

Model-Based Calibration Toolbox

For Use with MATLAB® and Simulink®

- Computation
- Visualization
- Programming
- Simulation

Model Browser User's Guide

Version 3



How to Contact The MathWorks:



www.mathworks.com
comp.soft-sys.matlab
www.mathworks.com/contact_TS.html

Web
Newsgroup
Technical Support



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

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



508-647-7000 (Phone)



508-647-7001 (Fax)



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

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

Model-Based Calibration Toolbox Model Browser User's Guide

© COPYRIGHT 2001–2005 by The MathWorks, Inc.

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

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

Trademarks

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

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

Patents

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

Revision History

December 2001	Online only	New for Version 1.0 (Release 12.1)
August 2002	Online only	Revised for Version 1.1 (Release 13)
May 2003	Online only	Revised for Version 2.0 (Release 13+)
June 2004	Online only	Revised for Version 2.1 (Release 14)
June 2004	Online only	Revised for Version 2.1.1 (Release 14+)
November 2005	Online only	Revised for Version 3.0 (Release 14SP3+)

Getting Started

1

What Is the Model Browser?	1-3
Starting the Model Browser	1-3
How to Use This Manual	1-4
Modeling with the Model-Based Calibration Toolbox ..	1-5
How To Set Up a One-Stage Model	1-6
How To Set Up a Two-Stage Model	1-6
Creating Multiple Models To Compare	1-8
Project Level: Startup View	1-11
All Models Pane	1-12
Data Sets Pane	1-12
Notes Pane	1-13
Test Plans List Pane	1-13
Project Level: Toolbar	1-15
Project Level: Menus	1-16
Model Tree	1-19
Tree Structure	1-20
Icons: Curves, Worlds, and Houses	1-21
Icons: Blue Backgrounds and Purple Worlds	1-22
Test Plans	1-25
Creating a New Test Plan	1-25
Creating New Test Plan Templates	1-26
Test Plan Level	1-29
Summary Statistics	1-32
Test Plan Level: Toolbar and Menus	1-34

The Design Editor	2-3
Opening the Design Editor	2-3
Design Styles	2-4
Design Editor Displays	2-5
Design Editor Toolbar and Menus	2-9
Creating a Classical Design	2-16
Classical Design Browser	2-17
Creating a Space-Filling Design	2-19
Space-Filling Design Styles	2-20
Setting Up a Space-Filling Design	2-23
Creating an Optimal Design	2-25
Optimal Design: Initial Design Tab	2-26
Optimal Design: Candidate Set Tab	2-29
Optimal Design: Algorithm tab	2-31
Averaging Optimality Across Multiple Models	2-33
Adding Design Points	2-35
Merging Designs	2-37
Fixing, Deleting, and Sorting Design Points	2-39
Exporting and Importing Designs	2-42
Applying Constraints	2-43
Constraint Types	2-44
Importing Constraints	2-51
Prediction Error Variance Viewer	2-54
Display Options	2-56
Prediction Error Variance	2-57
Prediction Error Variance for Two-Stage Models	2-59

Design Evaluation Tool	2-62
Table Options	2-63
Design Matrix	2-64
Full FX Matrix	2-64
Model Terms	2-64
Z2 Matrix	2-65
Alias Matrix	2-65
Z2.1 Matrix	2-65
Regression Matrix	2-65
Coefficient Information	2-66
Standard Error	2-68
Hat Matrix	2-68
$ XX $	2-68
Raw Residual Statistic	2-69
Degrees of Freedom Table	2-69
Design Evaluation Graphical Displays	2-71
Export of Design Evaluation Information	2-71

Data

3

Data	3-3
Tip for Fast Data Editor Access	3-4
The Data Editor	3-6
Data Editor Toolbar and Menus	3-9
Data Loading and Merging	3-13
Loading Data from File	3-13
Data Editor Views	3-19
2-D Plots	3-20
3-D Plots	3-21
Multiple Data Plots	3-21
Cluster Plots	3-22
List Views — Variables, Filters, Test Filters, Test Notes, and Clusters	3-23
Table View	3-23
Notes View	3-24

The Variable Editor	3-25
New Variables	3-25
The Filter Editors	3-27
New Filters	3-27
Storage	3-29
Test Groupings	3-33
Data Wizard	3-36
Step 1: Select Data Set	3-36
Step 2: Select Input Signals	3-36
Step 3: Select Response Models	3-38
Step 4: Set Tolerances	3-38
Matching Data to Designs	3-41
How to Use the Cluster Plot View	3-43
The Tolerance Editor	3-45
What Will Happen to My Data and Design?	3-47
Data Loading Application Programming Interface	3-49

Modeling

4

Setting Up Models	4-3
What Models Are Available?	4-4
Setting Up Inputs	4-4
Global Model Setup	4-6
Global Linear Models: Polynomials and Hybrid Splines ..	4-8
Global Model Class: Radial Basis Function	4-14
Global Model Class: Hybrid RBF	4-20
Global Model Class: Multiple Linear Models	4-21
Global Model Class: Free Knot Spline	4-22
Global Model Class: Neural Network	4-24
Global Model Class: User-Defined and Transient Models	4-26

