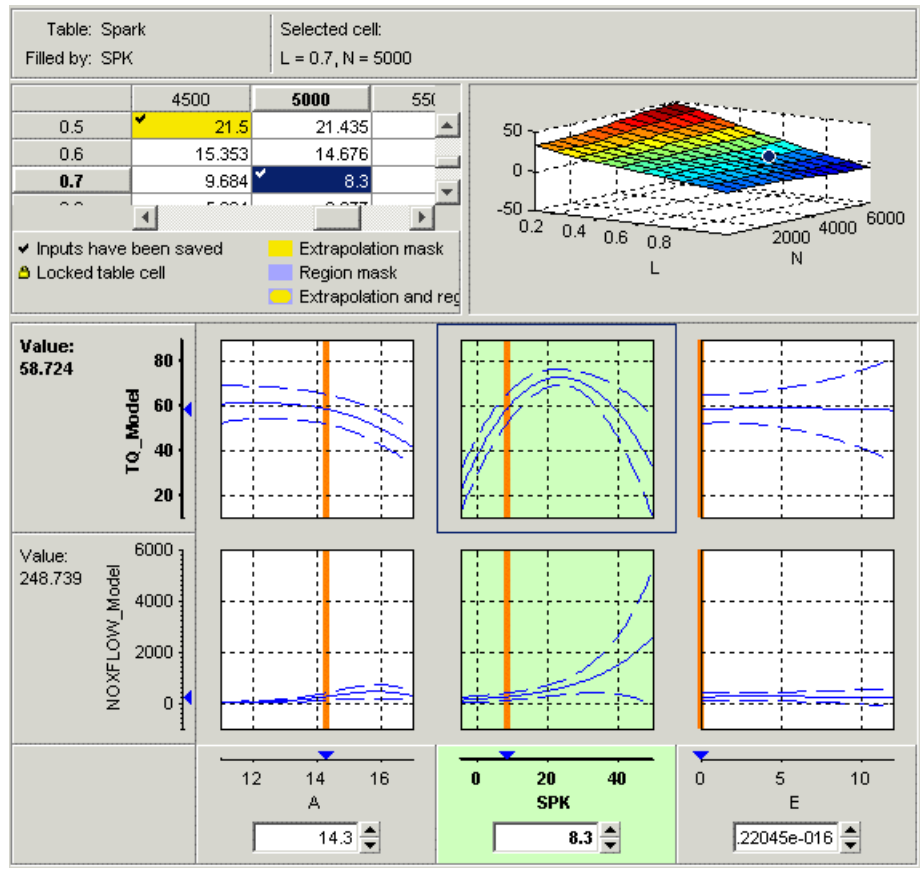

Note Not all the points in the lookup table will necessarily fulfill the requirements of maximizing torque and restricting the NOX emissions.

You could use these techniques to further improve the calibration and trade off torque and NOX to find the best values for each cell in the spark table.

For a more detailed description of tradeoff calibrations, see “Tradeoff Calibrations” in the CAGE documentation.

You can now export this calibration to file.

Exporting Calibrations

To export your table and its normalizers,

1. Select the `Spark` node in the branch display.
2. Select **File > Export > Calibration**.
3. Choose the file type you want for your calibration. For the purposes of this tutorial, select `Comma Separated Value (.csv)`.
4. Enter `tradeoff.csv` as the file name and click **Save**.

This exports the spark angle table and the normalizers, `Speed` and `Load`.

You have now completed the tradeoff calibration tutorial.

10

Tutorial: Data Sets

This section includes the following topics:

- "Setting Up the Data Set" on page 10-2
- "Comparing the Items in a Data Set" on page 10-8
- "Reassigning Variables" on page 10-16

Setting Up the Data Set

In this section...

"Tutorial Overview" on page 10-2

"Opening an Existing Calibration" on page 10-3

"Importing Experimental Data into a Data Set" on page 10-3

"Adding an Item to a Data Set" on page 10-6

Tutorial Overview

You can use the **Data Sets** view in CAGE to compare features, tables, and models with experimental data. You can use data sets to plot the features, tables, etc., as tabular values or as plots on a graph.

Data sets enable you to view the data at a set of operating points. You can determine the set of operating points yourself, using Build Grid. Alternatively, you can import a set of experimental data taken at a series of operating points. These operating points are not the same as the breakpoints of your tables.

This tutorial takes you through the basic steps required to compare a completed feature calibration to a set of experimental data.

Start CAGE by typing

```
cage
```

at the MATLAB prompt.

To set up the data set tutorial, you need to

1. Open an existing calibration.
2. Import the experimental data.
3. Add the `Torque` feature to the data set.

Your data set contains all the input factors and output factors required. As the imported data contains various operating points, this information is also included in the data set.

The next sections describe these processes in more detail.

Opening an Existing Calibration

For this tutorial, use the file `datasettut.cag`, found in the `matlab\toolbox\mbc\mbctraining` directory.

To open this file,

1. Select **File > Open Project**.
2. In the file browser, select `datasettut.cag` and click **Open**.

   This opens a file that contains a complete calibrated feature with its associated models and variables. This particular feature is a torque calibration, using a torque table (labeled `T1`) and modifiers for spark (labeled `T2`) and air/fuel ratio (labeled `T3`).

   For information about completing a feature calibration, see "Feature Calibrations" in the CAGE documentation.
3. Select **File > New > Data Set** to add a new data set to your session.

This automatically switches you to the **Factor Information** pane of the data set display.

Importing Experimental Data into a Data Set

To import data into a data set,

1. Select **File > Import > Data**.
2. In the file browser, select `meas_tq_data.xls` from the `mbctraining` directory, and click **Open**.

   This set of data includes six columns of data, the test cell settings for engine speed (`RPM`), and the measured values of torque (`tqmeas`), engine

speed (nmeas), air/fuel ratio (afrmeas), spark angle (spkmeas), and load (loadmeas).

3 The Data Set Import Wizard asks which of the columns of data you would like to import. Click **Next** to import them all.

The following screen asks you to associate variables in your project with data columns in the data.

4 Highlight afr in the **Project Assignments** column and afrmeas in the **Data Column**, then click the assign button, shown.

5 Repeat this to associate load with loadmeas, n with RPM, and spk with spkmeas. The dialog box should be the same as shown.

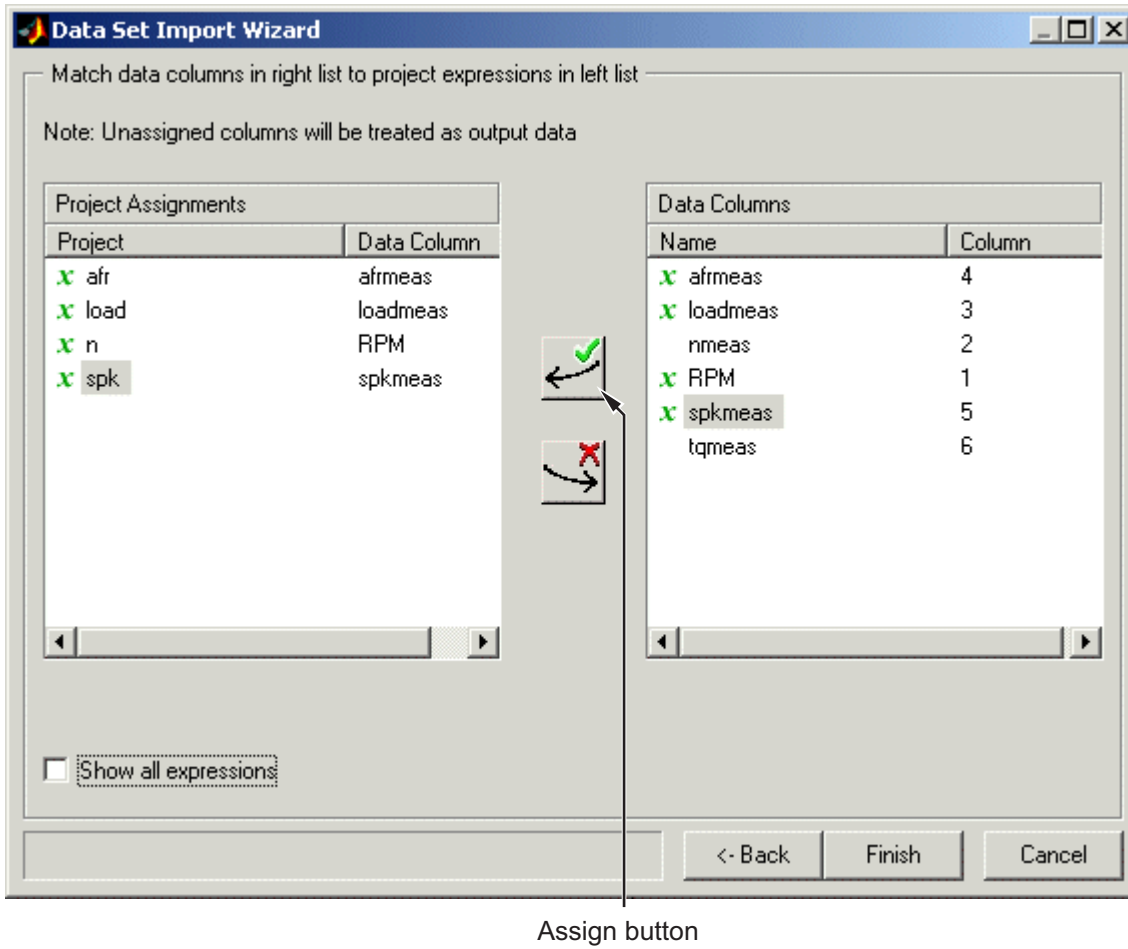

6 Click **Finish** to close the dialog box.

Note If you need to reassign any inputs after closing this dialog you can click in the toolbar or select **Data > Assign**.

Adding an Item to a Data Set

To add the `Torque` feature to the data set,

1. Highlight the `Torque` feature in the lower list of **Project Expressions**.
2. Select **Data > Factors > Add to Data Set**.

This adds two objects to the data set: `Torque: Model` and `Torque: Strategy`. These two objects make up the `Torque` feature.

- `Torque: Model` is the model used as a reference point to calibrate the feature.
- `Torque: Strategy` is the values of the feature at these operating points.

When these steps are complete, the list of factors includes four input factors and four output factors, as shown.

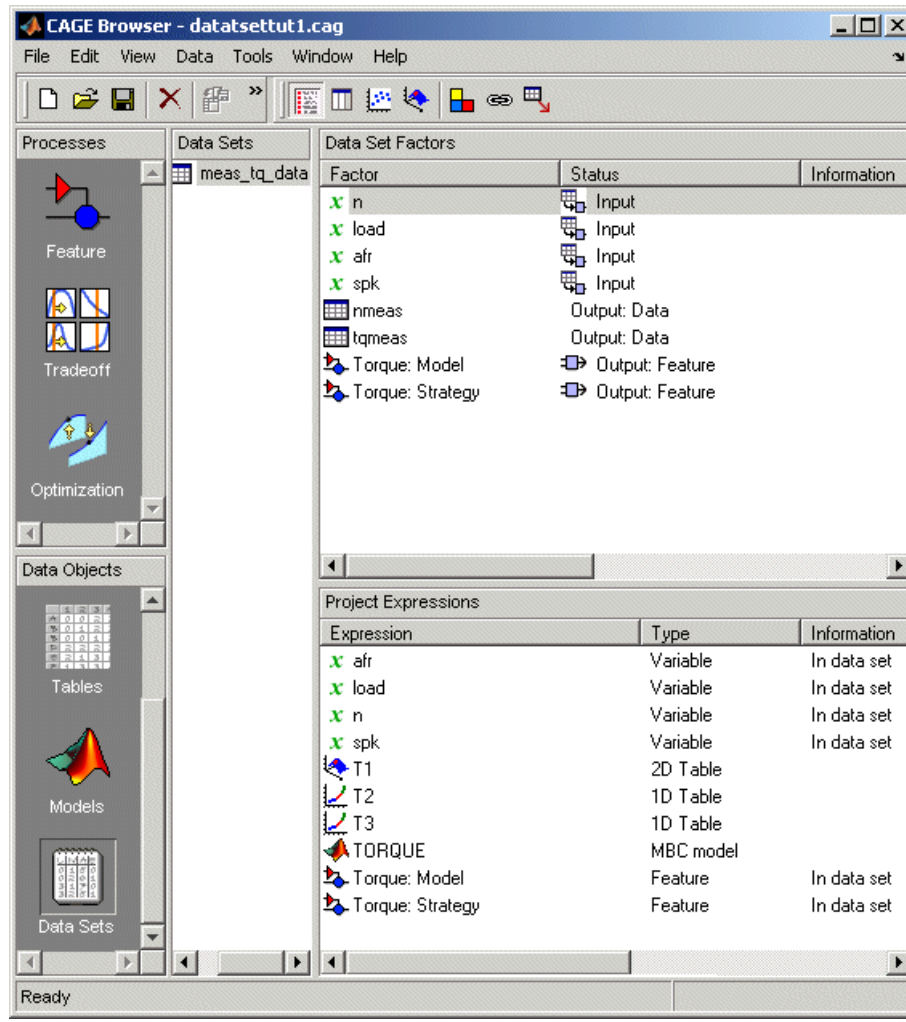
CAGE Browser - datatsettut1.cag
File Edit View Data Tools Window Help
Processes
Feature
Tradeoff
Optimization
Data Objects
Tables
Models
Data Sets
Data Sets
meas_tq_data
Data Set Factors
Factor
Status
Information
n
Input
load
Input
afr
Input
spk
Input
nmeas
Output: Data
tqmeas
Output: Data
Torque: Model
Output: Feature
Torque: Strategy
Output: Feature
Project Expressions
Expression
Type
Information
afr
Variable
In data set
load
Variable
In data set
n
Variable
In data set
spk
Variable
In data set
T1
2D Table
T2
1D Table
T3
1D Table
TORQUE
MBC model
Torque: Model
Feature
In data set
Torque: Strategy
Feature
In data set
Ready

Comparing the Items in a Data Set

In this section...

"Viewing the Data Set as a Table" on page 10-8
"Viewing the Data Set as a Plot" on page 10-10
"Displaying the Error" on page 10-11
"Coloring the Display" on page 10-13

Viewing the Data Set as a Table

By viewing the data set, you can compare experimental data with calibrations or models in your project.

Click in the toolbar to view the data set as a table of values.

	n	load	afr	spk	nmeas	tqmeas	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

In the table, the input cells are white and the output cells are grey. Select the `Torque: Strategy` column header to see the view shown. The selected column turns blue and the column headers of the strategy's inputs (`n`, `load`,

`afr` and `spk`) turn cream. Column headers are always highlighted in this way when they are associated with the currently selected column (such as model inputs, strategy inputs or linked columns).

In addition to viewing the columns, you can use data sets to create a column that shows the difference between two columns:

1. Select the `tqmeas` and `Torque: Strategy` columns by using **Ctrl**+click.
2. Select **Create Error** from the right-click menu on either column header.

This creates another column that is the difference between `tqmeas` and `Torque: Strategy`. Note that all the columns that are inputs to this new column have highlighted headers.

	n	load	afr	spk	nmeas	tqmeas	Torque: Model	Torque: Strategy	tqmeas_minus_Torque
1	2235	0.549	9.5	0.1	2247	66.7	71.666	66.079	-0.621
2	3591	0.454	13.2	0.1	3613	54.1	47.163	46.891	-7.209
3	4946	0.651	12	0.1	4974	73.7	47.573	79.256	5.556
4	881	0.648	11.9	5.7	881	75.8	99.23	80.211	4.411
5	2234	0.441	13.3	0.1	2247	55.9	51.256	45.152	-10.748
6	3591	0.747	10.9	0.1	3612	90	92.837	105.586	15.586
7	4947	0.541	9.7	0.1	4973	62.8	57.76	57.587	-5.213
8	881	0.622	9.9	0.1	884	72.1	76.198	60.926	-11.174
9	1219	0.333	14	0.1	1224	41.8	33.226	21.318	-20.482
10	1558	0.382	12	0.1	1567	49.4	40.487	31.957	-17.443
11	1896	0.209	10.7	3.3	1906	28.5	3.492	4.197	-24.303
12	2234	0.284	9.8	3.2	2245	36	23.063	19.891	-16.109
13	2574	0.407	13.4	3	2588	49.9	49.629	44.794	-5.106
14	2914	0.595	11.5	3.1	2929	70.5	84.68	82.229	11.729
15	3251	0.781	12.3	3.1	3268	90.5	117.424	117.259	26.759
16	3589	0.668	13.5	3	3608	77.1	87.987	96.408	19.308
17	3930	0.452	11.9	3.1	3952	52.7	46.511	51.722	-0.978
18	4268	0.235	10.9	3	4293	27.7	5.253	3.085	-24.615
19	4606	0.194	12	3.2	4633	21.3	-2.088	-5.771	-27.071

The error column is simply the difference between `tqmeas` and `Torque: Strategy`. This provides a simple way of comparing the feature and the measured data.

Viewing the Data Set as a Plot

1 Click or select **View > Plot** to view the data set as a plot.

The lower pane lists all the output expressions in the data set and in the project.

2 Use **Ctrl**+click to select `tqmeas` and `Torque: Strategy` from the lower list.

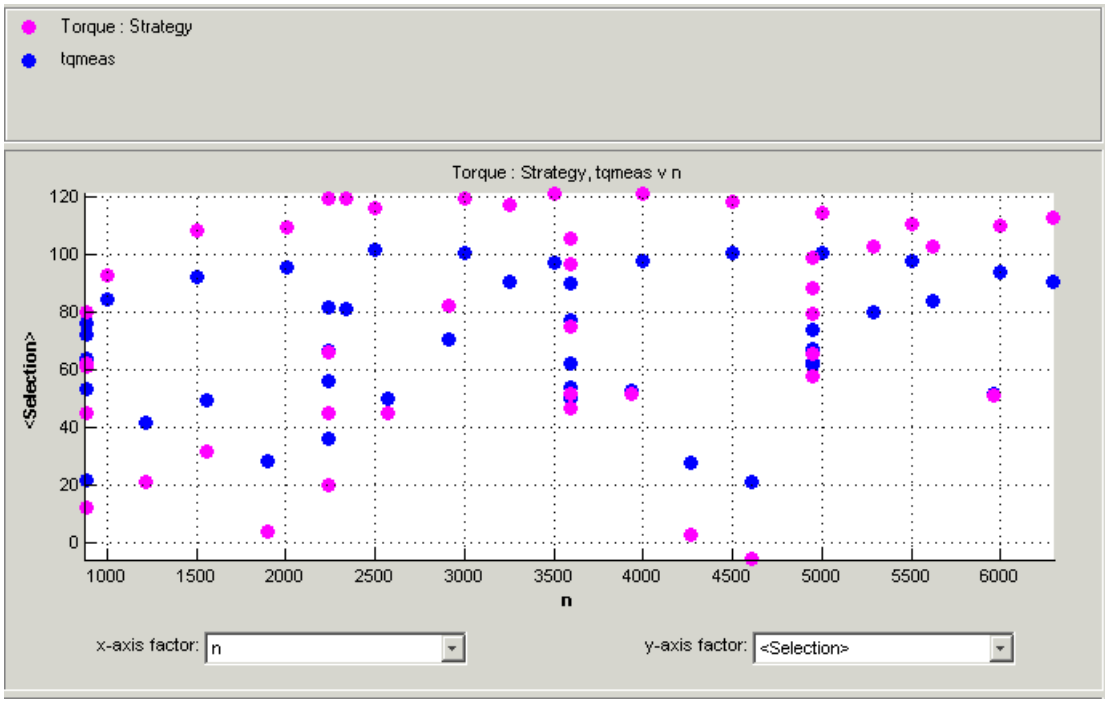

3 Change the **x-axis factor** to `n` from the drop-down menu.

This displays the calibrated values of torque from the feature, and the measured values of torque from the experimental data, against the test cell settings for engine speed.

Clearly there is some discrepancy between the two.

Displaying the Error

View the error between the calibrated and measured values of torque.

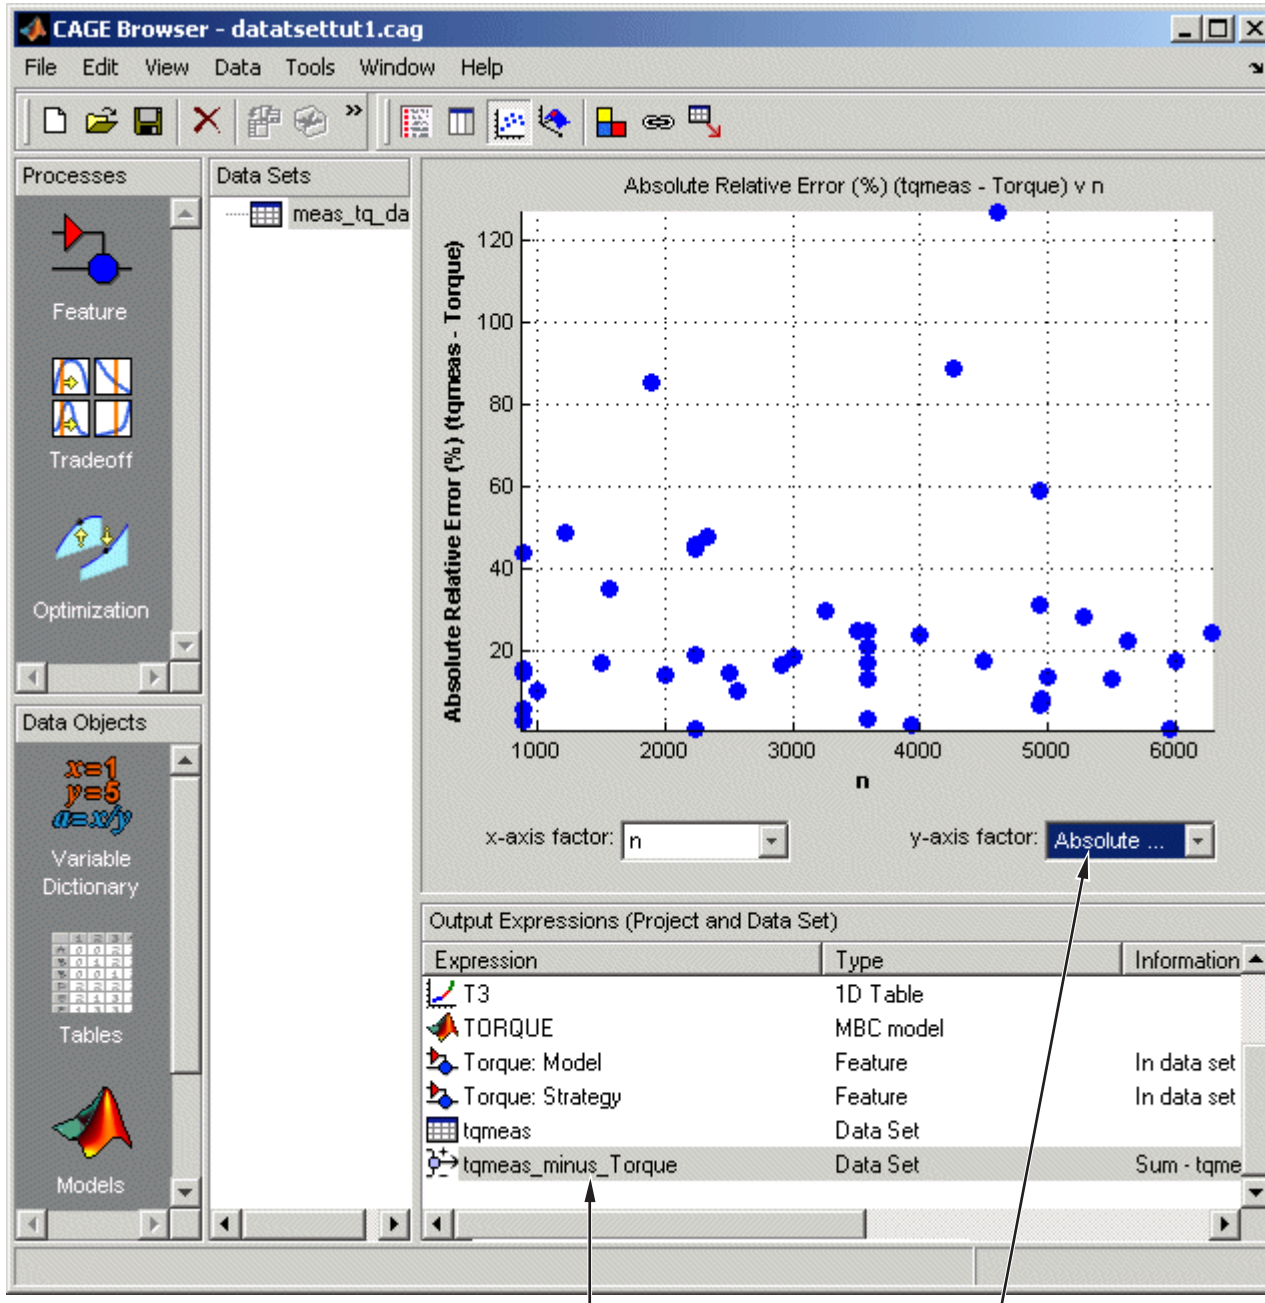

1. Select tqmeas_minus_Torque.

2. Select Absolute Relative Error (tqmeas - Torque).

1. Select `tqmeas_minus_Torque` from the lower list (**Output Expressions**).
2. For the **y-axis factor**, select `Absolute Relative Error (tqmeas - Torque)` from the drop-down menu.

As you can see, there seems to be no particular correlation between engine speed and the error in the calibration.

Coloring the Display

1. Select **Color by Value** from the right-click menu on the graph.
2. From the **Color by** drop-down menu, select `load`.

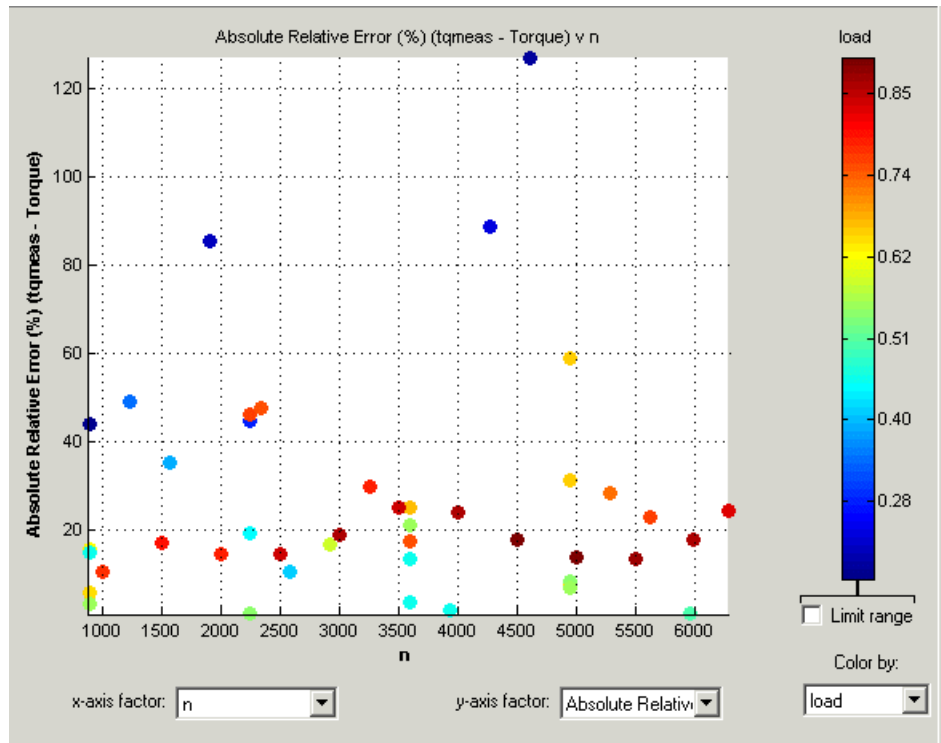

In this display, you can see that some of the low values of load display a high error.

Limiting the Range of the Colors

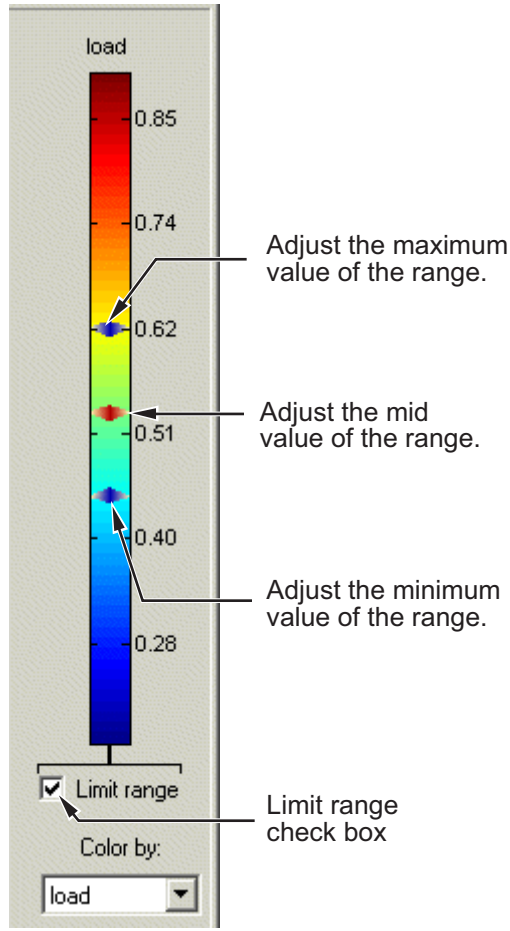

To view the colors in more detail, you can limit the range of the colors:

1 Select the **Limit range** box (or you could right-click the graph and select **Restrict Color to Limits**).

2 Set the minimum value of the color range to be as low as possible by dragging the minimum value down.

3 Set the maximum value of the color range to be around 0.4.

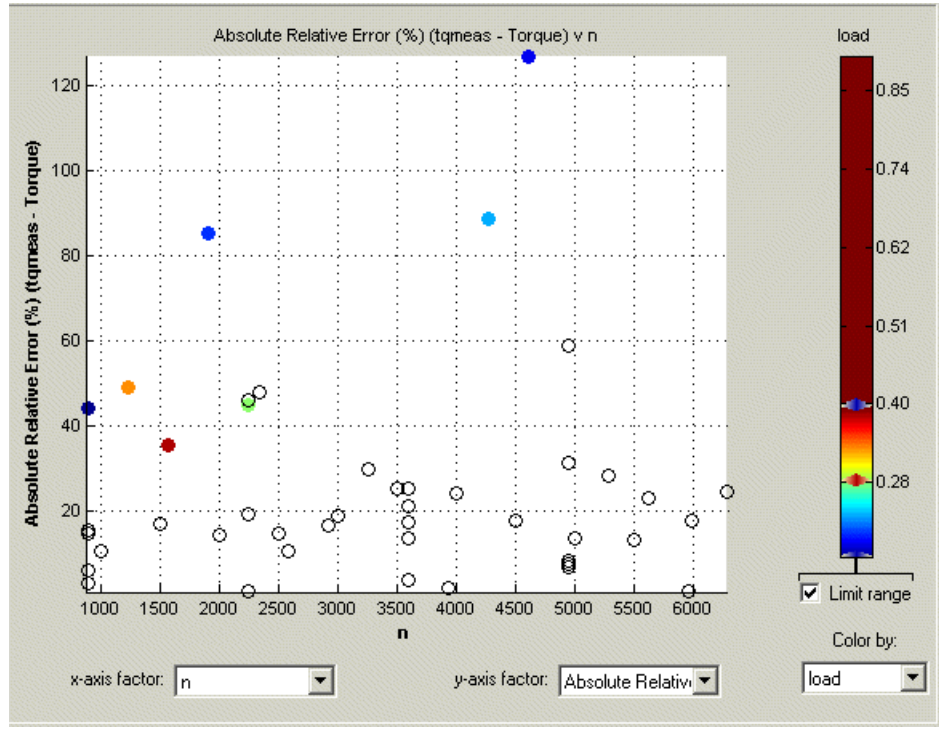

As the low values of load are causing large errors, it would be wise to reexamine the calibration, particularly at small values of `load`.

Reassigning Variables

Instead of using the test cell settings for the engine speed (`RPM`), you might want to use the measured values of engine speed (`nmeas`). So you have to reassign the variable `n` to `nmeas`.

To reassign `n`,

1. Click or select **Data > Assign**.
2. In the dialog that appears, select `n` from the **Project Assignments** pane and `nmeas` from the **Data Columns** pane.
3. Click the assign button.

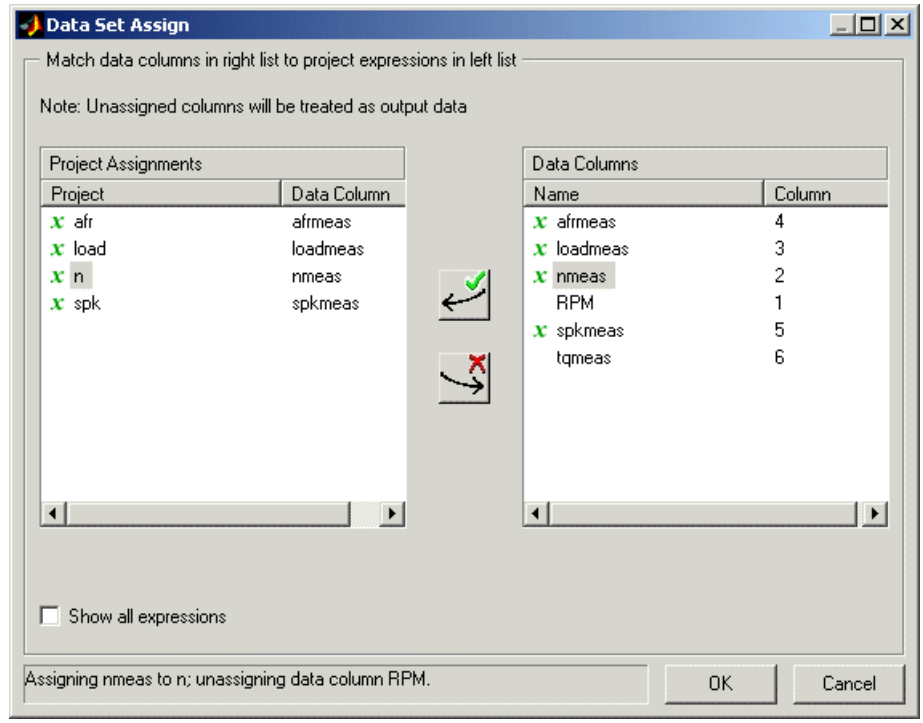

You can now compare your calibration with your experimental data again, using the techniques described.

For more information about the complete functionality of data sets, see “Data Sets” in the CAGE documentation.

You have now completed the data sets tutorial.

11

Tutorial: Filling Tables from Data

This section includes the following topics:

- "Setting Up a Table and Experimental Data" on page 11-2
- "Filling the Table from the Experimental Data" on page 11-9
- "Selecting Regions of the Data" on page 11-13
- "Exporting the Calibration" on page 11-15

Setting Up a Table and Experimental Data

In this section...

"Tutorial Overview" on page 11-2

"Adding Variables" on page 11-3

"Adding a New Table" on page 11-4

"Importing Experimental Data" on page 11-6

Tutorial Overview

If you are considering a straightforward strategy, you might want to fill tables directly from experimental data. For example, a simple torque strategy fills a lookup table with values of torque over a range of speed and relative air charge, or load. You can use CAGE to fill this strategy (which is a set of tables) by referring to a set of experimental data. You can also fill tables with the output of optimizations.

This tutorial takes you through the steps of calibrating a lookup table for torque, based on experimental data.

- This section describes the steps required to set up CAGE in order to calibrate a table by reference to a set of data.
- "Filling the Table from the Experimental Data" on page 11-9 describes the process of filling the lookup table.
- "Selecting Regions of the Data" on page 11-13 describes how you can select some of the data for inclusion when you fill the table.
- "Exporting the Calibration" on page 11-15 describes how to export your completed calibration.

Start CAGE by typing

```
cage
```

at the MATLAB prompt.

If you already have a CAGE session open, select **File > New Project**.

First you will set up a blank table ready for filling using experimental data or optimization output.

The steps that you need to follow to set up the CAGE session are

1. Add the variables for speed and load by importing a variable dictionary.
2. Add a new table to your session.
3. Import your experimental data.

The next sections describe each of these processes in detail.

Adding Variables

Before you can add tables to your session, you must add variables to associate with the normalizers or axes.

To add a variable dictionary,

1. Select **File > Import > Variable Dictionary**.
2. Select `table_filling_tutorial.xml` from the `matlab\toolbox\mbc\mbctraining` directory.

This loads a variable dictionary into your session. The variable dictionary includes the following:

- `N`, the engine speed
- `L`, the relative air charge
- `A`, the air/fuel ratio (AFR)
- `stoich`, the stoichiometric constant

You can now add a table to your session.

Adding a New Table

You must add a table to fill.

To add a new table,

1. Select **File > New > 2D table**.

   This opens a dialog box that asks you to specify the variable names for the normalizers. As you can see in the dialog controls, accepting the defaults will create a table with ten rows and ten columns with an initial value of `0` in each cell.

2. Change the number of columns to `7`.

3. Select `L` as the variable for normalizer **Y** and `N` as the variable for normalizer **X**, then click **OK**.

Note In CAGE, a 2-D table is defined as a table with two inputs.

CAGE takes you to the **Tables** view, where you can see the following.