Local Model Setup	4-26
Local Model Class: Polynomials and Polynomial Splines ..	4-28
Local Model Class: Truncated Power Series	4-32
Local Model Class: Free Knot Spline	4-34
Local Model Class: Growth Models	4-36
Local Model Class: Linear Models	4-42
Local Model Class: Average Fit	4-44
Local Model Class: Multiple Models	4-45
Local Model Class: User-Defined Models	4-49
Local Model Class: Transient Models	4-55
Covariance Modeling	4-66
Correlation Models	4-67
Transforms	4-68
Designing Experiments	4-69
New Response Models	4-71
Datum Models	4-72
Local Level	4-74
Local Special Plots	4-76
Local Scatter Plots	4-77
Response Features List	4-78
Diagnostic Statistics Pane	4-79
Pooled Statistics	4-79
Test Notes Pane	4-81
Data Tab	4-81
Local Level: Toolbar	4-82
Local Level: Menus	4-83
Outlier Selection Criteria	4-89
Global Level	4-93
Global Special Plots	4-95
Global Scatter Plots	4-96
Removed Data Pane	4-98
Models List	4-98
Global Level: Toolbar	4-99
Global Level: Model-Specific Tools	4-100
Global Level: Menus	4-103
Build Models Dialog	4-106
Selecting Models	4-111
Select Button	4-111

Model Selection Guide	4-113
Using Information Criteria to Compare Models	4-118
Model Selection Window	4-120
Tests View	4-122
Predicted/Observed View	4-125
Response Surface View	4-127
Likelihood View	4-131
RMSE View	4-133
Residuals View	4-135
Cross Section View	4-136
MLE	4-141
Calculating MLE	4-141
MLE Settings	4-143
Response Level	4-145
Response Level: Toolbar and Menus	4-147
Model Evaluation Window	4-148
Model Evaluation Using Other Data	4-150
Boundary Constraint Modeling	4-153
Creating A Boundary Constraint	4-154
Boundary Constraint Toolbar	4-157
Boundary Constraint Menus	4-159
Complete Constraint Fit Dialog	4-161
Stepwise	4-166
Using the Stepwise Regression Window	4-166
Automatic Stepwise	4-169
Stepwise in the Model Building Process	4-170
PRESS statistic	4-174
Box-Cox Transformation	4-176
Maximum Likelihood Estimation for Nonlinear	
Repeated Measures	4-180
Two-Stage Models for Engines	4-181
Local Models	4-182

Global Models	4-185
Two-Stage Models	4-186
Global Model Selection	4-188
References	4-189
Linear Regression	4-191
Toolbox Terms and Statistics Definitions	4-192
Exporting Models	4-195
What Is Exported?	4-198
Evaluating Models in the Workspace	4-200

Radial Basis Functions

5

Guide to Radial Basis Functions for Model Building ..	5-3
Types of Radial Basis Functions	5-5
Gaussian	5-5
Thin-Plate Spline	5-6
Logistic Basis Function	5-7
Wendland's Compactly Supported Function	5-7
Multiquadrics	5-9
Reciprocal Multiquadrics	5-10
Linear	5-10
Cubic	5-10
Fitting Routines	5-11
Center Selection Algorithms	5-12
Rols	5-12
RedErr	5-13
WiggleCenters	5-13
CenterExchange	5-13
Lambda Selection Algorithms	5-15
IterateRidge	5-15

IterateRols	5-16
StepItRols	5-18
Width Selection Algorithms	5-19
TrialWidths	5-19
WidPerDim	5-19
Tree Regression	5-21
Prune Functionality	5-25
Statistics	5-28
GCV Criterion	5-28
GCV for Ridge Regression	5-29
GCV for Rols	5-30
References	5-31
Hybrid Radial Basis Functions	5-32
Width Selection Algorithm: TrialWidths	5-32
Lambda and Term Selection Algorithms: Interlace	5-32
Lambda and Term Selection Algorithms: Two-Step	5-33
Tips for Modeling with Radial Basis Functions	5-34
Plan of Attack	5-34
How Many RBFs to Use	5-35
Width Selection Algorithms	5-36
Which RBF to Use	5-37
Lambda Selection Algorithms	5-38
Center Selection Algorithms	5-38
General Parameter Fine-Tuning	5-38
Hybrid RBFs	5-38
How to Find RBF Model Formula	5-39

Index

Getting Started

The following sections introduce the Model Browser part of the Model-Based Calibration Toolbox.

- | | |
|--|---|
| What Is the Model Browser? (p. 1-3) | The Model Browser part of the toolbox is a powerful tool for design of experiment and statistical modeling. |
| How to Use This Manual (p. 1-4) | A guide to the Model Browser reference chapters. |
| Modeling with the Model-Based Calibration Toolbox (p. 1-5) | An introduction to the process of making and using models in the Model-Based Calibration Toolbox. This section contains quick overview guides to setting up and comparing models, with links to the reference sections. |
| Project Level: Startup View (p. 1-11) | This section describes the functionality available at project level, where you begin. The different views of the Model Browser are project, test plan, local, global and response level. |
| Model Tree (p. 1-19) | It is very important to understand the structure and functions of the model tree to navigate the views of the Model Browser. |

Test Plans (p. 1-25)

You must use test plans as a basis for making and comparing models. This section explains what test plans are and how to use them.

Test Plan Level (p. 1-29)

This section describes the functionality available at test plan level, where you start building models.

What Is the Model Browser?

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

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

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 MATLAB®, Simulink®, or CAGE.

Starting the Model Browser

To start the application, type

```
mbcmodel
```

at the MATLAB® command prompt.