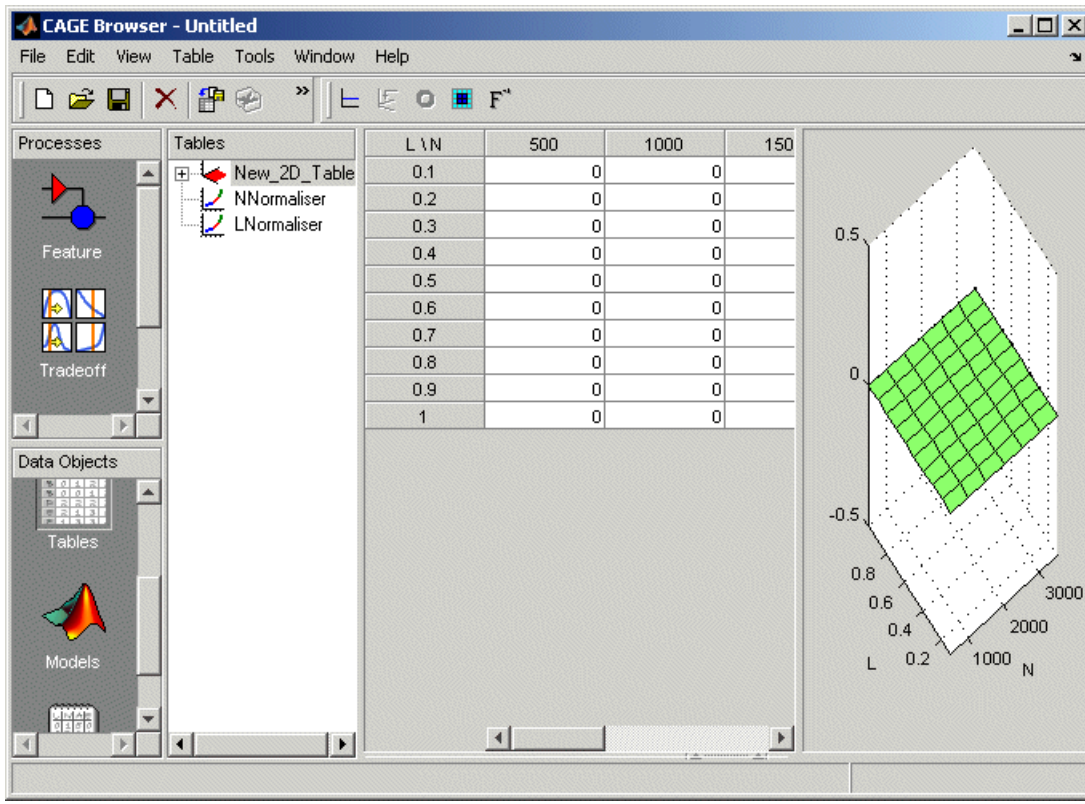

Inspecting the Values of the Normalizers

CAGE has automatically initialized the normalizers by spacing the breakpoints evenly across the range of values for the engine speed (N) and load (L). The variable ranges are found in the variable dictionary. Switch to the Normalizer view to inspect the normalizers.

Expand the table branch by clicking ⊞, and select `NNormalizer` as shown.

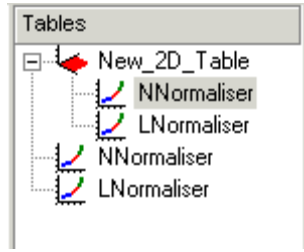

This displays the two normalizers for the table.

You have an empty table with breakpoints over the ranges of the engine speed and load, which you can fill with values based on experimental data.

Importing Experimental Data

To fill a table with values based on experimental data, you must add the data to your session. If you want to fill a table with the output of an optimization, the output appears automatically in the Data Sets view as a new data set called `Exported_Optimization_Data` when you select the Export to Data Set toolbar button. For this tutorial you need to import some experimental data.

CAGE uses the **Data Sets** view to store grids of data. Thus, you need to add a data set to your session as well.

Select **File > New > Data Set** to add a data set to your session. This changes the view to the **Data Set** view.

You can now import experimental data into the data set:

1. Select **File > Import > Data**.
2. In the file browser, select `meas_tq_data.csv` from the `matlab\toolbox\mbc\mbctraining` directory and click **Open**.

   This set of data includes six columns of data: the test cell settings for engine speed (`RPM`), and the measured values of torque (`tqmeas`), engine speed (`nmeas`), air/fuel ratio (`afrmeas`), spark angle (`spkmeas`), and load (`loadmeas`).

3 This opens the Data Set Import Wizard. The first screen asks which of the columns of data you want to import. Click **Next** to import them all.

   The following screen asks you to associate variables in your project with data columns in the data.

4 Highlight `N` in the **Project Assignments** column and `nmeas` in the **Data Column**, then click the assign button, shown.

5 Repeat this to associate `L` with `loadmeas`. The dialog box should be the same as the following.

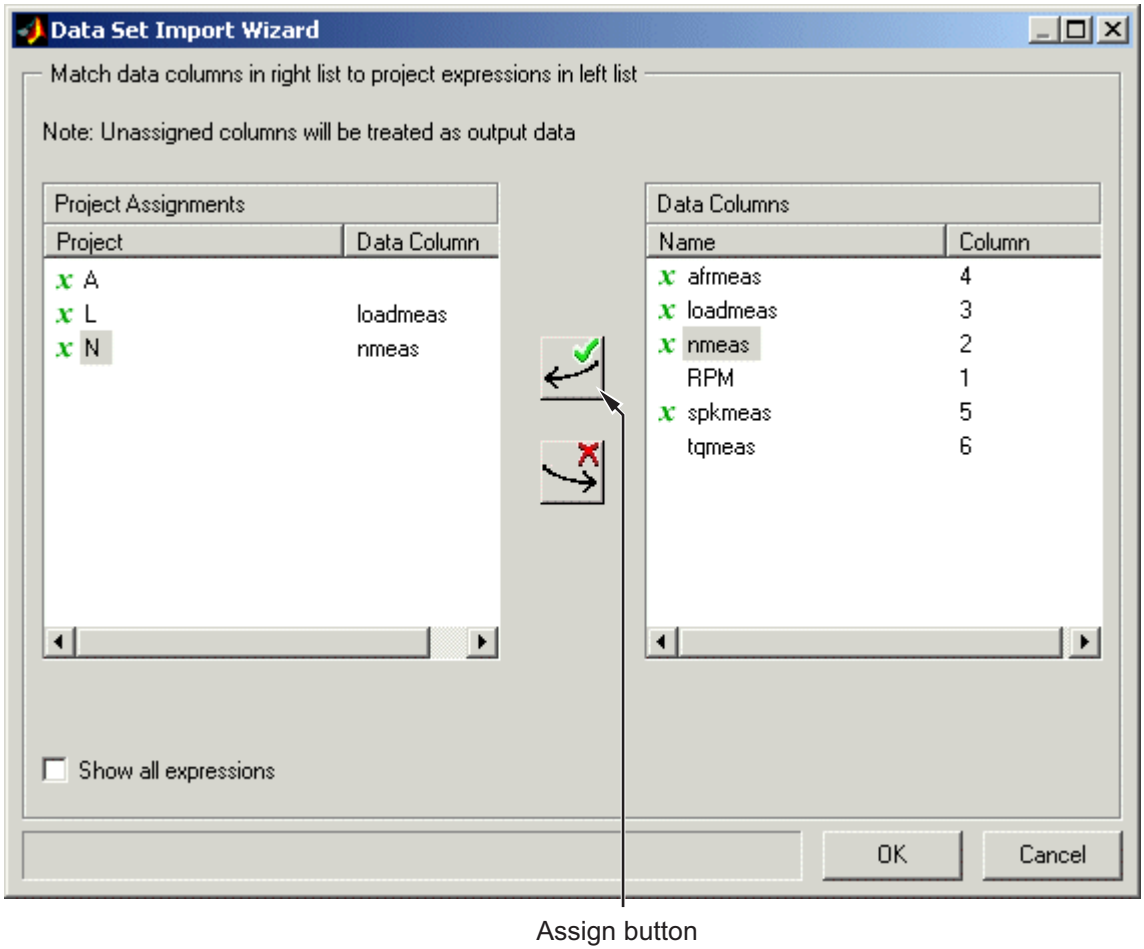

6 Click **Finish** to close the dialog box.

You now have an empty table and some experimental data in your session. You are ready to fill the table with values based on this data.

Filling the Table from the Experimental Data

You have an empty table and the experimental data in your session. You can now fill the table with values based on your data.

The data that you have imported is a series of measured values of torque at a selection of different operating points. These operating points do not correspond to the values of the breakpoints that you have specified. The lookup table has a range of engine speed from 500 revolutions per minute (rpm) to 3500 rpm. The range of the experimental data is far greater.

CAGE extrapolates the values of the experimental data over the range of your table. Then it fills the table by selecting the torque values of the extrapolation at your breakpoints.

To fill the table with values based on the experimental data,

1. To view the **Table Filler** display, click (Fill Table From Data Set) in the toolbar in the **Data Sets** view; or select **View > Table Filler**.

   You can use this display to specify the table you want to fill and the factor you want to use to fill it.

2. In the lower pane, select `New_2D_Table` from the **Table to fill** list.

3. Select `tqmeas` from the **Factor to fill table** list. This is the data that you want to use to fill the table.

4. Select `N` from the **x-axis factor** list and `L` from the **y-axis factor** list. Your session should be similar to the following display.

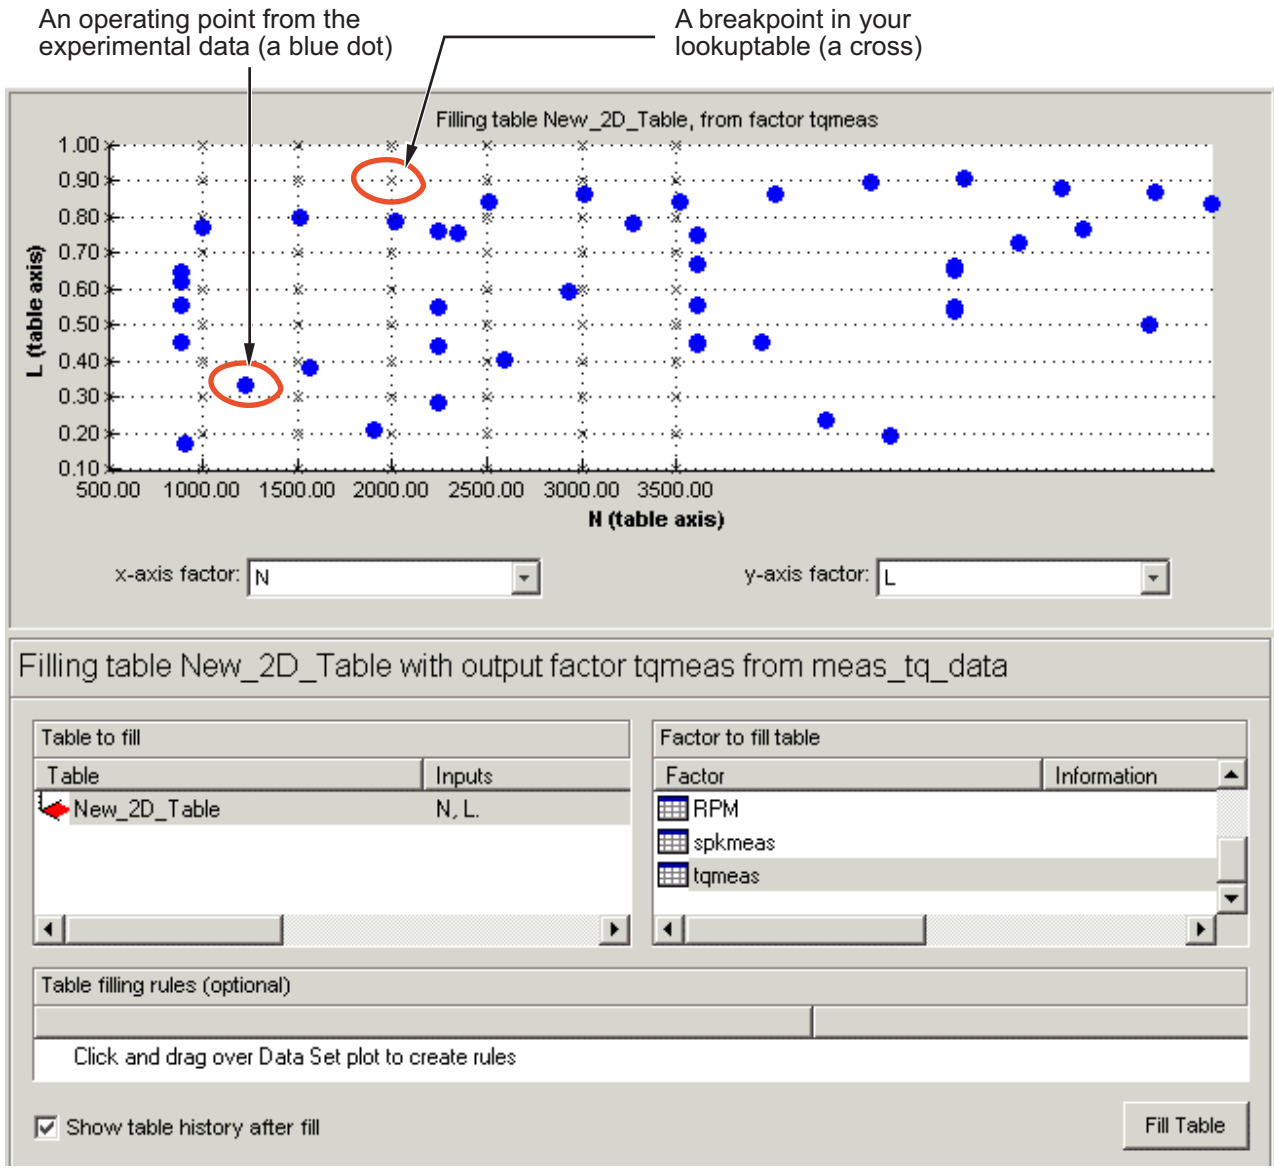

The upper pane displays the breakpoints of your table as crosses and the operating points where there is data as blue dots. Data sets display the

points in the experimental data, not the values at the breakpoints. You can inspect the spread of the data compared to the breakpoints of your table before you fill the table.

5 To view the table after it is filled, ensure that the **Show table history after fill** box, at the bottom left, is selected.

6 To fill the table with values of `tqmeas` extrapolated over the range of the normalizers, click **Fill Table**.

This opens the History dialog box, shown.

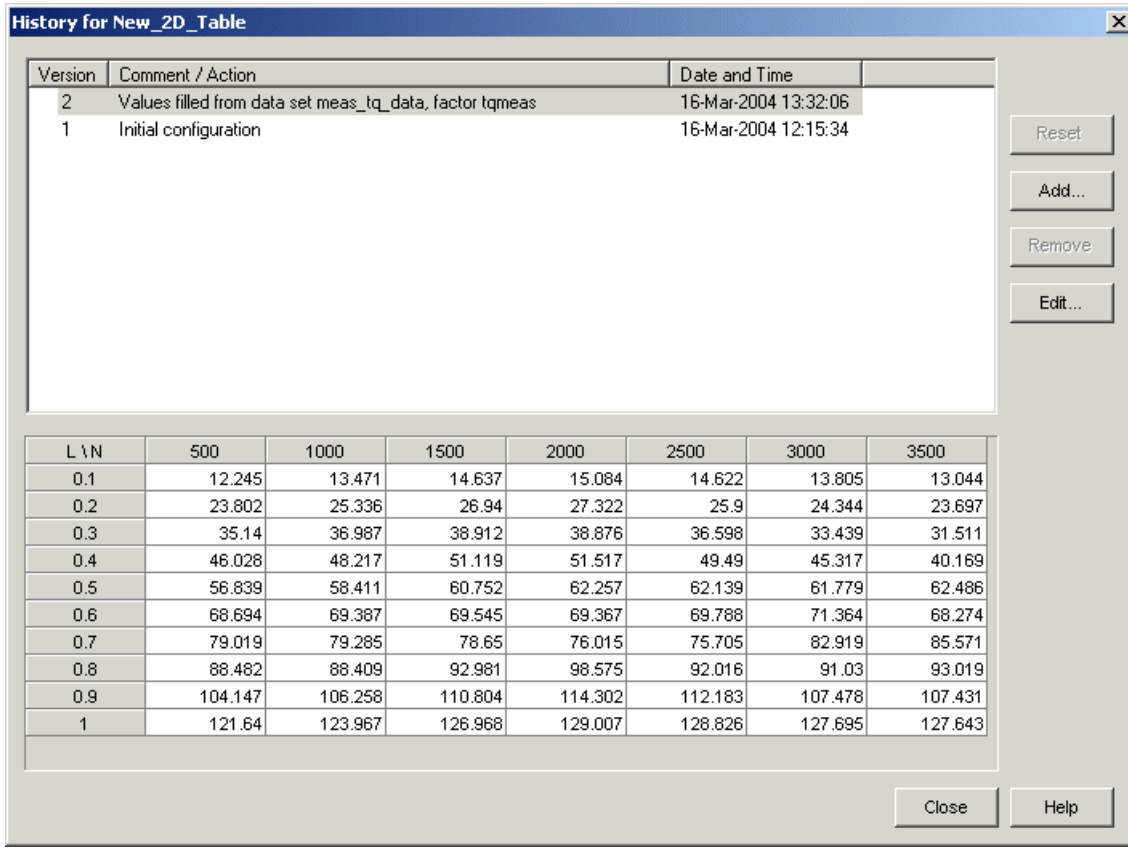

7 Click **Close** to close the History dialog box and return to the **Table Filler** display.

8 To view the graph of your table, as shown, select **Data > Plot > Surface**.

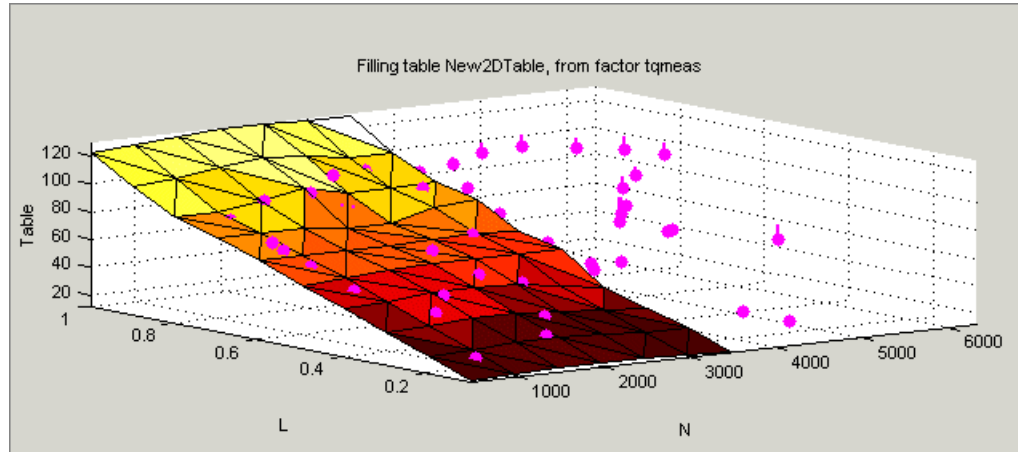

This display shows the table filled with the experimental points overlaid as purple dots.

The table has been calibrated by extrapolating over the values of your data and filling the values that the data predicts at your breakpoints.

Notice that the range of the table is smaller than the range of the data, as the table only has a range from 500 rpm to 3500 rpm.

The data outside the range of the table affects the values that the table is filled with. You can exclude the points outside the range of the table so that only points in the range that you are interested in affect the values in the table.

Selecting Regions of the Data

You can ignore points in the data set when you fill your lookup table.

For example, in this tutorial the experimental data ranges over values that are not included in the lookup table. You want to ignore the values of engine speed that are greater than the range of the table.

To ignore points in the data set,

1. Select **Data > Plot > Data Set**. This returns you to the view of where the breakpoints lie in relation to the experimental data.

2. To define the region that you want to include, left-click and drag the plot. Highlight all the points that are included in your table range, as shown.