How to Use This Manual

This manual is the Model Browser User's Guide. See also the CAGE User's Guide for information on the other main interface of the Model-Based Calibration Toolbox.

Learning the Model Browser

See the Getting Started guide for tutorials and case studies.

Using the Model Browser

- Chapter 1, “Getting Started” provides an overview of how to navigate the Model Browser and where to find information, covering projects, test plans and the model tree.
- Chapter 2, “Designs” is a guide to constructing experimental designs using the Design Editor. You can set up, view, and compare design types: classical, space-filling, and optimal. You can define and apply constraints and export designs.
- Chapter 3, “Data” describes how to load, merge, filter, and view data using the Data Editor, including matching data to experimental designs. You can load data from files, or the workspace, or custom 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. You can customize the data loading interface with your own functions.
- Chapter 4, “Modeling” is a complete guide to setting up, evaluating, and verifying models using the Model Browser. This includes all the functionality available in the different model views of the Model Browser, descriptions and illustrations of model types, boundary modeling, and exporting models.
- Chapter 5, “Radial Basis Functions” is a guide to all aspects of using radial basis functions in modeling, from setup to the mathematical basis.

Modeling with the Model-Based Calibration Toolbox

Functionality is described in the order you see it during the process of model building. For a quick guide to setting up models, then searching for the best fit, see these overview pages:

- “How To Set Up a One-Stage Model” on page 1-6
- “How To Set Up a Two-Stage Model” on page 1-6
- “Creating Multiple Models To Compare” on page 1-8

The different views of the Model Browser are described in these sections:

- “Project Level: Startup View” on page 1-11
- “Test Plan Level” on page 1-29
- “Local Level” on page 4-74
- “Global Level” on page 4-93
- “Response Level” on page 4-145

To construct models, you must navigate using the model tree, use test plans, and load and manipulate data and designs. These topics are covered in these sections:

- “Model Tree” on page 1-19
- “Test Plans” on page 1-25
- Chapter 3, “Data”
- Chapter 2, “Designs”

Model construction, evaluation and export are covered in these sections:

- “Selecting Models” on page 4-111
- “Model Evaluation Window” on page 4-148
- “Exporting Models” on page 4-195

How To Set Up a One-Stage Model

The following steps are necessary to set up a one-stage model:

- 1** From the project node, create a new one-stage test plan. See “Creating a New Test Plan” on page 1-25.
- 2** Select the new node in the model tree to change to the test plan level.
- 3** Set up the inputs and model type by double-clicking the Inputs block and the Model block in the test plan diagram. See “Setting Up Inputs” on page 4-4 and “Setting Up Models” on page 4-3.
- 4** At this point, you might want to design an experiment. See “The Design Editor” on page 2-3.
- 5** From the test plan node, load a new data set to use. Choose **TestPlan > Select Data** which opens the Data Wizard. Click **Load new data set** and the Data Import Wizard appears. See “Loading Data from File” on page 3-13. Use the Data Wizard to match up data signals with model inputs and outputs.
- 6** Dismissing the Data Wizard opens the Data Editor, where you can select data for modeling and match data to designs. See “The Data Editor” on page 3-6.

On closing the Data Editor the model fit is calculated and the new model node appears in the model tree. Select the new node to view the model fit.

Functionality available for viewing and refining the model fit is described in “Global Level” on page 4-93 and “Selecting Models” on page 4-111.

Once you have built a single model, you should create more models for comparison, to search for the best fit. You can follow the guidelines in “Creating Multiple Models To Compare” on page 1-8.

How To Set Up a Two-Stage Model

The following steps are necessary to set up a two-stage model:

- 1 From the project node, create a new two-stage test plan. See “Creating a New Test Plan” on page 1-25.
- 2 From the test plan node, set up the inputs and models at the local and global stages. See “Setting Up Inputs” on page 4-4 and “Setting Up Models” on page 4-3.
- 3 At this point, you might want to design an experiment. See “The Design Editor” on page 2-3.
- 4 From the test plan node, load the data set you want to use. Select **TestPlan > Select Data**. See “Loading Data from File” on page 3-13.

This opens the Data Wizard, where you can load a data set, match data signals to model variables and then set up the response model.

- 5 On completing the Data Wizard, the Data Editor opens. Here you can select data for modeling and match data to designs. See “The Data Editor” on page 3-6. Close the Data Editor and click **Yes** to accept the data for modeling.

Note On closing the Data Editor, the local and global models are calculated.

- 6 At the local node, you can view the fit of the local models to each test, and you can also view the global models at the response feature nodes.
- 7 The two-stage model is not calculated until you use the **Select** button (from the local node, in the **Response Features** pane) and choose a model as best (even if it is the only one so far), *unless* you go straight to MLE. See below.

See “Selecting Models” on page 4-111.

Note At this point, the two-stage model is calculated, and the icon changes at the local node to reflect this. See “Icons: Curves, Worlds, and Houses” on page 1-21.

- 8 You are prompted to calculate the maximum likelihood estimate (MLE) at this point. You can do this now, or later by selecting **Model > Calculate MLE**. See “MLE” on page 4-141 for a detailed explanation.

Note If there are exactly enough response features for the model, you can go straight to MLE calculation without going through the **Select** process. The **MLE toolbar** button and the **Model > Calculate MLE** menu item are both active in this case. If you add new response features, you cannot calculate MLE until you go through model selection to choose the response features to use.