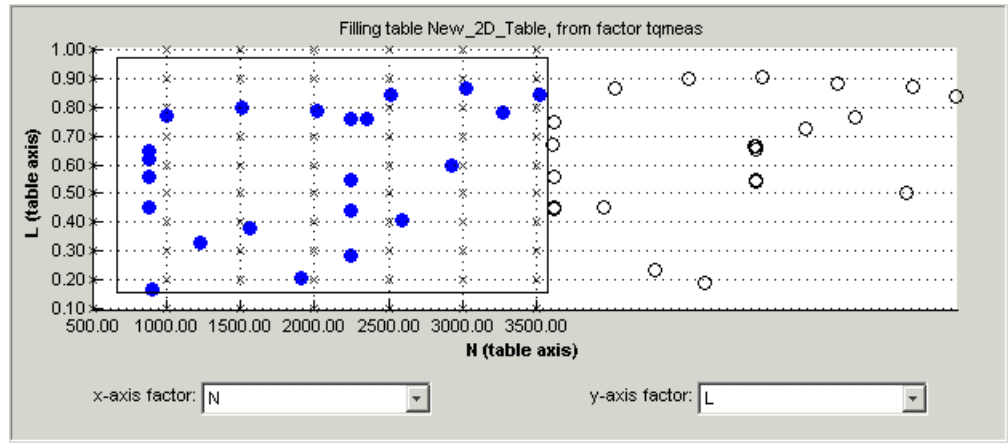

3. To fill the table based on an extrapolation over these data points only, click **Fill Table**. This opens the History display again.

4. In the History display, select version 2 and 3, using **Ctrl**+click. The following display shows a comparison between the table filled with two different extrapolations.

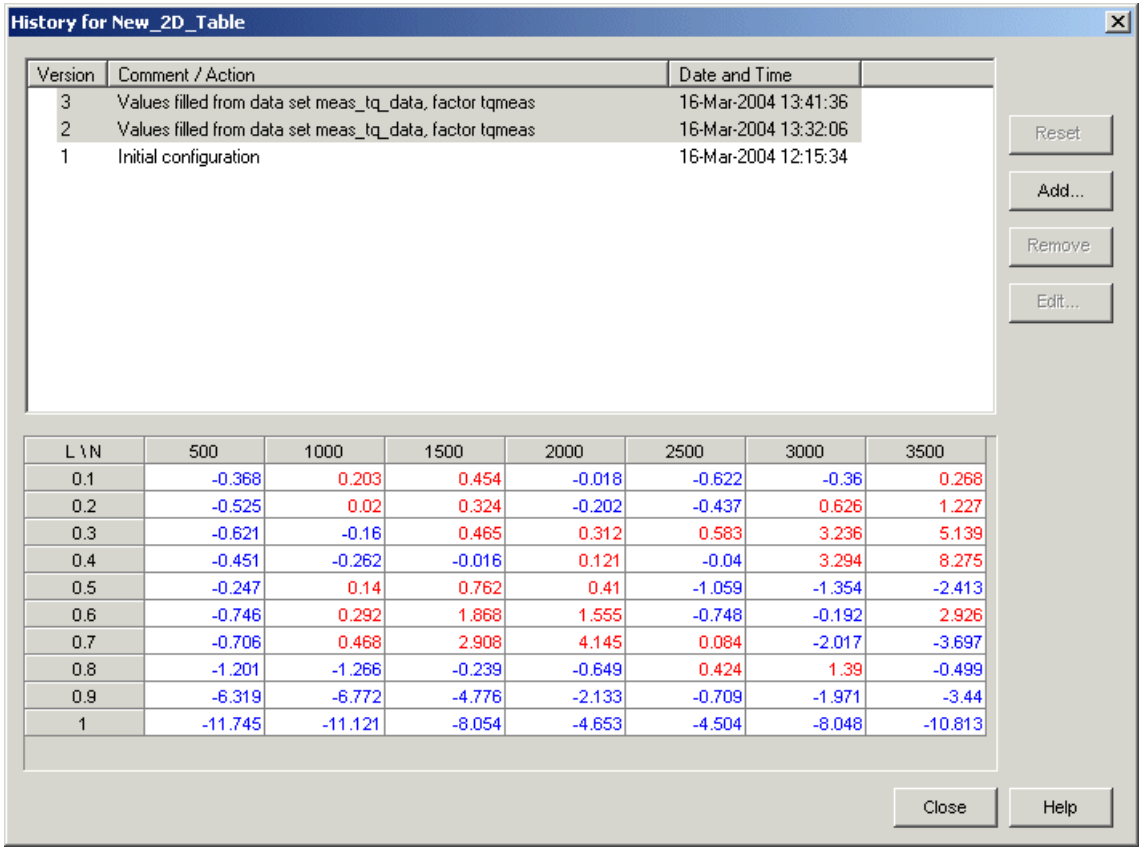

5 Click **Close** to close the History viewer.

6 Select **Data > Plot > Surface** to view the surface again.

The display of the surface now shows the table filled only by reference to the data points that are included in the range of the table.

You have filled a lookup table with values taken from experimental data.

Exporting the Calibration

To export the calibration,

1. To highlight the table that you want to export, you must first click **Tables**, shown.

2. Highlight the `New_2D_Table`.

3. Select **File > Export > Calibration > Selected Item**.

4. Choose the type of file you want to save your calibrations as. For the purposes of this tutorial, select `Comma Separated Value (.csv)`.

5. Enter `table_filling_tutorial.csv` as the file name and click **Save**.

This exports the calibration.

You have now completed this tutorial.

12

Tutorial: Optimization and Automated Tradeoff

This section includes the following topics:

- "Tutorial Overview" on page 12-2
- "Single-Objective Optimization" on page 12-5
- "Multiobjective Optimization" on page 12-19
- "Sum Optimization" on page 12-33
- "Automated Tradeoff" on page 12-38
- "Worked Example Optimization" on page 12-41
- "Creating an Optimization from Your Own Algorithm" on page 12-48

Tutorial Overview

In this section...

"Tutorial Problem Definitions" on page 12-2

"The Optimization View" on page 12-2

"Setting Up Your Optimization Session" on page 12-4

Tutorial Problem Definitions

In this tutorial you will use optimization to find solutions to the following problems:

- A single-objective optimization to find maximum values of torque, subject to a constraint to keep NOX emissions below a specified level. You will export the output and use it to fill a table. See "Single-Objective Optimization" on page 12-5.
- A multiobjective optimization to maximize torque and minimize NOX emissions. See "Multiobjective Optimization" on page 12-19.
- A sum optimization to maximize torque while minimizing NOX, weighted to give more importance to idle speed. See "Sum Optimization" on page 12-33.
- Using any of your optimizations to run an automated tradeoff. Once you have set up an optimization you can apply it to a tradeoff. See "Automated Tradeoff" on page 12-38.

The Optimization View

You can use the **Optimization** view to set up, run, view, and export optimizations. You must also set up optimizations here in order to use them for automated tradeoff.

Start the CAGE Browser part of the Model-Based Calibration Toolbox product by typing

```
cage
```

at the MATLAB prompt.

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

When you first open the **Optimization** view both panes are blank until you create an optimization. After you set up your optimizations, the left **Optimization** pane shows a tree hierarchy of your optimizations, and the right hand panes display details of the optimization selected in the tree, as with other CAGE processes.

As for other CAGE processes, you must set up your session for an optimization. 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. In this case you also need to set up this set of points using the **Data Sets** view. The steps required are as follows:

1. Import a model or models.
2. Set up a new optimization.
3. Define an operating point set if required.

The following tutorial guides you through this process to evaluate this optimization problem:

MaxTQ (SPK, N, L)

That is, find the maximum of the torque model (TQ) as a function of spark (SPK), engine speed (N), and load (L). You will use the `NOXFLOW` model to constrain these optimization problems.

Setting Up Your Optimization Session

Before you can set up the optimization, you must set up your session.

1. Select **File > Open Project** (or the toolbar button) to choose the `tradeoffInit.cag` file, found in the `matlab\toolbox\mbc\mbctraining` directory, then click **OK**.

   The `tradeoffInit.cag` project contains two models and all the variables necessary for this tutorial. For more information about how to set up models and variables, see "Variables and Models" in the CAGE documentation.

2. Select the Models view by clicking the **Models** button in the **Data Objects** pane.

   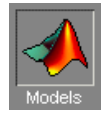

   Observe that the project you have opened contains two models: `TQ_Model` and `NOXFLOW_Model`. In this tutorial you use these models to optimize torque values subject to emissions constraints.

3. To view the items in the Variable Dictionary, click the **Variable Dictionary** button in the **Data Objects** pane.

   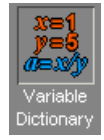

   The **Variable Dictionary** view appears, displaying the variables, constants, and formulas in the current project. The project already has the relevant variables defined, so you do not need to import a variable dictionary. Note that the variables have ranges and set points defined.

Single-Objective Optimization

In this section...
"Process Overview" on page 12-5
"Using the Optimization Wizard" on page 12-5
"Setting Objectives, Constraints, and Operating Points" on page 12-8
"Running the Optimization" on page 12-13
"Using Optimization Results to Fill Tables" on page 12-14
"Using a Custom Fill Routine to Fill Tables" on page 12-17

Process Overview

The following sections describe these stages:

1. Using the Optimization Wizard to choose
   - Your optimization algorithm
   - How many objectives and constraints
   - What free variables to use
2. Using the Optimization view to choose
   - A model for your objective
   - A model, type, and value for your constraint
   - An operating point set for your optimization
3. Running the optimization, examining the output, exporting to a data set, and using the output to fill a table

Using the Optimization Wizard

To create a new optimization,

1. Select **File > New > Optimization**.

   This opens the Optimization Wizard.

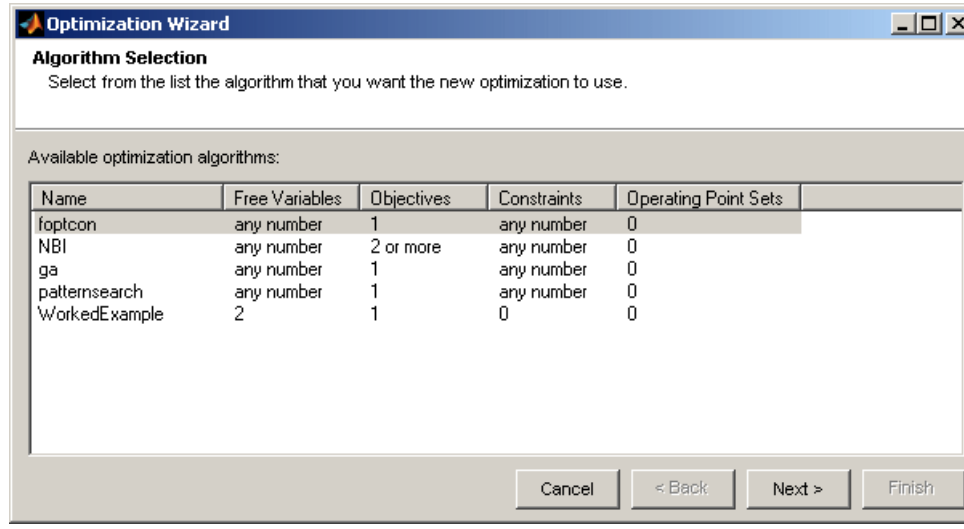

2 `foptcon` is selected by default, and this is the optimization algorithm you will use for this example. Note that this algorithm specifies a single objective in the **Objectives** column. Click **Next**.

3 On the next screen, set the number of constraints to 1, as shown.

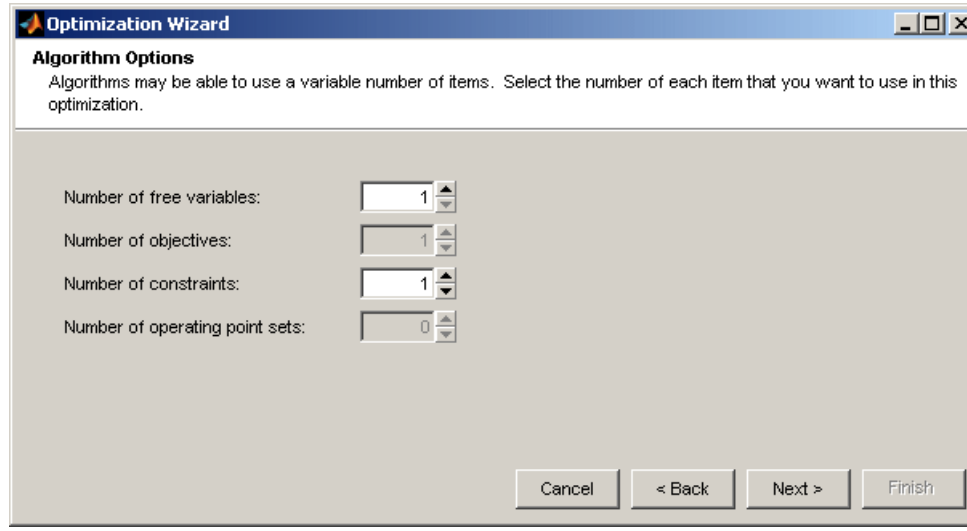

Leave the number of free variables at `1` (spark will be the free variable). Click **Next**.

4 On the next screen, choose spark as your free variable for this optimization by clicking `SPK` in the list on the right, then click the button to match it up with `FreeVariable1`, as shown. Click **Finish**.

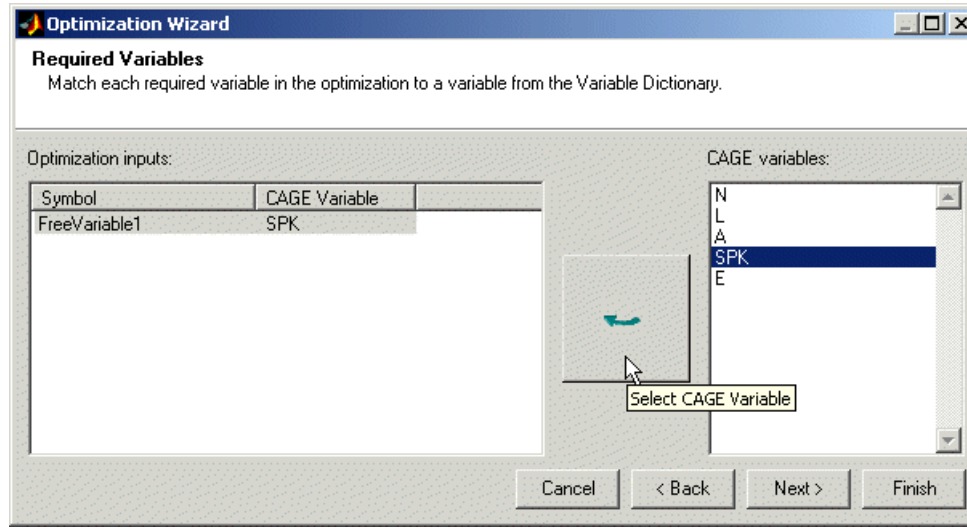

A new branch named `Optimization` appears in the Optimization tree. Your CAGE browser should look like the following example. In the **Optimization Information** pane you can see listed the algorithm name `mbcOSfmincon`, free variable `SPK`, and the description `Single objective optimization subject to constraints`. In the **Objectives** and **Constraints** panes there are status messages informing you that you need to specify a model for an objective, and valid constraint settings.

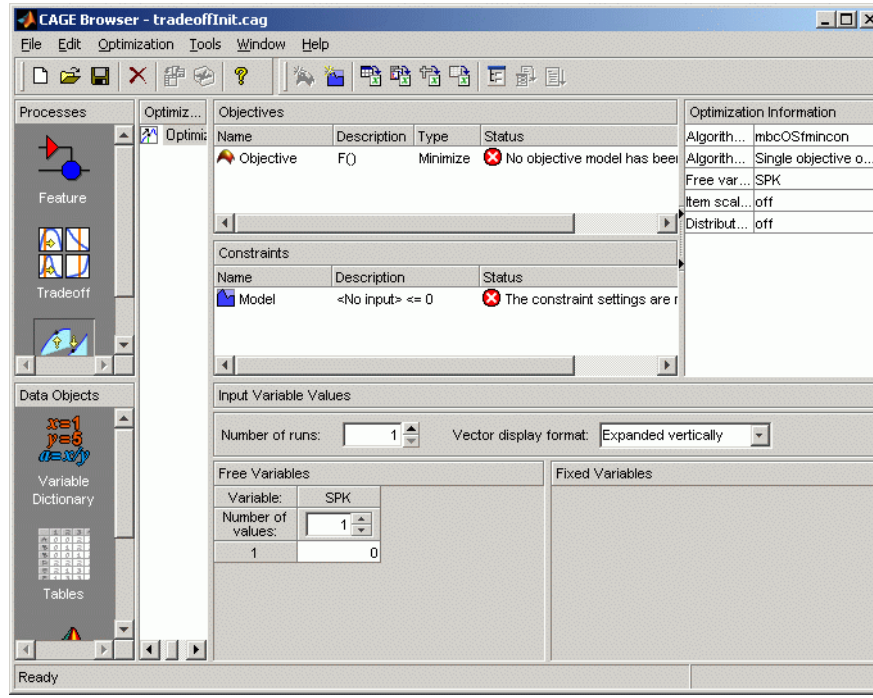

Setting Objectives, Constraints, and Operating Points

1 Double-click `Objective` in the **Objectives** pane.

The Edit Objective dialog appears.

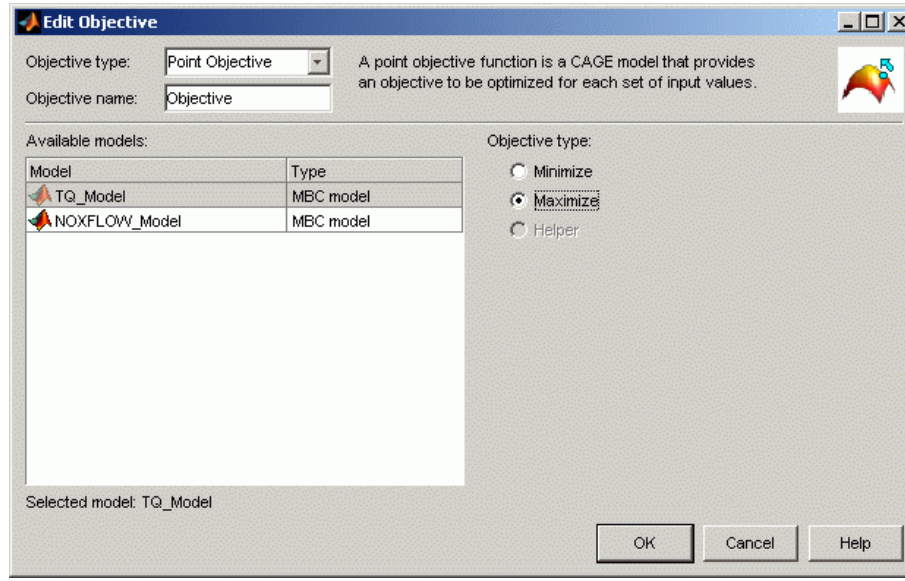

2 Click to select `TQ_Model` and select **Maximize** from the radio buttons on the right. Click **OK**.

You return to the CAGE Browser Optimization view. The **Description** `TQ_Model(SPK,L,N,A,E)` appears in the **Objectives** pane.

3 Double-click `Model` in the **Constraints** pane.

The Edit Constraint dialog appears.