See “Two-Stage Models for Engines” on page 4-181 for a detailed explanation of two-stage models.

Double-click the Model blocks of the block diagram or select the **TestPlan > Set Up Model** menu item. MBC supports a wide range of models.

For model descriptions, see “Global Model Setup” on page 4-6 and “Local Model Setup” on page 4-26.

Once you have built a single model, you should create more models for comparison, to search for the best fit. You can follow the guidelines in the next section, “Creating Multiple Models To Compare” on page 1-8.

Creating Multiple Models To Compare

Once you have fitted and examined a single model (either one- or two-stage), you will normally want to create more models to search for the best fit. You can:

- Create individual new models.
 - You can create new child nodes by clicking the **New** button from any modeling node. Then select the new child node and choose **Model > Set Up Model** to reach the Model Setup dialog and change the type and settings. You can do this for multiple child nodes to create a selection of different model types fitted to the same data.
- Create a template to save a variety of model settings for reuse.

- From any global or one-stage model with child nodes, select **Model > Make Template**. You can save the child node model types of your currently selected modeling node as a template. You can then use the Build Models dialog to find your user-defined templates and quickly build all those model types again for any global model you choose (see below).
- From any global model node (before calculating MLE), click **Build Models** in the toolbar. You can save a template containing whatever models you choose by selecting **New** and adding the model types you want.
- Use the **Build Models** function to create a selection of models at once. You can use the **Browse** button to find your stored templates that are not in the default directory, and there are predefined templates for polynomials and RBF kernels. All the child node model types in the template you select are built as child nodes of the currently selected global model.

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 that fit well to the data for each of your global models.

We provide a detailed tutorial example to guide you through using these techniques to make a number of models to compare. See the modeling tutorial section “Creating Multiple Models to Compare” in the Getting Started documentation. This tutorial section requires you to have completed the previous steps in the Quick Start tutorial, so that you have an initial model as a starting point.

Note that you can click **New** from any modeling node to create new child nodes.

- For all global model nodes (one- and two-stage) a new child node appears that is a copy of the parent, and you can select **Model > Set Up Model** to change the type and settings.
- For the other two-stage modeling nodes (local and response models) you also get a new child node but the procedure is slightly different — see the tutorial example for details.

Once you have a variety of models to compare, you should use the diagnostic statistics and powerful visual aid plotting capabilities of the Model Browser to help you decide which models are best. For a description of the views and statistics available in each modeling view, see:

- “Selecting Models” on page 4-111
- “Model Evaluation Window” on page 4-148
- “Local Level” on page 4-74
- “Global Level” on page 4-93
- “Response Level” on page 4-145

For linear models, make use of the Stepwise functions (open the Stepwise window for existing models, and/or choose a **Stepwise** option during model setup) to refine your models and remove less useful model terms. Make sure you examine outliers but do not automatically remove them without good reason. Pay attention to the diagnostic statistics to help you find the best models. The following sections describe how to use Stepwise to make better models and how to understand the diagnostic statistics:

- “Model Selection Guide” on page 4-113 for guidelines.
- “Stepwise” on page 4-13 for an introduction to Stepwise.
- “PRESS statistic” on page 4-174 — See this section for guidelines on what to look for in the statistics to indicate a good fit.

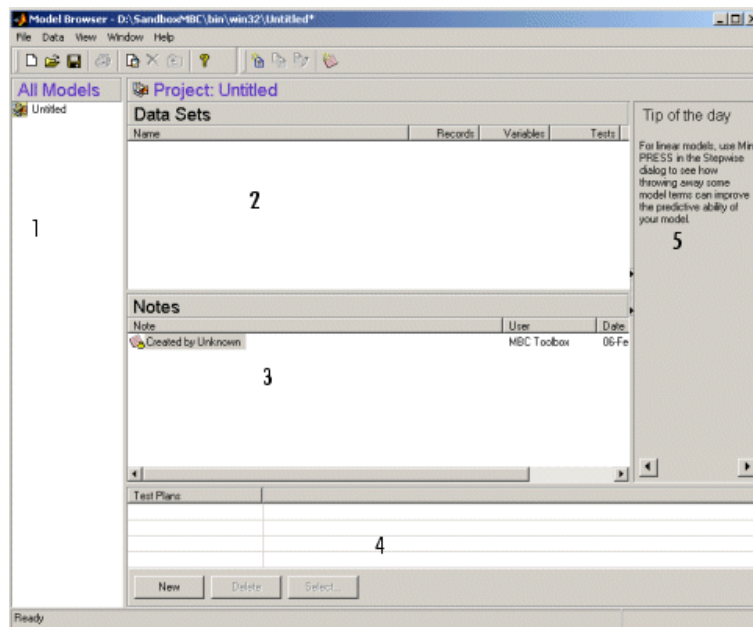
Project Level: Startup View

When you open the Model Browser part of the Model-Based Calibration Toolbox there is a single node, the project (named Untitled), in the model tree. This node is automatically selected.

When the project node in the model tree is selected at any time, the following functionality is available. This state is called *project level*. When you start you are automatically at project level, as there are not yet any other nodes to select.

See “Modeling with the Model-Based Calibration Toolbox” on page 1-5 to find information about using the Model Browser.

Note The node selected in the model tree determines what appears in the menus and panes of the rest of the Model Browser.