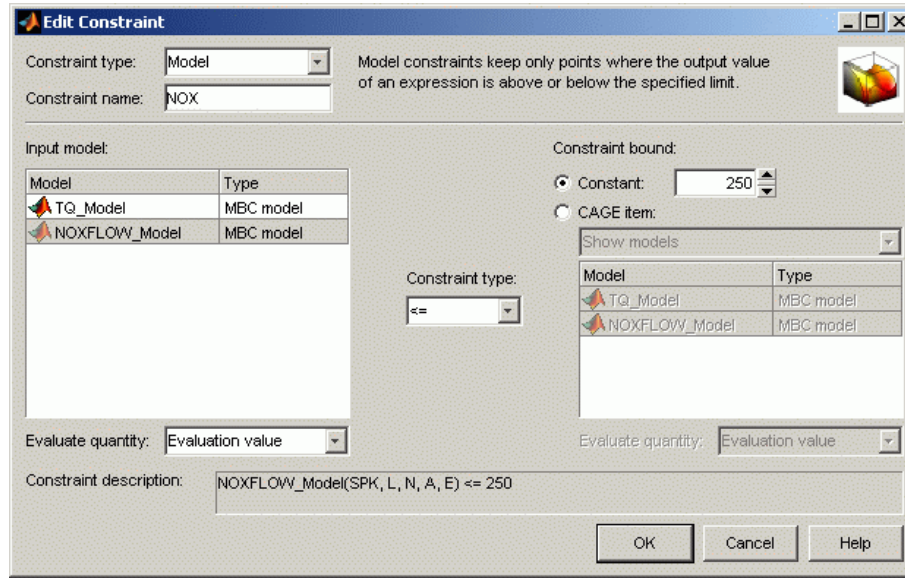

a Leave the **Constraint type** drop-down menu at the default, `Model`.

b Edit the **Constraint name** to `NOX`.

c Select `NOXFLOW_Model` from the **Input model** list.

d Make sure the inequality is `<=`, and enter `250` in the **Constant** edit box as the maximum value for the constraint, as shown above.

e Press **Enter**.

You return to the CAGE Browser Optimization view. Notice that the **Description** `NOXFLOW_Model(SPK, L, N, A, E) <= 250` appears in the **Constraints** pane.

Note that the toolbar button Run Optimization ( ) is now enabled, because your optimization setup has provided enough information to start an optimization.

4 You can use the **Input Variable Values** panes to define a set of operating points for the optimization. Note that you do not have to have an operating point set; if you do not, the optimization will run at a single point of your choosing (the set points of variables is the default).

Running the optimization requires the selected models to be evaluated (many times over) and hence values are required for all the model input factors (L, N, A, E, and SPK). The defaults of the fixed variables (L, N, A, E) are their set points, as shown in the **Fixed Variables** pane. You have chosen SPK as a free variable, so the optimization will choose different values for SPK in trying to find the best. The default initial value for a free variable is the set point, as shown in the **Free Variables** pane.

To define the set of operating points for the optimization,

a In the **Input Variable Values** pane, increase the **Number of runs** to 6. Notice 6 rows appear in both fixed and free variables panes, all containing the default set point values of each variable.

b Enter, or copy and paste, these values into the N column of the **Fixed Variables** pane:

N
1000
1000
3000
3000
6000
6000

c Enter, or copy and paste, these values into the L column of the **Fixed Variables** pane:

L
0.1
0.8
0.1
0.8
0.1
0.8

The **Fixed Variables** pane should look as shown.

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

Leave the other fixed variables and the free variable values at the defaults. If you wished to restrict the range of the free variables, you could select **Optimization > Edit Free Variable Ranges**. The default is the range of the variable as defined in the Variable Dictionary. For this example, leave the default.

5 Your CAGE Browser should now look like the following example, with an objective, constraint, and set of operating points. The optimization is ready to run.

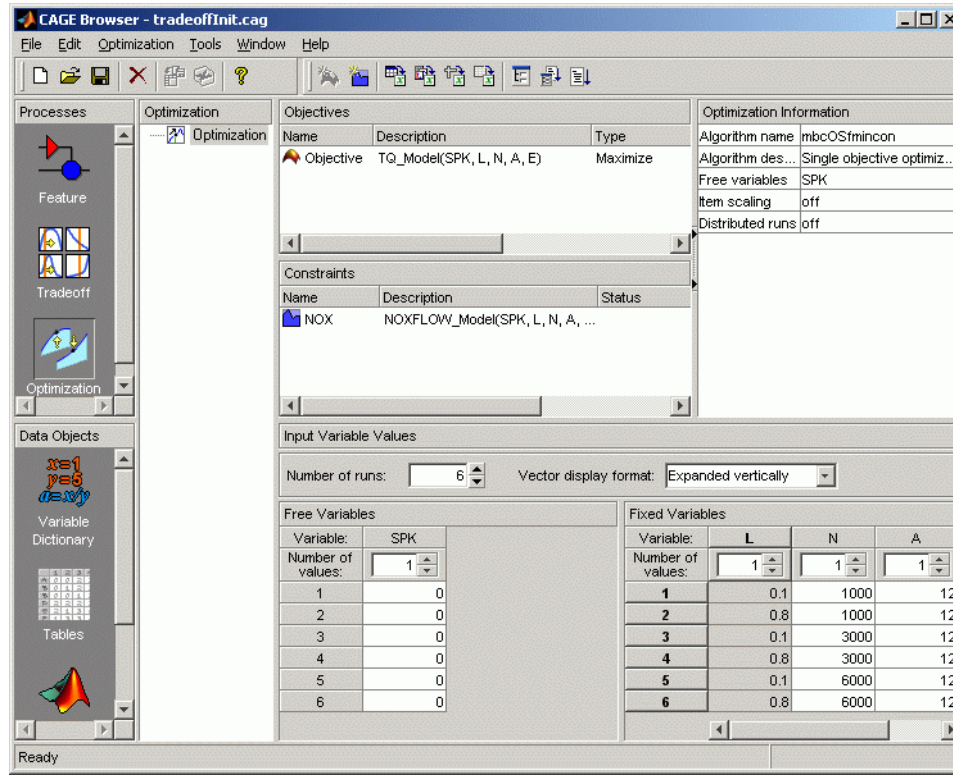

Running the Optimization

1 Click Run Optimization ( ) in the toolbar.

The optimization runs, showing progress messages as each point is evaluated until the optimization is complete. On completion of the optimization, a new node appears in the Optimization tree.

2 The view switches to the new node `Optimization_Output` where you can view the optimization results.

This single-objective optimization produces one best solution for each point in the operating point set. Click the cells of the table, or the points in the Results Surface, to view solutions at those operating points.

3 The optimization output view retains a memory of previous layout. If you have not used these views before, try the buttons and right-click context menus in the view title bars to add constraint graphs to examine your results.

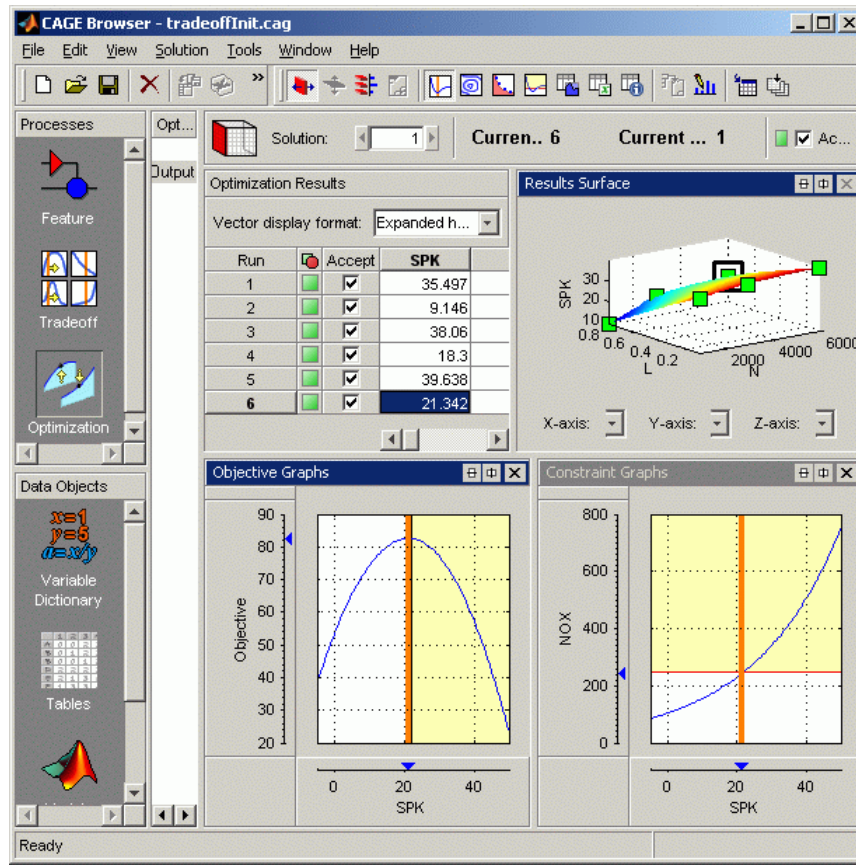

For more information, see “Analyzing Point-by-Point Optimization Output” in the CAGE User’s Guide documentation.

Using Optimization Results to Fill Tables

As an example, to use these optimization results to fill a table, first create a new table as follows:

1. Build a `SPK` table in `N` and `L`. `Select` **File > New > 2-D Table**.
2. Leave `10` in the **Rows** and **Columns** edit boxes and `0` in the **Initial Value** edit box.
3. Use the drop-down menus to select `L` and `N` for the Y and X inputs.
4. Rename the table to `SPK_Table`.
5. Click **OK**. Your CAGE browser switches to the **Tables** view. CAGE has automatically initialized the normalizers to space breakpoints evenly over the ranges of `N` and `L`.

There are two methods for filling tables with optimization results.

1. Click the **Optimization** button in the **Processes** pane to return to the Optimization view
2. Click the plus to expand the Optimization node, and select the `Optimization_Output` node.
3. Select **Solution > Fill Tables** (or the toolbar button **Fill tables using optimal settings**).

   The Table Filling wizard appears.
4. Select the `SPK_Table` table and click the button to add it to the list of tables to be filled. Click **Next**.
5. Select the `SPK_Table` table, and double-click `SPK` in the list of optimization results to select it to fill the table, as shown.

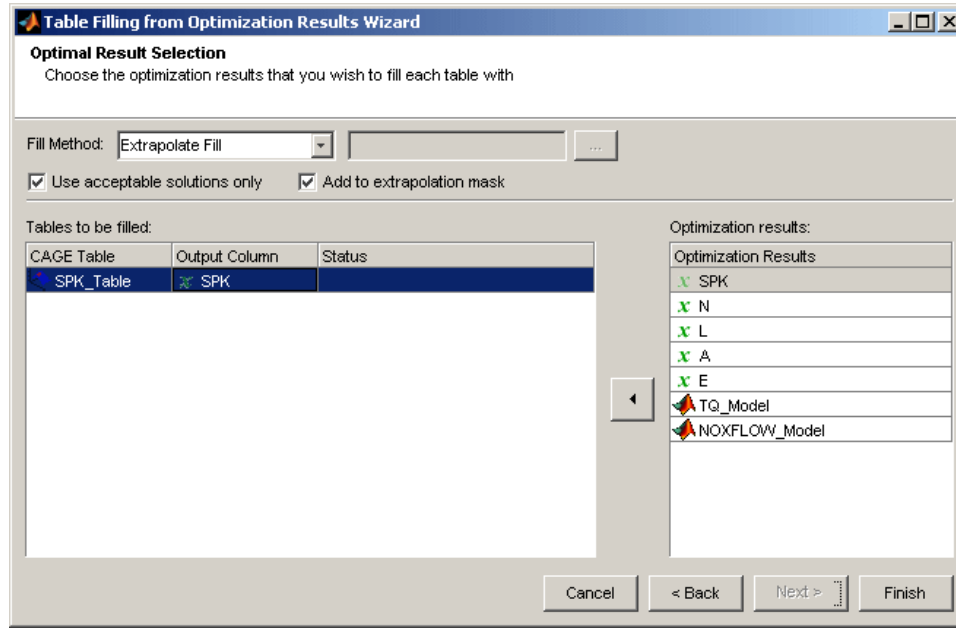

6 Click **Finish**.

You will see a dialog reporting successful table filling. Switch to the **Tables** view to examine the new spark table.

The other method of filling tables with optimization output uses Data Sets.

1 From the `Optimization_Output` optimization output node, click Export to Data Set ( ) in the toolbar (or select **Solution > Export to Data Set**). Click **OK** in the Export to Data Set dialog box to accept the defaults.

2 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 `New_Dataset`.

You can now use this data set (or any optimization results) to fill tables, as you can with any data set.

3 Select the data set and click (Fill Table From Data Set) in the toolbar.

4. Clear the check box at the bottom to **Show table history after fill**.

5. Choose to fill the spark table with the `SPK` optimization output by selecting them in the two lists, then click the button **Fill Table** at the bottom right.

6. Right-click the display and select **Surface** to see the filled table surface and the optimization output spark values.

See also Chapter 11, "Tutorial: Filling Tables from Data" for more details on using data sets to fill tables.

In the next section you will use a custom fill routine to fill the table.

Using a Custom Fill Routine to Fill Tables

It can be useful to create your own custom fill function to fill tables from the results of an optimization. Some example situations are:

- You have your own smoothing strategy for certain regions of your look-up tables
- Implementation of an alternative method to the two fill methods supplied
- You want to produce some customized output

You can use a custom fill routine to fill the `SPK_Table` table from the optimization results.

1. Create a custom fill function. For this example, you can use the supplied example, `griddataTableFill.m`, which can be found in the `mbctraining` directory. Copy `griddataTableFill.m` to a directory away from the MATLAB root directory, and make sure this directory is on the MATLAB path (or change the current directory to the location where you copied the file).

2. At the optimization output node, select **Solution > Fill Tables**.

3. The wizard retains a memory so the `SPK_Table` table is already selected to be filled. Click **Next**.

4. Similarly, `SPK` is already selected from the list of optimization results to fill the table.

5 Select `Custom` from the **Fill Method** drop down menu. Use the file selector, or enter the name of the fill function you wish to use to fill your tables. In this case, select or enter `griddataTableFill`. Note that this function must be on the MATLAB path.

6 Click **Finish** to fill the `SPK_Table` table.

In the next section you will add a multiobjective optimization to this project.

Multiobjective Optimization

In this section...

"Setting Up and Running the Multiobjective Optimization" on page 12-19

"Optimization Output View" on page 12-23

"Selecting Best Solutions" on page 12-29

Setting Up and Running the Multiobjective Optimization

In this optimization you will construct a search for values of spark that maximize values of torque while minimizing values of NOX at a series of (`L`, `N`, `A`, `E`) points.

1. Select **File > New > Optimization**.

   This opens the Optimization Wizard.

2. Click to select `NBI` in the list. This is the optimization algorithm you will use for this example. Note that this algorithm specifies two or more objectives in the **Objectives** column. Click **Next**.

3. On the next screen, set the number of constraints to `1`. Leave the number of free variables at `1` (this will be spark) and objectives at `2` (these will be the `TQ` and `NOXFLOW` models). Click **Next**.

4. On the next screen, select `SPK` as your free variable and click the button to match it up to `FreeVariable1`. Click **Next**.

5. Here you must select models for your objectives.

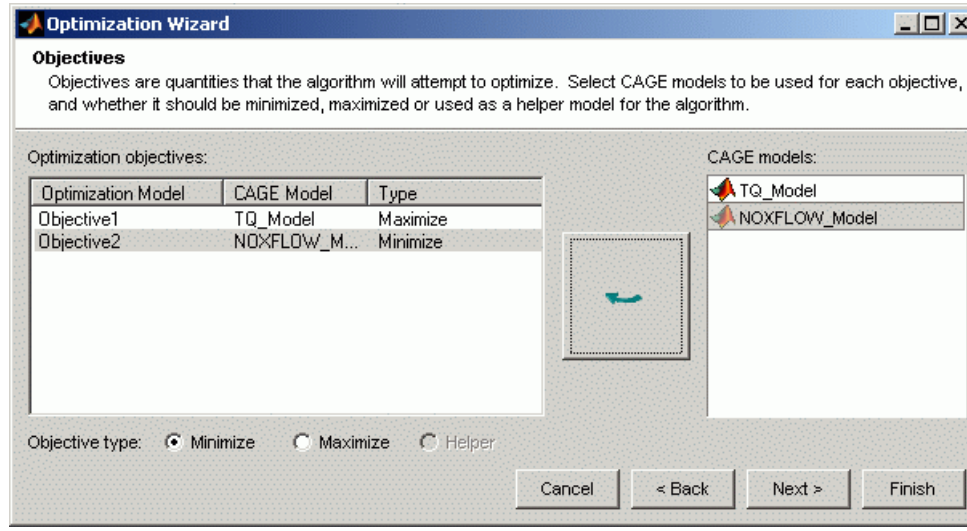

a Select `TQ_Model` and click the button to match it up to `Objective1`, then click the **Maximize** radio button. Check `Maximize` appears next the `TQ_Model` as shown.

b Select `Objective2`, then `NOXFLOW_Model` and click the button to match them up. Leave the `NOXFLOW_Model` objective set to **Minimize**.

Note that this stage was skipped for the first simple example (single-objective optimization) by clicking **Finish** on the previous screen.

- You can set up objectives, constraints, and operating point sets in the Optimization Wizard. You can change any of these settings later.
- You can also click **Finish** on any of these screens of the wizard and set up any or all of these later from the main **Optimization** view in the CAGE Browser, as in the first tutorial example.

Click **Next**.

6 On the next screen you can set up a constraint.

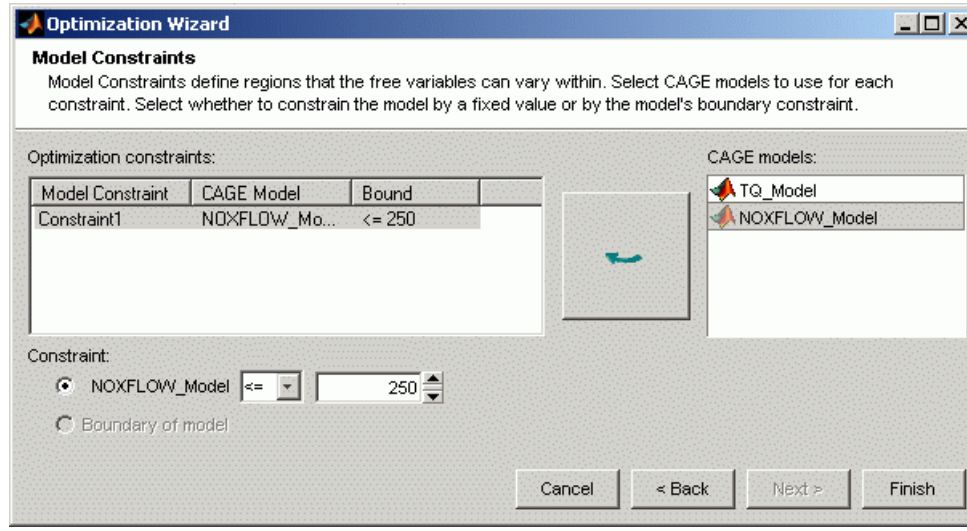

Select the `NOXFLOW_Model` and click the button to match it to `Constraint1`. Enter `250` in the edit box and press **Enter**. Leave the operator at `<=` to constrain the `NOXFLOW_Model` to a maximum of 250. Click **Finish**.

A new node, `Optimization_1`, appears in the Optimization tree. Your CAGE browser should look like the following example. Your optimization has objectives and a constraint set up and is ready to run. However unless you edit the fixed variable values it will run at a single point, the set point of the variables.

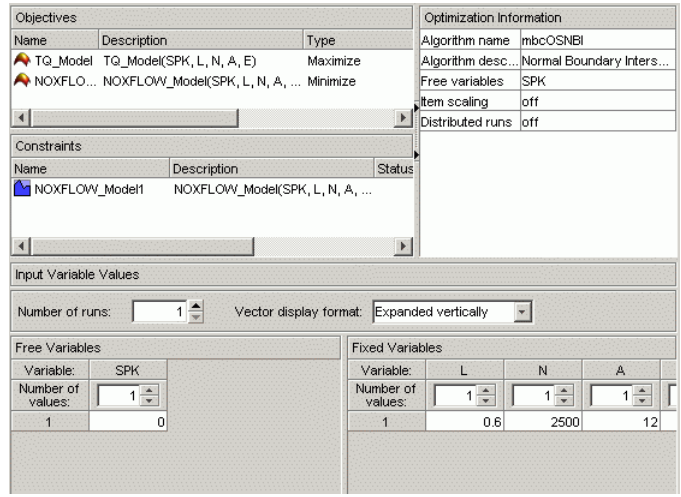

7 In the **Input Variable Values** pane, increase the **Number of runs** to 6. Notice 6 rows appear in both fixed and free variable values panes, all containing the default set point values of each variable.

8 Select **Optimization > Import From Output**. The Import From Output dialog box appears.

a Select the previous single objective optimization node, `Optimization_Output`, in the top list to import values from this optimization.

b Clear the check boxes for `SPK`, `A`, and `E`, to leave only `N` and `L` selected for import, and click **OK**.

9 View these values in the N and L columns in the **Fixed Variables** pane.

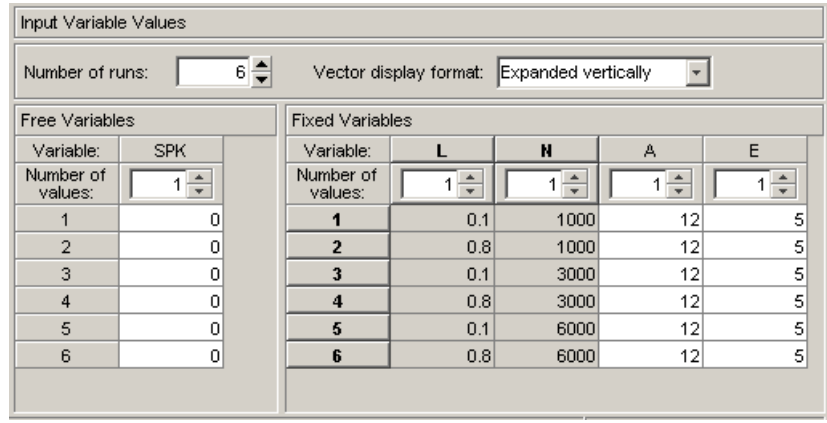

10 Click Run Optimization ( ) in the toolbar.

The optimization runs, showing progress messages as each point is evaluated until the optimization is complete. A new node, `Optimization_1_Output`, appears under `Optimization_1` in the Optimization tree.

Optimization Output View

The view switches to the `Optimization_1_Output` node in the Optimization tree where you can examine the optimization output.

The toolbar buttons determine which view is displayed. The default is the Solution Slice ( ). The Solution Slice shows one solution at all operating points. That is, you can see a table (and surface plot) of all operating points at once, and you can scroll through the solutions using the **Solution** buttons at the top. At the start all 6 operating points show solution 1. Change solution to 2, and you see the second solution for all 6 operating points, and so on. As this is a multiobjective optimization, there are several solutions for each operating point.

Change solution number to `10`.

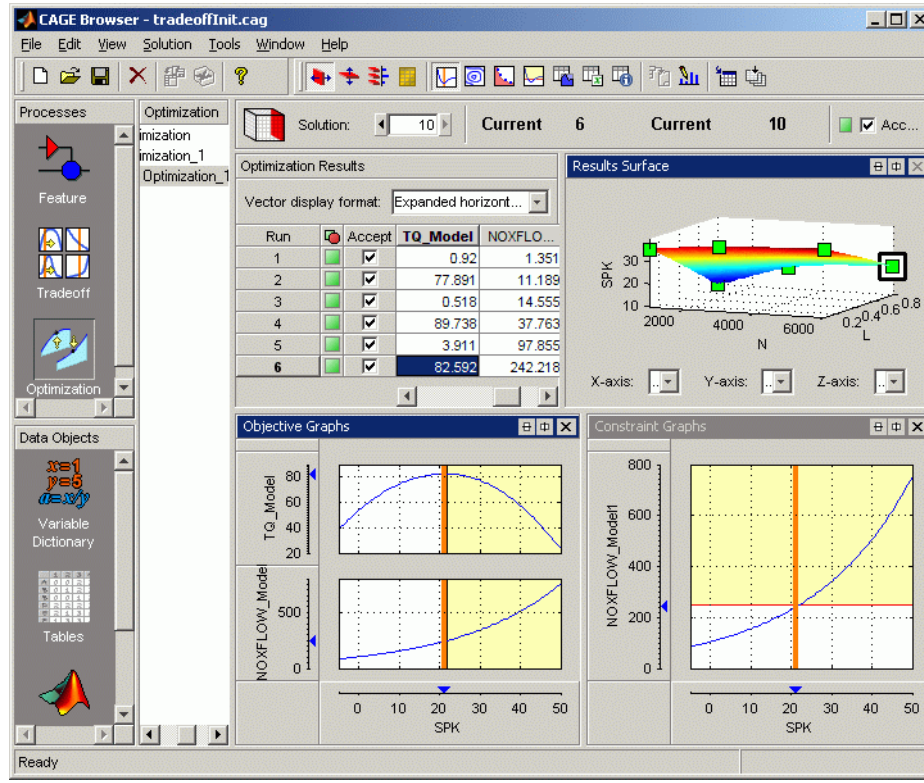

The graphs show the objective functions at the currently selected operating point (highlighted in the table and Results Surface view), with the solution value shown in red.

Note that before you run an optimization you can specify how many solutions you want the optimization to find, using the Set Up and Run Optimization toolbar button.

1 For an example point, click in the table to select operating point 6, and enter 10 in the **Solution** edit box. Observe the constraint you applied in the objective function graphs, as shown in the example. Areas in yellow are excluded by constraints. Similarly, if you use a boundary constraint model exported from the Model Browser as a constraint, areas outside the boundary appear in optimization graphs as yellow areas. Note that for some

problems the optimization might fail to find a value within the constraints (depending on the constraints and starting values) in which case you might need to run the optimization again to find valid solutions. Choosing more suitable starting values and changing your settings to make constraints less stringent can help in these cases. See “Analyzing Point-by-Point Optimization Output” in the CAGE User’s Guide documentation.

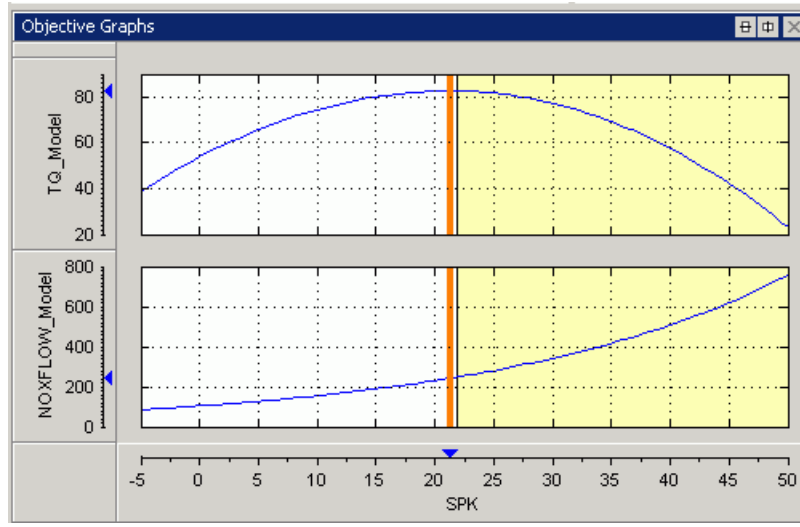

2 Right-click the Objective Graphs view and select **Split View > Pareto Graphs**.

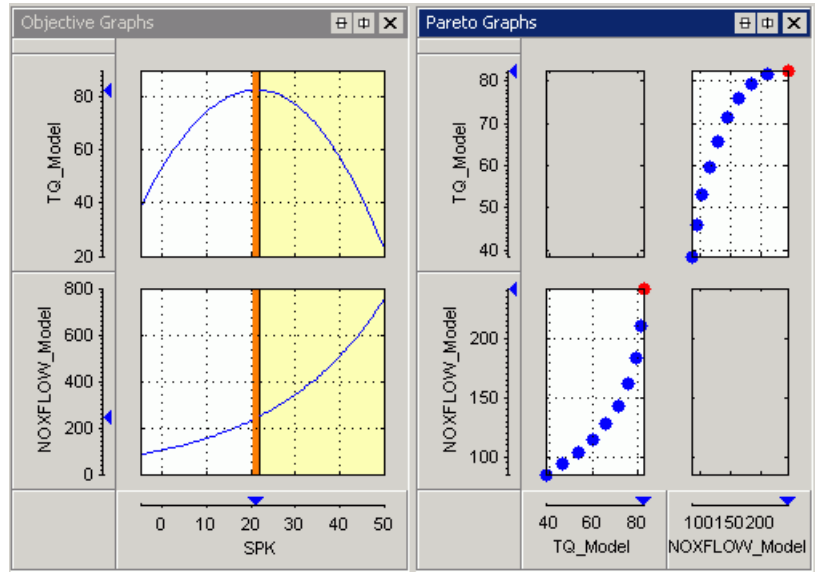

The view splits to show both objective and pareto graphs. You can right-click and select **Graph Size** to adjust how many plots can be displayed. In the Pareto Graphs view you can see all solutions found by the optimization at the selected operating point (the selected solution is highlighted in red). Try clicking different points in the pareto graph to see the different solutions in the objective graphs.

3 Click Pareto Slice ( ) in the toolbar. This changes the table to display all solutions at a single operating point. You can scroll through the operating points using the **Run** buttons at the top. The pareto graphs always show the currently selected solution in red. Click in the table or the graph to select different solutions.

Recall that the first example, a single-objective optimization, produced a single solution at each point, so you could not view the Pareto Slice. The Pareto Slice is useful to show you the set of optimal tradeoff solutions when using multiobjective optimizations, as in this case. You can use these plots to help you select the best solution for each operating point. As you can see, this example trades off NOX emissions for torque, so it is a judgment call to choose the best depending on your priorities. You will select best solutions in a later section, “Selecting Best Solutions” on page 12-29.

4 Click Weighted Pareto Slice ( ) in the toolbar.

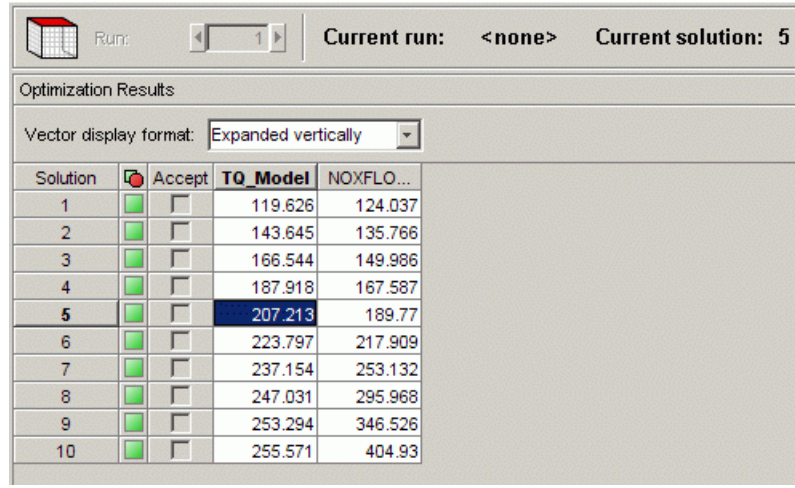

Solution	Accept	TQ_Model	NOXFLO...
1		119.626	124.037
2		143.645	135.766
3		166.544	149.986
4		187.918	167.587
5		207.213	189.77
6		223.797	217.909
7		237.154	253.132
8		247.031	295.968
9		253.294	346.526
10		255.571	404.93

This table view displays a weighted sum objective output across all operating points for each solution.

The value in the `NOXFLOW_Model` column in the first row shows the weighted sum of the solution 1 values of NOX across all 6 operating points. The second row shows the weighted sum of solution 2 NOX values across all 6 operating points, and so on. This can be useful, for example, for evaluating total emissions across a drive cycle. The default weights are unity (1) for each operating point.

5 You can alter these weights by clicking Edit Pareto Weights ( ) in the toolbar. The Pareto Weights Editor appears.

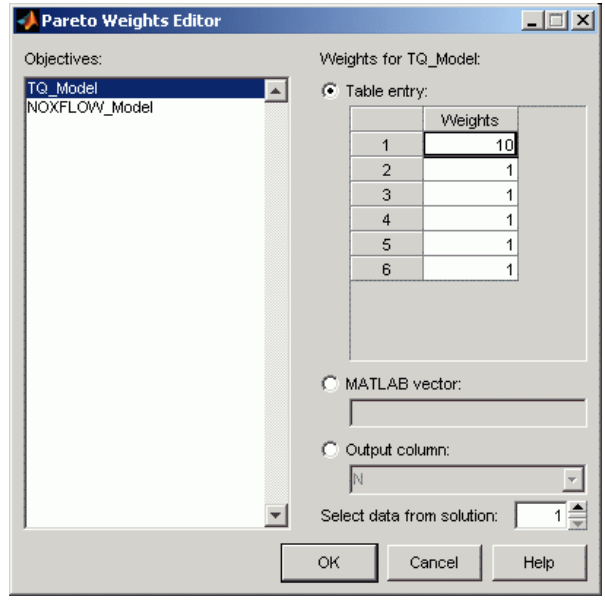

Here you can select models, and select weights for any operating point, by clicking and editing, as shown in the example above. The same weights are applied to each solution to calculate the weighted sums. Click **OK** to apply new weights, and the weighted sums are recalculated.

You can also specify weights with a MATLAB vector or any column in your optimization output by selecting the other radio 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 Slice 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.

Selecting Best Solutions

In a multiobjective optimization, there is more than one possible optimal solution at each operating point. You can use the Selected Solution Slice to collect and export those solutions you have decided are optimal at each operating point.

Once you have enabled the Selected Solution Slice, you can use the plots in the Pareto Slice and Solution Slice to help you select best solutions for each operating point. These solutions are saved in the Selected Solution Slice. You can then export your chosen optimization output for each point from the Selected Solution Slice, or use your chosen optimization output to fill tables.

1. In order to choose a single solution at each operating point, you need to enable the Selected Solution Slice. Select **Solution > Selected Solution > Initialize**.

   A dialog called Initialize Selected Solution appears. Click **OK**.

   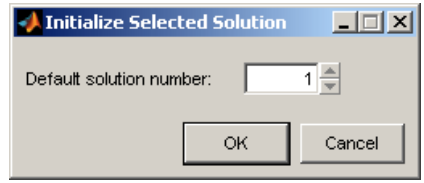

   The default initializes the first solution for each operating point as the selected solution.

2. Click the Selected Solution Slice button ( ) which is now enabled in the toolbar. Observe that the **Solution** number at the top is not editable, and is initially solution 1 for each operating point you click in the table. You must select the solutions you want using the Pareto Slice and Solution Slice to decide which solution is best for each point.

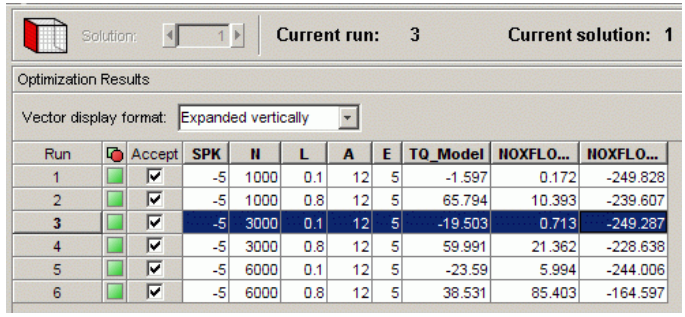

Run	Accept	SPK	N	L	A	E	TQ_Model	NOXFLO...	NOXFLO...
1	☑	-5	1000	0.1	12	5	-1.597	0.172	-249.828
2	☑	-5	1000	0.8	12	5	65.794	10.393	-239.607
3	☑	-5	3000	0.1	12	5	-19.503	0.713	-249.287
4	☑	-5	3000	0.8	12	5	59.991	21.362	-228.638
5	☑	-5	6000	0.1	12	5	-23.59	5.994	-244.006
6	☑	-5	6000	0.8	12	5	38.531	85.403	-164.597

3 Return to the Pareto Slice and select a solution for run 6. Enter `6` in the **Run** edit box, and click in the table or in the pareto graphs to select a solution. An example is shown with solution 7 highlighted. To select this solution as best, do one of the following:

- Click Select Solution ( ) in the toolbar.
- Select the menu item **Solution > Selected Solution > Select Current Solution**.

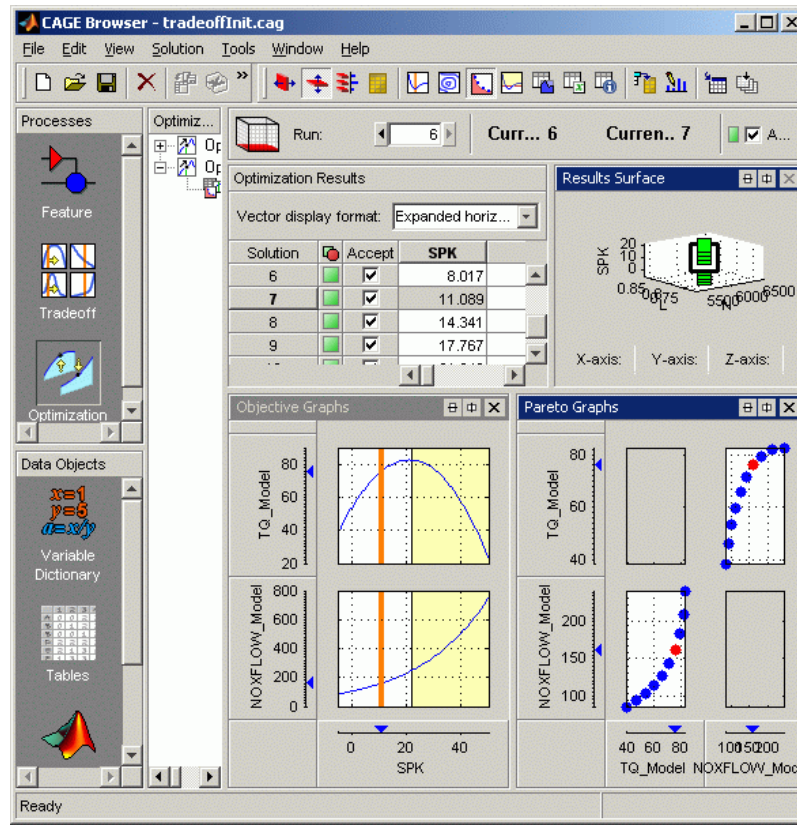

4 If you return to the Selected Solution Slice you can now see that solution 7 is now present for operating point 6, while all the other operating points remain at the initial solution 1. This view collects all your selected solutions together in one place. For example, you might want to select solution 7 for the first operating point, and solution 6 best for the second, and so on.

5 In order to use your optimization output to fill tables, you should repeat this process to select a suitable solution for all operating points. Then use the Table Filling From Optimization Results Wizard (**Solution > Fill Tables**) as before — the table is filled with your selected solution for each run.