All Models Pane

Labeled 1, see “Project Level: Startup View” on page 1-11. This pane contains a hierarchical structure showing all the models created within the current project. See “Model Tree” on page 1-19 for a detailed description of the information contained in this pane.

Data Sets Pane

Labeled 2 “Project Level: Startup View” on page 1-11. All data sets loaded in the current project are displayed in the **Data Sets** pane (whether in use for modeling or not).

You can select a data set (by clicking it) and then

- Delete it by pressing the **Delete** key.
- Rename it, by clicking again or pressing **F2** (as when selecting to rename in Windows Explorer), then editing the name.
- Open it by double-clicking. Double-clicking a data set opens the Data Editor; unless it is already associated with a test plan, see below. See “The Data Editor” on page 3-6.

Note All data sets loaded are visible at the project node and appear in the **Data Sets** pane. However, they are not necessarily used by any test plan child nodes of that project until you select them within a particular test plan. For example, with a data set loaded at the project node, when you switch to a new test plan node, the **Data Sets** pane at top right displays ‘No Data is selected’ until you use the Data Wizard to attach data to that test plan. See “Data Wizard” on page 3-36.


The same data set can be used by many test plans within a project, although each individual test plan can only ever use one data set.

When you have associated a data set with a test plan, a new data set icon (with a name specific to that test plan) appears here in the Data Sets pane. The same data set may be used by several test plans, at which you may have applied different filters, groupings or edits, and so each time you associate a data set with a new test plan a new icon appears here. You cannot edit

these test plan-specific data sets at project level, you must edit them from the associated test plan.

Notes Pane

Labeled 3, see “Project Level: Startup View” on page 1-11. The **Notes** pane contains a list box showing all previous notes on the current project. You use notes to record changes within a project, perhaps by different users, over time.

- You add new notes by clicking the **Add new note** button  in the toolbar, or by pressing the **Insert** key after selecting the **Notes** pane by clicking. Notes automatically have the user login name and the date.
- You edit notes (only the user that created them can edit them; user names must match) by select-clicking or by pressing **F2** when selected (as when selecting to rename in Windows Explorer). Edited notes have updated time and date fields.
- You remove notes by selecting them and pressing **Delete** (but only the same user that created them can delete them).
- Notes are automatically added to the project when it is modified (for example, the initial “Created by <username>” note). These notes (listed as user “info”) cannot be deleted or edited.

Test Plans List Pane

Labeled 4, see “Project Level: Startup View” on page 1-11. You generate new test plans from the **Test Plans** list pane by clicking the **New** button. See “Test Plans” on page 1-25.

This pane is the **Test Plans** list pane at startup but changes depending on the level in the model tree that is selected. The list box always displays all the child nodes of whichever node is currently selected in the tree in the **All Models** pane, and always contains three buttons: **New**, **Delete**, and **Select**.

Note Double-clicking any item within this pane changes the view directly to that node. (This is equivalent to selecting that node in the model tree.) You can also use the **Delete** and **Insert** keys to remove or add new test plans (select a test plan first).

- The **Test Plans** list becomes the **Response Models** list, the **Local Models** list, the **Response Features** list, and the **Models** list as you select the nodes at subsequent levels of the model tree. In each case this pane displays the immediate child nodes of the current node selected. You can use the buttons to delete selected nodes or create new nodes.
- The feature added by clicking **New** always corresponds to the list items. For example, clicking **New** when the pane shows a list of test plans adds a new test plan. Clicking **New** when the pane shows a list of response features opens the Response Features dialog, as shown in the following example. The response features you can add are model-specific. This example shows the response features available for a polyspline model.



For example, if you choose $f(x + \text{datum})$ and enter 10 in the **Value** edit box, the new response feature tracks the datum +10. For a torque/spark polyspline model, the datum is MBT (maximum brake torque); so the new response

feature is MBT + 10 degrees of spark angle. This allows you to create response features that have engineering interest.

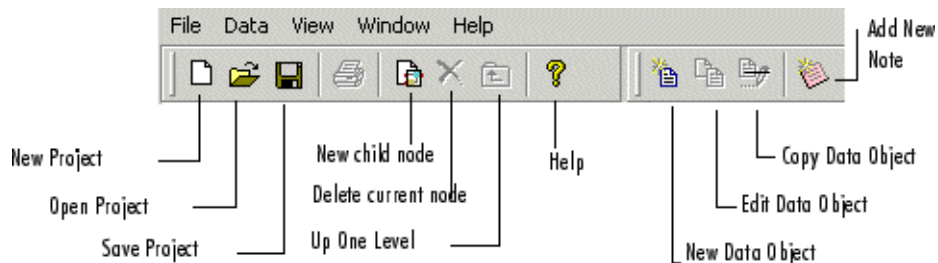
The response features available depend on the model type. For more details on which response features are available, see each model type under “Local Model Setup” on page 4-26.

- You can use the **Select** button to select the best child node, but only when the child nodes are local models, response features, or submodels of response features. In each case clicking the **Select** button takes you to a selection window where you can compare the child nodes. See “Selecting Models” on page 4-111.

Tip of the Day

Labeled 5, see “Project Level: Startup View” on page 1-11. Hints about using the Model Browser appear here. You can scroll through more tips using the buttons at the bottom, and you can snap this pane closed or open by clicking the “snapper point” where the cursor changes if you roll the mouse over it.

Project Level: Toolbar



This is how the toolbar appears when you first start the toolbox. The last two Data buttons are grayed out; the **Edit data object** and **Copy data object** buttons are not enabled until a data set has been loaded.