Alternatively you could fill from a data set containing the selected solutions. To do this, click Export to Data Set ( ) and click **OK** in the

dialog box. Go to the **Data Sets** view (click **Data Sets** in the **Data Objects** pane) to see that the table of optimization results is contained in a new data set. You could use these optimization results to fill tables. Both these table-filling methods are described in "Using Optimization Results to Fill Tables" on page 12-14.

Note that the table in the current view is exported to the data set. If you want to export your selected best solutions for each operating point, make sure you display the Selected Solution Slice before exporting the data. If you export from the Pareto Slice, the new data set contains all solutions at the single currently selected operating point set. If you export from the Solution Slice the new data set will contain the current solution at all operating points.

Recall that the previous example was a single-objective optimization and therefore only had one solution per operating point. In that case the optimization results could be exported directly from the Solution Slice, as there was no choice of solutions to be selected. See "Single-Objective Optimization" on page 12-5.

In the next tutorial section you will duplicate this NBI optimization example and alter it to create a sum optimization.

Sum Optimization

In this exercise you will use a copy of the NBI optimization from the last example to create a sum optimization.

Up to this point, you have found the optimal values of each objective function at each point of an operating point set individually. A sum optimization finds the optimal value of a weighted sum of each objective function over all free variables simultaneously. The weighted sum is taken over each operating point in the run, and the weights can be edited.

A sum optimization problem without constraints, for M operating points, is the same as M individual optimizations; the minimum of the sum is the same as the sum of the minima.

To illustrate this, consider the following example:

Say the objective function is $f(a,b)$. Consider a to be the free variable and b to be the fixed variable. To set up a sum optimization, at different values of b, we want to find $[a_1, a_2, a_3 .. a_M]$ which minimizes:

$$w_1*f(a_1,b_1) + w_2*f(a_2,b_2) + w_3*f(a_3,b_3) + ... + w_M*f(a_M,b_M), \quad [1]$$

where w_1, w_2 etc. are the weights.

There are two ways of viewing this problem. It can be viewed as a big optimization problem in an M-dimensional space for the vector $[a_1, a_2, a_3, ..a_M]$ as shown in [1]. Alternatively, as each element of the sum depends on its own subset of the free variables, the problem can be written as M separate optimization problems, as in [2]:

$$min\ w_i*f(a_i,b_i)\ for\ i=1:M, \quad [2]$$

Once the M individual problems are solved, the weighted sum can be constructed to get the answer.

When there are sum constraints present, it is not true that a M-point sum optimization problem is equivalent to solving M individual optimizations. In this case, all points must be evaluated together to find the optimal solution that meets the sum constraints across all of the points.

In a sum optimization, the objectives are typically sum objectives. You can use a mixture of point and model sum constraints in a sum optimization. The following instructions describe these settings.

1. Right-click the `Optimization_1` node in the Optimization tree and select **Duplicate Optimization_1**.

   A copy of the `Optimization_1` node called `Optimization_2` appears in the tree. You do not need an existing NBI optimization to create a sum optimization. This example is used for convenience and also to illustrate copying optimizations. This feature can be useful when you are trying different settings to improve your optimizations while keeping previous attempts for comparison.

   You use this copy to create a sum optimization. This means that instead of performing the optimization at each point individually, the optimization takes the sum of solutions at all points into consideration. You can apply different weights to operating points, allowing more flexibility for some parts of the optimization.

2. Select the node `Optimization_2` and select **Edit > Rename** (or press **F2**). Edit the name to read `SUM_NBI`.

   You will edit all the objectives in your existing optimization to be sum objectives.

3. Double-click `TQ_Model` (or right-click and select **Edit Objective**). The Edit Objective dialog appears.

   To set up your new objective,

   a. Select `Sum Objective` from the **Objective Type** drop-down menu.

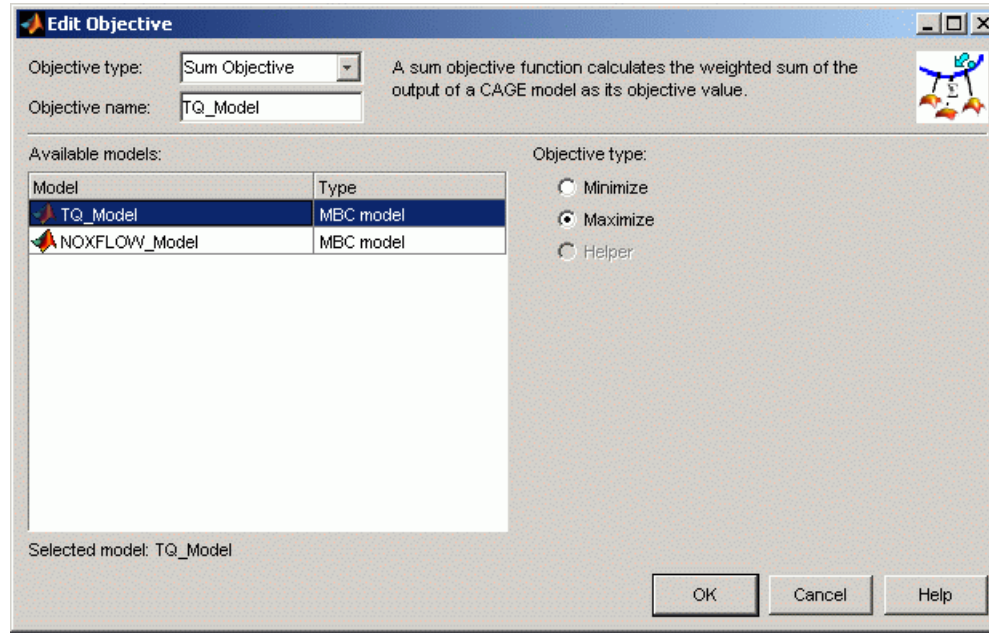

b Select `TQ_Model` and make sure **Maximize** is selected. Torque has a strong correlation with fuel consumption so this sort of problem could be useful for a fuel consumption study.

c Click **OK** to finish editing the objective.

4 Repeat to edit `NOXFLOW_Model`. Set up a sum objective to minimize NOXFLOW.

5 You can also edit your constraint to be a sum constraint. You can use a mixture of point and sum constraints. Double-click `NOXFLOW_Model1` and the Constraint Editor appears.

To set up your sum constraint,

a Select `Sum Constraint` from the **Constraint Type** drop-down menu, then select `NOXFLOW_Model` under **Input Model**.

b Edit the **Constraint name** to `NOXFLOW_Constraint` to help you later distinguish it from the `NOXFLOW_Model` in the Optimization view fixed variables pane.

c Make sure the inequality is <=, and enter 450 in the constraint bound edit box as shown.

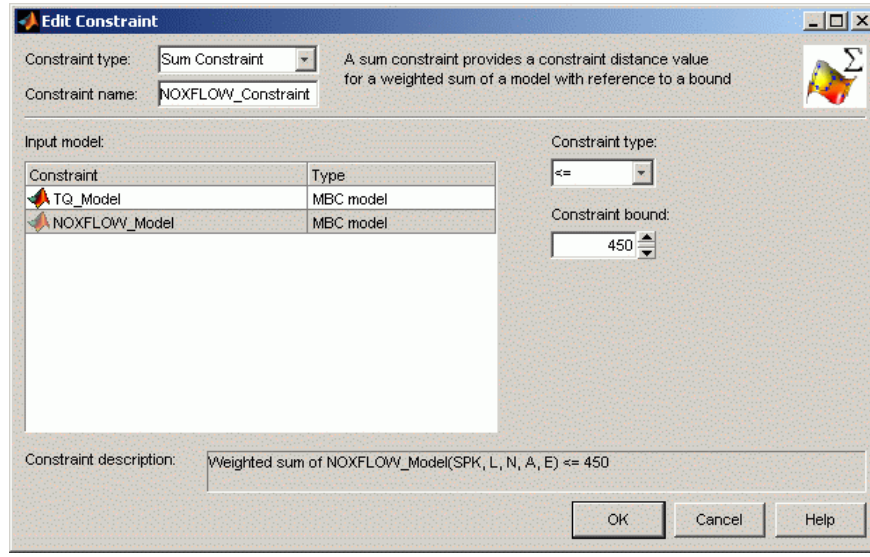

d Press **Enter** to dismiss the dialog box and return to the Optimization view.

6 You want to perform a sum optimization across a series of operating points. Select **Optimization > Convert to Single Run**. Look at the change in the **Variable Values** panes. Instead of 6 separate runs you now have one run containing 6 operating points. **Number of runs** is now 1, and **Number of Values** is 6 for all variables (you use these controls to set the number of operating points within each run). Each solution will be a sum across all the points in the run.

7 Look in the **Fixed Variables** pane for the weights columns, last columns on the right. Enter 5 in the first row (run 1, point 1) for **TQ_Model_weights**, **NOXFLOW_Model_weights**, and **NOXFLOW_Constraint_weights**, to weight the first operating point by five, as shown.

Fixed Variables

Variable:		TQ_Model_weights	L	N	A	E	NOXFLOW_Model_weights	NOXFLOW_Constraint_wei...
Number of values:		6	6	6	6	6	6	6
1	(1)	5	0.1	1000	12	5	5	5
	(2)	1	0.8	1000	12	5	1	1
	(3)	1	0.1	3000	12	5	1	1
	(4)	1	0.8	3000	12	5	1	1
	(5)	1	0.1	6000	12	5	1	1
	(6)	1	0.8	6000	12	5	1	1

8 You have modified your objectives and constraint for a sum optimization, which is ready to run. Click Run Optimization ( ) in the toolbar.

9 There is a wait notice as the optimization runs. There are no progress messages as points are evaluated because sum optimizations do not evaluate points individually.

When the optimization is complete, examine the results at the output node. Select the Pareto Slice table. Look at the objective and constraint results for the ten solutions (number of solutions is specified in the Optimization Parameters dialog). In the Pareto Slice table, negative constraint values are within the constraint. Look at the constraint summary view, where the left and right information corresponds to the constraint inequality; the left value is the distance from the constraint.

For more information see “Interpreting Sum Optimization Output” in the CAGE User’s Guide documentation.

Automated Tradeoff

Once you have set up an optimization you can fill tables in a tradeoff using automated tradeoff. You can select cells and fill them from the results of an optimization. The cells you select in the tradeoff table define the operating point set for the optimization.

Set up a tradeoff as follows (also described in "Setting Up a Tradeoff Calibration" on page 9-3).

1. Select **File > New > Tradeoff**.

   This takes you to the **Tradeoff** view. You need to add tables to the tradeoff.

2. Click (Add New Table). This opens the Table Setup dialog.

3. Enter `Spark` as the table **Name**.

4. Select `L` as the **Y name** and `N` as the **X name**.

5. Click **Select** to open the Select Filling Item dialog.

   a. Select the radio button to **Display variables**.

   b. Click to select `SPK`.

   c. Click **OK** to return to the Table Setup dialog.

6. Leave `10` as the size of the rows and columns (the speed and load axes), and 0 as the initial value, and click **OK**.

   A new `Spark` table appears in the **Tradeoff** tree. CAGE has automatically spaced the normalizers evenly over the ranges of `N` and `L`.

7. Click to expand the `New_Tradeoff` tree and select the `Spark` table node to view the new table.

You need to select the cells where you want to apply automated tradeoff. Create a region within the table:

1 Highlight a rectangle of cells in the SPK table by clicking and dragging. Note that a large region can take a very long time to evaluate. Try four cells to start with.

2 Click (or right-click and select **Extrapolation Regions > Add Selection**) to define the region. The cells become light blue.

To use automated tradeoff on the cells in the defined region,

1 Select **Inputs > Automated Tradeoff** or click the toolbar button .

The Automated Tradeoff dialog appears, showing a list of available optimizations in your session that are set up and ready to run.

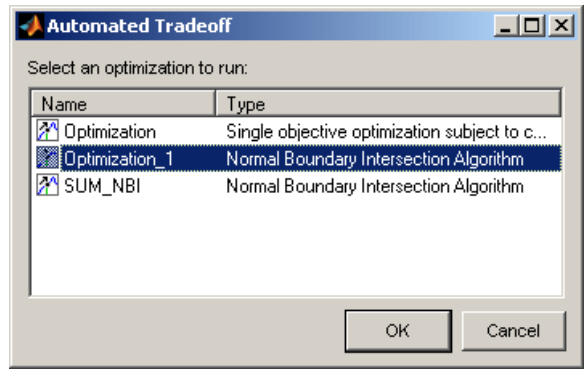

2 Select your `Optimization_1` multiobjective optimization to apply to the tradeoff and click **OK**.

3 Click **OK** in the following dialog to optimize only the table cells that are in the region, rather than all cells.

The automated tradeoff optimization runs. The results appear in the selected cells in the table, as shown in the example. You could optimize several regions and then use these results to extrapolate across the whole table.

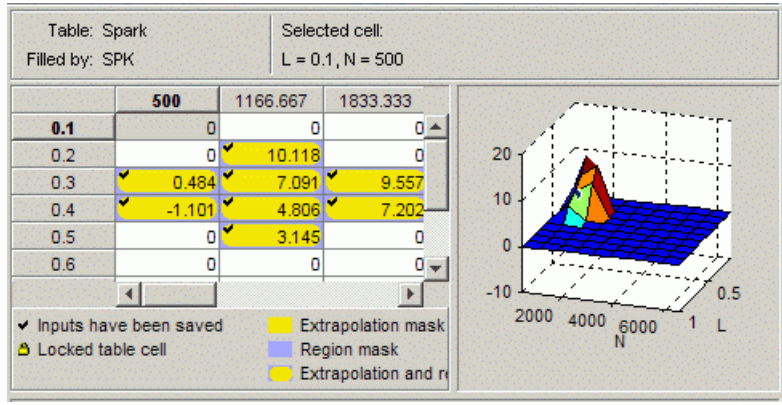
Table: Spark
Filled by: SPK
Selected cell:
L = 0.1, N = 500
500
1166.667
1833.333
0.1
0.2
0.3
0.4
0.5
0.6
0
10.118
0.484
7.091
9.557
-1.101
4.806
7.202
3.145
Inputs have been saved
Locked table cell
Extrapolation mask
Region mask
20
10
0
-10
2000
4000
6000
N
1
0.5
L

Worked Example Optimization

In this section...

"Example Overview" on page 12-41

"Using the Worked Example Optimization" on page 12-42

Example 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 M-file. This worked example modifies the template provided to show you how to use your own algorithms within CAGE. You can find detailed information on all the available CAGE optimization interface functions in "User-Defined Optimization" in the CAGE documentation.

- To view the worked example M-file, at the command line, type

```
edit mbcOSworkedexample
```

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.

The next section, "Using the Worked Example Optimization" on page 12-42, demonstrates how to use the example within CAGE.

The section "Creating an Optimization from Your Own Algorithm" on page 12-48 is a detailed tutorial example explaining how to incorporate an example user-defined optimization algorithm into a CAGE optimization function.

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` directory.
   - The `tq` model was fitted to the Holliday engine data and exported from the Model Browser quick start tutorial (also used in the CAGE feature calibration tutorial). It can be found in `tutorial.exm` in the `mbctraining` directory. To view this model in your current session, click the **Models** button in the **Data Objects** pane. There is also another model in the session that you will use later.
   - 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 > 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**.

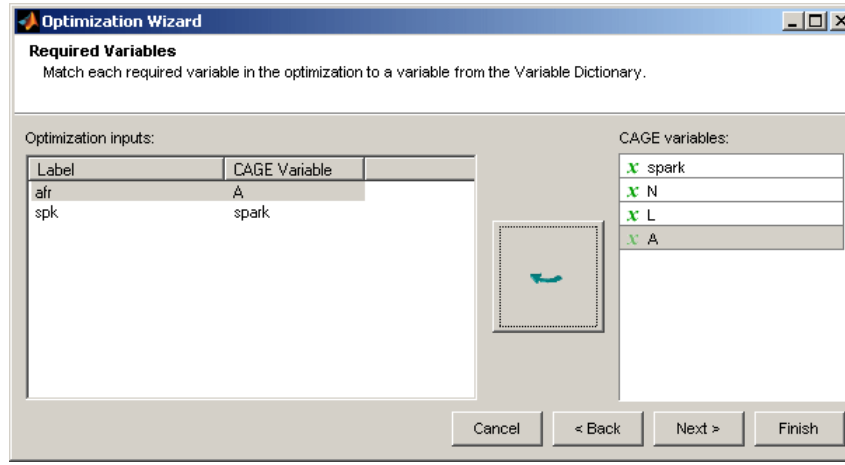

5 The next screen of the wizard automatically shows the `Torque` model selected and **Maximize** chosen; these are specified in the function. Select `tq` in the CAGE model list and click the button to match it with the `Torque` optimization model, then click **Finish**.

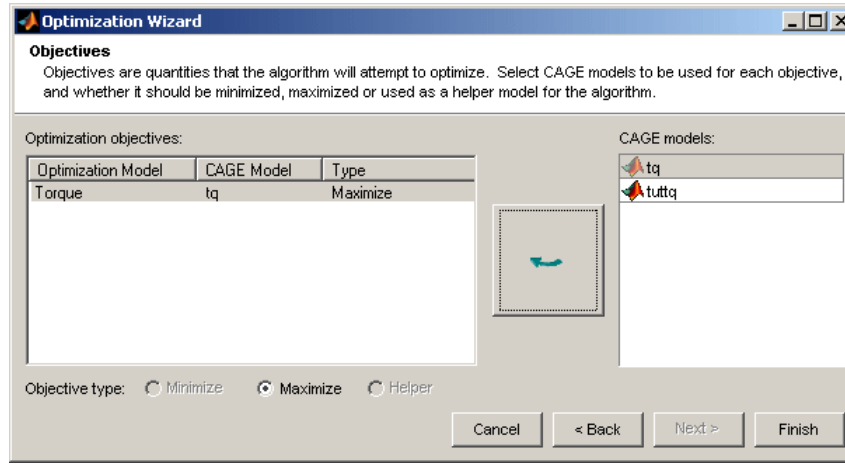

CAGE switches to the **Optimization** view and the new `Optimization` node appears in the tree.

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**.

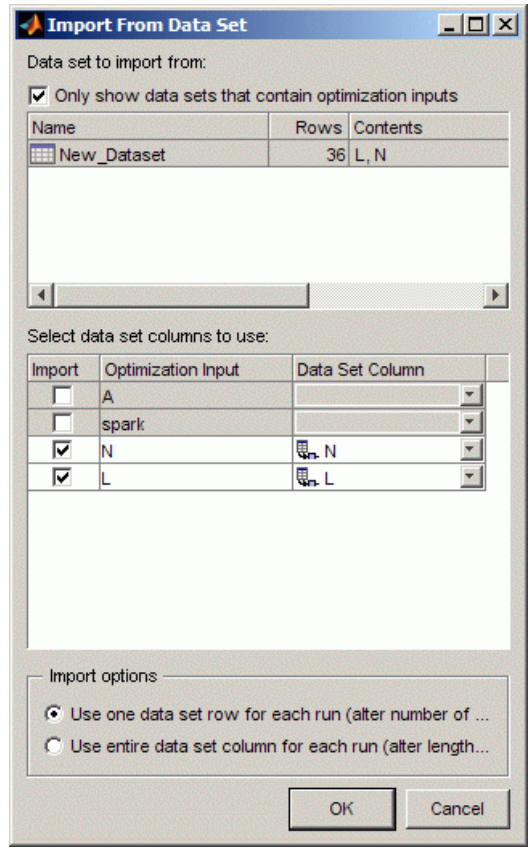

The project file contains a data set with N and L values, and these are automatically selected. Click **OK** to import.