- All the toolbar items are duplicated under the menus except New Note.
- For the **Project** buttons, see the File menu.
- For the **Data** buttons, see the Data menu.

- New Note adds a note to the **Notes** pane.
- The **New** and **Delete** node buttons are duplicated in the **File** menu. In both cases, their function depends on the node selected in the model tree. In every case, **New** generates a new child node of the one currently selected, and **Delete** removes the current node (and all its children).
- The **Up One Level** button moves the current selection (and hence all the views) one level up the model tree. For example, if a test plan node is selected, clicking this button moves one level up to the project node.

Two buttons, **Delete** and **Up One Level**, are grayed out at startup because the default selection in the model tree is the project node, so there are no levels above, and you cannot delete the project node (although you can replace it with a new one).

- The print icon is only enabled in views with plots, for example, the local node, response feature nodes, and response nodes *after* selection of a best two-stage models (response nodes are blank until then).
- The **Help** button opens the Help Browser and displays the appropriate help documentation for the current view in your Model Browser.

Project Level: Menus

File Menu

Note The **File** menu remains constant in each Model Browser view. The **New child node** function always creates a new child node, and the **Delete current node** function always deletes the current node. These change according to which node in the model tree is currently selected.

- **New Project** opens a new project file. You are prompted to save or lose the current project.
- **Open Project** opens a file browser to select the project to open.
- **Save Project** and **Save Project As** save the project with all the models it contains as a .mat file.

- **New Test Plan** opens a dialog with the choice of One-Stage or Two-Stage test plans, or you can browse for other test plans. The **New** (child node) menu option always creates a new child node of whichever node is selected in the model tree.

At startup, the project node is automatically selected, so the appropriate child node is a new test plan.

Note File > New changes depending on which node in the model tree is selected. In each case the option offers a new child node immediately below the one currently selected, that is, a **New Test Plan** (if a project node is selected), a **New Response Model** (from a test plan) or a **New Model** child node (from a one-stage response). For two-stage models you can add a **New Local Model** (from a response node), a **New Response Feature** (from a local node) and a **New Model** from a response feature node.

- **Export Models** brings up the Export Models dialog. This allows you to export any models selected in the tree (along with their child nodes, in some cases) to the MATLAB workspace, to file for importing into CAGE, or to Simulink. See “Exporting Models” on page 4-195.
- **Delete “Untitled”** Like the **New** item in this menu, this option changes depending on which node in the model tree is selected. This menu item deletes whichever node is currently selected in the model tree (along with any child nodes), and the appropriate name appears here.
- **Clean Up Tree** From any modeling node where a best model has been selected (from the child nodes), you can use this to delete all other child nodes. Only the child nodes selected as best remain.
- **Preferences** brings up the MBC File Preferences dialog, in which you can specify default locations for projects, data, models, and templates. You can also edit and save user information: name, company, department and contact details. This information is saved with each project level note, and you can use this to trace the history of a project.
- **Print** is only enabled in views with plots — not test plan or project level.
- **1,2,3,4:** A list of the four most recent project files, including their pathnames.

- **Close** Exits from the Model Browser part of the toolbox (CAGE and MATLAB are unaffected).

Data Menu

- **New Data** — Opens the Data Editor. See “The Data Editor” on page 3-6.
- **Copy Data** — Copies the selected data set.
- **Edit Data** — Opens the Data Editor to enable data editing.
- **Delete Data** — Deletes the selected data set.

View Menu

- **Tip of the Day** — Choose to display or hide the tips pane.
- **Note Information** — Opens a dialog where you can decide which categories of information to display for each project note. You can specify user information for display with notes by using **File > Preferences**.

Window Menu

Depending on which toolbox windows are open, a list appears under this menu and whichever window is selected is brought to the front. The **Window** menu remains throughout the Model Browser.

Help Menu


The **Help** menu remains consistent throughout the Model Browser.

- **MBC Help** — Opens the Model-Based Calibration Toolbox Roadmap with links to the help tutorials and the indexed help pages.
- **Context Help** — Depending on what part of the Model Browser is currently active, **Context Help** links to different places in the Help files.
- **About MBC** — Displays version notes.

Model Tree

The tree in the **All Models** pane displays the hierarchical structure of the models you have built. Views and functionality within the browser differ according to which node is selected.

The following is an example of a model tree.



The elements of the tree consist of the following:


- 1 Project node
- 2 Test plan node
- 3 Response node
- 4 Local node
- 5 Global nodes — All one-stage model nodes are global models. For two-stage models, global models are fitted to response features of the local models. Each step down in the tree is a *child node* of the node above. Global models are child nodes of local nodes and so on.

Note The selected node in the tree governs the model that is displayed in the various other panes of the browser and which menu items are available. The selected node governs the level displayed: project level, test plan level, and so on. The functionality of each level is described in the Help.

You can rename all nodes, as in Windows Explorer, by clicking again or by pressing **F2** when a node is selected.

There is a context menu available. When you right-click any node in the model tree, you can choose to delete or rename that node, or create a new child node.

Tree Structure



The preceding example shows a more extensive model tree, with two two-stage models as child nodes of a single response model.

There can be many models within (or under, as *child nodes* in the tree) each two-stage global node, or any one-stage model node.

There can also be many different response nodes within a single test plan, and each project can contain several different test plans. However, there is only one project node in the tree at any time.


Note You can only have one project open at any one time; if you open another, you are prompted to save or discard your current project.

You can add child nodes to all global models — several candidate models can be tried at each global node and the best selected. There is an example showing this at the end of the section on “Icons: Blue Backgrounds and Purple Worlds” on page 1-22 and the process is illustrated in the “Tutorial: Model Quickstart ” in the Getting Started documentation.

Icons: Curves, Worlds, and Houses


The icons are designed to give visual reminders of their function.

- Test plan icons have a tiny representation of the test plan diagram. You can see the one-stage and two-stage icons in the following example.
- The local model icon shows a curve over a house.
- Global model icons show a curve over a globe. All one-stage models are global models and for two stage models, all nodes below the local node are global models.



- The response node (empty until a two-stage model is calculated) has an icon that combines the local and global motifs — a curve over a house *and* a globe — to symbolize the two-stage process.
- When a two-stage model has been calculated, the icon at the local node changes to show the combination of local and global motifs.

Icons: Blue Backgrounds and Purple Worlds



Icon changes convey information about modeling processes.

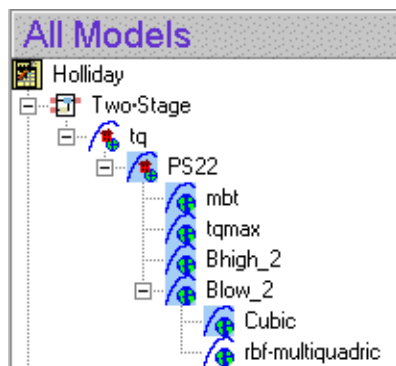
- When a model is selected as the *best model*, its icon changes color and gains a blue background, like the BSPLINE1 model in the preceding example.
- When the maximum likelihood estimate (MLE) is calculated and chosen as the best model, the associated model icon and all its child nodes (along with the plots of that model) become purple.

You can see this in the preceding example: the B Spline model and all its response features have purple curves, globes, and house, indicating that they are MLE models. The Poly3 model and its children have blue curves and globes and a red house, indicating that they are univariate models.

- Observe the other difference between the B Spline and the Poly3 icons: the B Spline has a blue background. This indicates that this is selected as best model, and is used to calculate the two-stage model at the response node, so the response node is also purple. If an MLE model (with purple worlds) is selected as best model and is used to create the two-stage model, the response node always reflects this and is also purple.
- Notice also that the response features all have blue backgrounds. This shows they are selected as best and are all being used to calculate the

two-stage model. In this case they are all needed. That is, a B Spline model needs six response features, and a Poly3 model requires four. If more response features are added, however, some combination must be selected as best, and the response features not in use do not have a blue background. An example is shown in “Tutorial: Model Quickstart ” in the Getting Started documentation.

In the following example you can see child nodes of a global model. You can try different models within a global model, and you must select one of the attempts as best. In this example you can see that Cubic is selected as best, because it has a blue background, so it is the model being used for the Blow_2 global model.



When a model is selected as best it is copied up a level in the tree together with the outliers for that model fit.

When a new global or local model is created the parent model and outliers are copied from the current level to the new child node. This gives a mechanism for copying outliers around the model tree.

A tree node is automatically selected as best if it is the only child, except two-stage models which are never automatically selected - you must use the Model Selection window.

If a best model node is changed the parent node loses best model status (but the automatic selection process will reselect that best model if it is the only child node).

Note Try the Quick Start tutorial in the Getting Started documentation to understand how to use the model tree. The last section, “Creating Multiple Models to Compare”, guides you through the process of creating a variety of models of different types and how to understand the information in the model tree. You need to complete the previous sections of the tutorial first, which guides you through setting up a single two-stage model to get started.

Test Plans

You need to select a test plan to construct one-stage or two-stage models.


You can select the one- or two-stage test plans provided, as described next. You can also use these to create your own test plan template so you can reuse the setup for one test plan with another set of data. See “Creating New Test Plan Templates” on page 1-26.

Creating a New Test Plan

To create a new test plan:

- Click **New** in the **Test Plans** pane (visible at startup and whenever the project node is selected in the model tree).


Alternatively, make sure the project node is selected first, and then do one of the following:

- Click the **New Test Plan** icon () in the toolbar.
- Select **File > New Test Plan**.
- Press **Insert** immediately after clicking the tree.

These steps all open a dialog with the choice of One Stage or Two Stage test plans, or you can browse for other test plans (as new templates can be created and saved). See “Local Level” on page 4-74.

A new test plan node appears in the model tree. To view the new test plan:

Change to test plan level:

- Select the node  in the tree by clicking it.


Alternatively, double-click the new test plan in the **Test Plans** pane, as in Windows Explorer.

The Model Browser changes to test plan level, showing the block diagram representations of models in the main display, and the **Test Plans** pane changes to the **Response Models** pane (empty until models are set up).

For the next steps in model construction, see

- “Setting Up Models” on page 4-3
- “Setting Up Inputs” on page 4-4
- “Loading Data from File” on page 3-13
- “Selecting Models” on page 4-111

Creating New Test Plan Templates

You build user-defined templates from existing test plans using the Make Template toolbar button  or **TestPlan > Make Template**.

The procedures for modeling engines for calibrations are usually repeated for a number of different engine programs. The test plan template facility allows you to reuse the setup for one test plan with another set of data. You can alter the loaded test plan settings without restriction.