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.

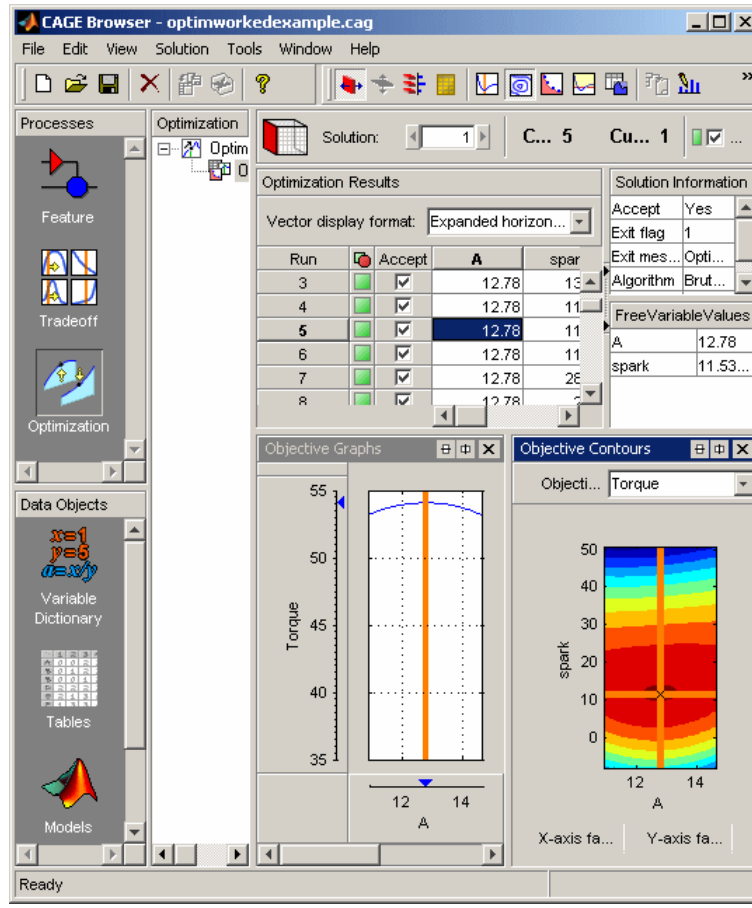

For a detailed walk-through of incorporating an example user-defined optimization algorithm into a CAGE optimization function, see the next tutorial section, “Creating an Optimization from Your Own Algorithm” on page 12-48.

Creating an Optimization from Your Own Algorithm

In this section...

"Process Overview" on page 12-48
"Step 1: Verify the Algorithm" on page 12-50
"Step 2: Create a CAGE Optimization Function" on page 12-52
"Step 3: Define the Optimization Options" on page 12-55
"Step 4: Add the Algorithm to the Optimization Function" on page 12-58
"Step 5: Register Your Optimization Function with CAGE" on page 12-62
"Step 6: Verify Your New Optimization" on page 12-63

Process Overview

The CAGE optimization feature allows you to use your own optimization algorithms as alternatives to the library routines `foptcon`, `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/Holliday.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 directory.

1. Create a new folder (for example, `C:\Optimization_Work`). We recommend that you place this directory outside your MATLAB folders to avoid interfering with toolbox files.

2. Copy the following six files from the mbctraining directory into your new working folder:

   ```
   currtutoptim.m
   mbcOStemplate.m
   mbcOStutoptimfunc_s1.m
   mbcOStutoptimfunc.m
   optimtut.mat
   optimtuteg.mat
   ```

3. Make sure your new working directory is on the MATLAB path; either change **Current Directory** in MATLAB to the new working folder, or add the folder to the path as follows:

   a. Select **File > Set Path**.

   b. Click **Add Folder** and browse to your working directory.

   c. Click **OK**.

   d. Click **Save**.

   e. Click **Close**.

Step 1: Verify the Algorithm

`currtutoptim.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 currtutoptim.m` at the command line, or if the **Current Directory** in MATLAB is your new working folder, select **Desktop > Current Directory**, then double-click `currtutoptim.m` . You should see the following code in the MATLAB editor.

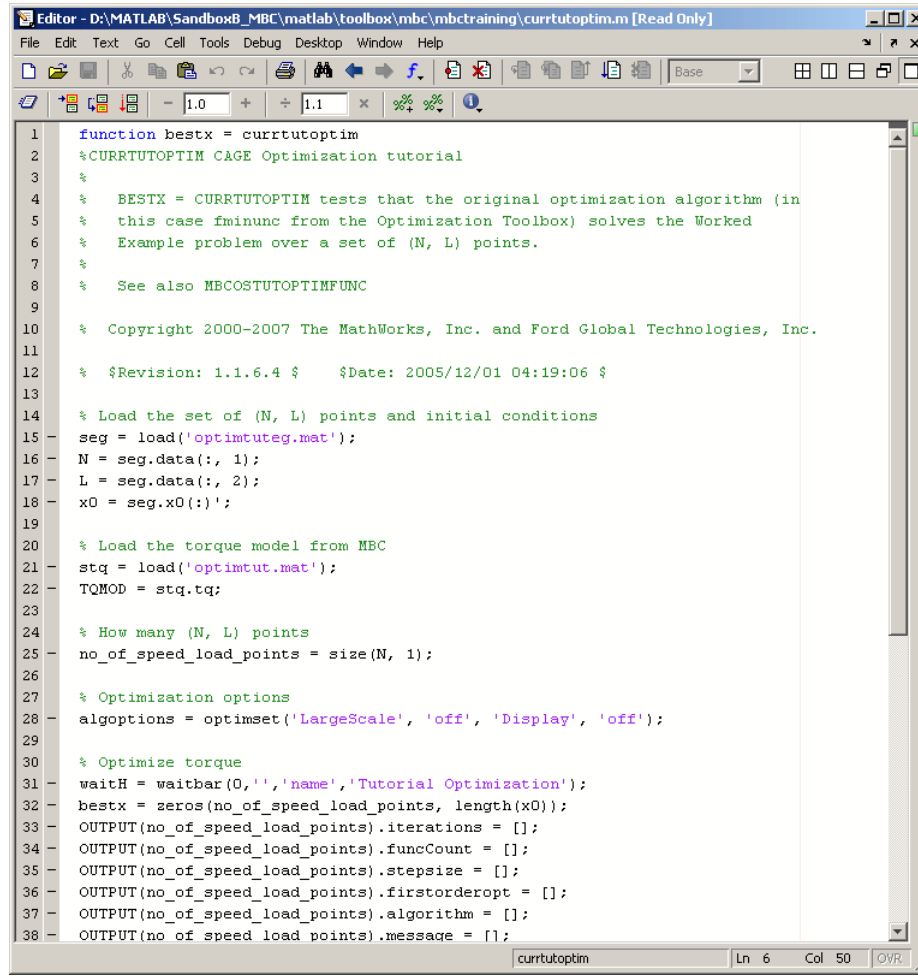

2 To verify that `fminunc` solves the worked example problem, type the following command at the MATLAB prompt:

```
bestX = currtutoptim
```

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 `mbcOStemplate.m`.

You should see the following M-file in the MATLAB editor.

```
function out = mbcOStemplate(action, in)
%MBCOSTEMPLATE CAGE Optimization template function
%   OUT = MBCOSTEMPLATE(ACTION, IN) is a template function for use with
%   CAGE Optimization. This function can be used to create user-defined
%   optimization functions that can subsequently be used in CAGE.

%  Copyright 2000-2007 The MathWorks, Inc. and Ford Global Technologies, Inc.

%  $Revision: 1.1 $    $Date: 2006/12/13 12:36:54 $

% Deal with the action inputs
if strcmp(action, 'options')

    options = in;

    %
    % Define optimization attributes here
    %

    out= options;

elseif strcmp(action, 'evaluate')

    optimstore = in;

    %
    % Put optimization algorithm here
    %

    out = optimstore;

else
    error('Incorrect action type specified');
end

end
```

Section 1

Section 2

`mbcOStemplate.m` is an empty CAGE optimization function. The two (currently empty) sections of the function are labeled above. Note that this M-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 `mbcOStutoptimfunc_s1.m`. In this M-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 `mbcOStutoptimfunc_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" in the CAGE documentation. The above code has used the `cgoptimoptions` 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`

Attribute	Value	Code Section Where Set
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 labelled **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 `cgoptimoptions` 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 `mbcOStutoptimfunc.m`.

This M-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 subfunction `tutoptimizer` can be found at the bottom of the `mbcOStutoptimfunc.m` file. 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, `currtutoptim`, 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 (`x0`) for the free variables.

Original code:

`x0` passed in from the MATLAB workspace.

Modified code:

```
x0 = getInitFreeVal(optimstore);
```

In the original algorithm, `x0` is passed into the algorithm from the MATLAB workspace. In CAGE, we invoke the `getInitFreeVal` function on the `optimstore` object to retrieve `x0`.

Algorithm Section 2

Perform the optimization (in Section 2 of the template).

Original code (from `currtutoptim`):

```
    [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 `currtutoptim`.

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 M-file.

3. Locate `mbcOStutoptimfunc.m` file (in the working directory you created) and click **Open**.

   This registers the optimization function with CAGE.

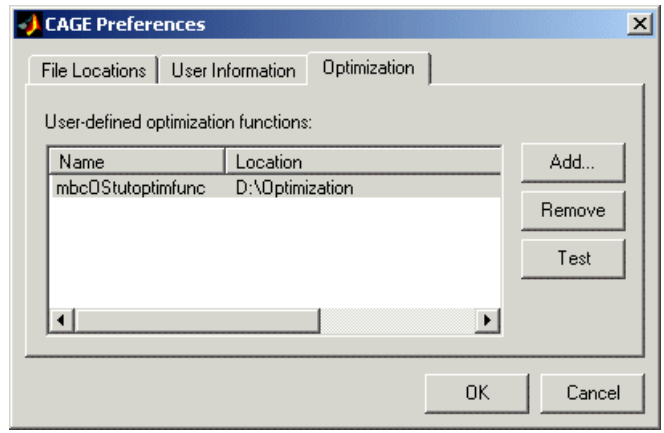

4 You can now test the function by clicking **Test**. This is a good check for any syntax errors in your optimization function. This is a very useful function when you use your own functions; if anything is incorrectly set up the test results will tell you where to start correcting your function.

   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.

5 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.

1 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` directory.

2 Select **File > New > Optimization**.

3 The newly registered optimization appears in the list of algorithm names. Select `Tutorial_Optimization` from the list. Click **Next**.

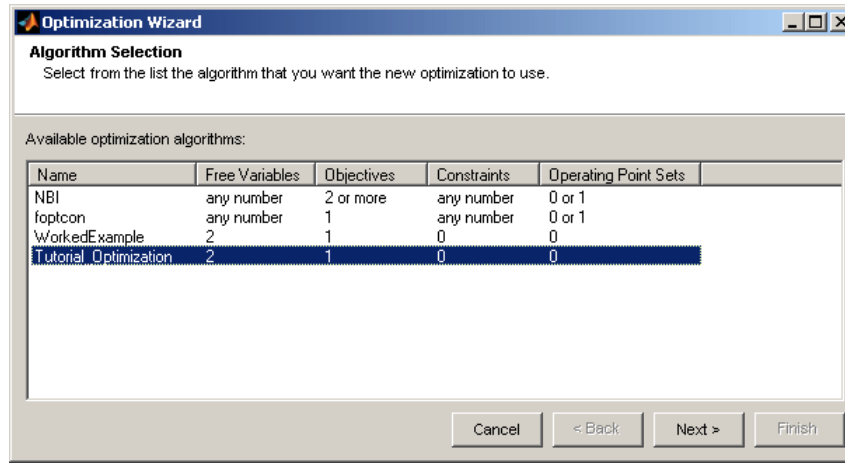

4 Match the variables as shown.

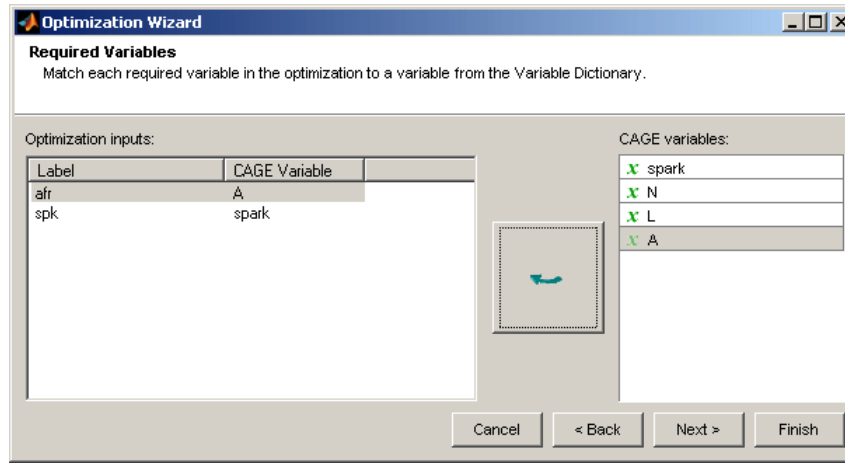

Click **Next**.

5 Match the `Torque` model to the `tuttq` CAGE model as shown.

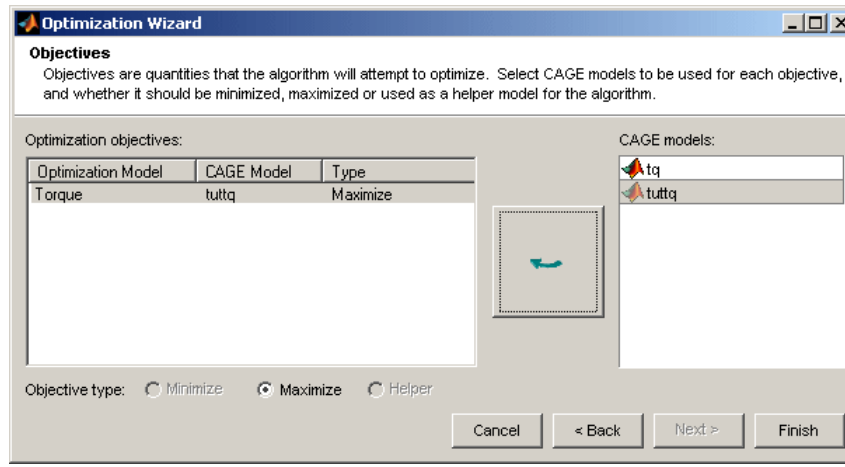

Click **Finish**.

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	☑	12.78	23.768	1000	0.2	9.07
2	☑	12.78	18.179	1000	0.3	20.319
3	☑	12.78	14.261	1000	0.4	31.567
4	☑	12.78	12.014	1000	0.5	42.816
5	☑	12.78	11.441	1000	0.6	54.064
6	☑	12.78	12.534	1000	0.7	65.313
7	☑	12.78	27.476	2000	0.2	9.742
8	☑	12.78	21.887	2000	0.3	20.99
9	☑	12.78	17.969	2000	0.4	32.238
10	☑	12.78	15.722	2000	0.5	43.487
11	☑	12.78	15.147	2000	0.6	54.735
12	☑	12.78	16.243	2000	0.7	65.984
13	☑	12.78	31.185	3000	0.2	9.342
14	☑	12.78	25.597	3000	0.3	20.591
15	☑	12.78	21.683	3000	0.4	31.839
16	☑	12.78	19.431	3000	0.5	43.087
17	☑	12.78	18.856	3000	0.6	54.336
18	☑	12.78	19.953	3000	0.7	65.584
19	☑	12.78	34.891	4000	0.2	7.872
20	☑	12.78	29.305	4000	0.3	19.12
21	☑	12.78	25.385	4000	0.4	30.369
22	☑	12.78	23.139	4000	0.5	41.617
23	☑	12.78	22.563	4000	0.6	52.866
24	☑	12.78	23.659	4000	0.7	64.114
25	☑	12.78	38.614	5000	0.2	5.331
26	☑	12.78	33.013	5000	0.3	16.58
27	☑	12.78	29.093	5000	0.4	27.828
28	☑	12.78	26.847	5000	0.5	39.077
29	☑	12.78	26.27	5000	0.6	50.325
30	☑	12.78	27.367	5000	0.7	61.574
31	☑	12.78	42.321	6000	0.2	1.72
32	☑	12.78	36.721	6000	0.3	12.969
33	☑	12.78	32.801	6000	0.4	24.217
34	☑	12.78	30.555	6000	0.5	35.465
35	☑	12.78	29.98	6000	0.6	46.714
36	☑	12.78	31.075	6000	0.7	57.962

A

Examples

Use this list to find examples in the documentation.

Case Studies

Chapter 2, “Gasoline Engine Calibration Case Study”
Chapter 3, “Diesel Engine Calibration Case Study”

Commandline MBC

Chapter 4, “Command-Line Interface to the Model-Based Calibration Toolbox Product”

Model Browser

Chapter 5, “Tutorial: Model Quickstart ”
Chapter 6, “Tutorial: Design of Experiment”
Chapter 7, “Tutorial: Data Editor”

CAGE

Chapter 8, “Tutorial: Feature Calibration”
Chapter 9, “Tutorial: Tradeoff Calibration”
Chapter 10, “Tutorial: Data Sets”
Chapter 11, “Tutorial: Filling Tables from Data”
Chapter 12, “Tutorial: Optimization and Automated Tradeoff”

Index

A

automated tradeoff tutorial 12-38

C

case study
  diesel 3-1
  gasoline 2-1
classical design tutorial 6-27
command-line interface to Model-Based Calibration Toolbox 4-1

D

data editor tutorial 7-1
data filtering
  data editor tutorial 7-2
  gasoline example 2-18
data loading
  data editor tutorial 7-3
  gasoline example 2-18
  quickstart example 5-4
data sets tutorial 10-1
design
  classical 6-27
  optimal 6-10
  space-filling 6-37
  tutorial 6-1
design of experiment tutorial 6-1
diesel case study 3-1

F

feature calibration tutorial 8-1

G

gasoline case study 2-1
global models
  specifying inputs to 5-15
  specifying type of 5-17

L

local models
  specifying inputs to 5-11
  specifying type of 5-12

M

models
  feature calibration tutorial 8-7
  import file to CAGE 2-67
  introducing two-stage 2-6
  quickstart two-stage 5-2
  tradeoff tutorial 9-6
  two-stage illustrations 2-27
multiobjective optimization tutorial 12-19

O

optimal design tutorial 6-10
optimization
  automated tradeoff tutorial 12-38
  multiobjective tutorial 12-19
  single-objective tutorial 12-5
  sum tutorial 12-33
  user-defined tutorial 12-41
  worked example tutorial 12-41
optimization constraints
  diesel example 3-42
optimization tutorial 12-1
optimized calibration
  diesel example 3-31
  gasoline example 2-64

Q

quick start modeling tutorial 5-1

R

requirements, system 1-8

S

single-objective optimization tutorial 12-5
space-filling design tutorial 6-37
sum optimization tutorial 12-33
system requirements 1-8

T

tradeoff calibration tutorial 9-1
tutorial 12-1
    data editor 7-1
    data sets 10-1
    design of experiment 6-1
    diesel case study 3-1
    feature calibration 8-1
    filling tables from data 11-1
    gasoline case study 2-1
    quick start 5-1
    tradeoff calibration 9-1
two-stage models
    gasoline case study 2-6
    illustrations 2-27
    quick start 5-2

U

user-defined optimization tutorial 12-41