A list of test plan templates is displayed when you build a new test plan. There are built-in templates for one- and two-stage models.

Test plan templates store the following information:

- Hierarchical model — Whether the model is one- or two-stage and the default models for each level.
- Designs — *If* they were saved with the template (check box in the Test Plan Template Setup dialog)

The design for one type of engine might or might not be appropriate for a similar type of engine. You can redefine or modify the design using the Design Editor.

- All response models (for example, torque, exhaust temperature, emissions) — *If* they were saved with the template (check box in the Test Plan Template Setup dialog)

- Numbers and names of input factors to models
- Model types (local and global)
- Summary Statistics for display (see “Summary Statistics” on page 1-32)
- No model child nodes are included, just the top level of the test plan (response models, and local and global models for two-stage models).

The response models are automatically built after you assign data to the test plan; see “Using Stored Templates” on page 1-27.

Saving a New Template

From the test plan node that you want to make into a template:

- Click the **Make Template** toolbar icon or choose **TestPlan > Make Template**.

The templates are stored in the directory specified in the **File > Preferences** dialog.

The Test Plan Template Setup dialog appears, in which you can change the name of the new template and choose whether to include designs and/or response models.

Using Stored Templates

- When you load a new test plan from the project node, any stored templates appear in the New Test Plan dialog.

From the project node, select **File > New Test Plan**, or use the **toolbar** button, or click **New** in the **Test Plans** pane.


- Selecting templates in the New Test Plan dialog displays all templates found in the directory specified in the Preferences dialog (**File** menu). Select by clicking to see the information on a particular template; the number of stages and factors is displayed in the **Information** pane. You can use the **Browse** button if the required template is not in the directory specified in the **File > Preferences** dialog.
- Click **OK** to use the selected test plan template. The new test plan node appears in the model tree.

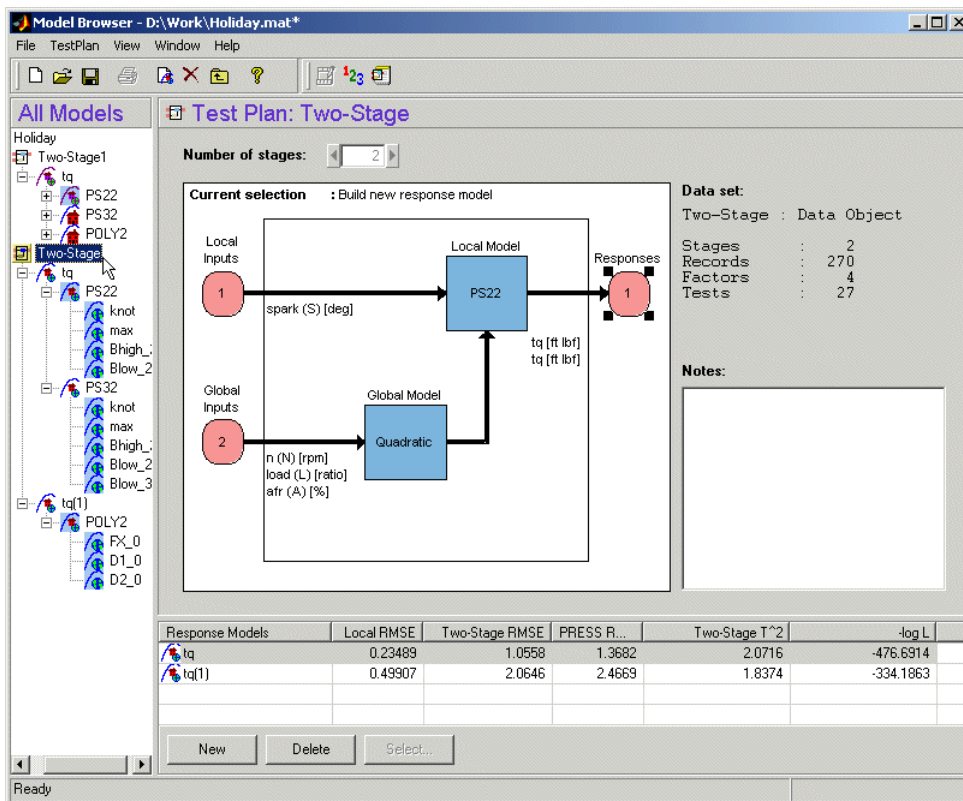
Use stored templates in exactly the same way as the default blank one- and two-stage templates. Models and input factors are already selected (including the response if that was saved with the template) so you can go straight to selecting new data to model. You can still change any settings and design experiments.

Double-click the Responses block to launch the Data Wizard and select data for the test plan. The response models are automatically built after selection of data.

Note The data selection process takes you through the Data Wizard. If any signal names in the new data do not match the template input factors, you must select them here, including the responses. If signal names match the factor names stored in the template, they are automatically selected by the Data Wizard, and you just click **Next** all the way to the end of the wizard. When you click **Finish** the response models are built automatically.

Test Plan Level

When you select a test plan node (with the icon ) in the model tree, this view appears.



Model Browser - D:\Work\Holiday.mat*

File TestPlan View Window Help

All Models

Holiday

- Two-Stage1
 - Two-Stage
 - PS22
 - PS32
 - POLY2
 - Two-Stage
 - PS22
 - PS32
 - POLY2

Test Plan: Two-Stage

Number of stages: 2

Current selection : Build new response model

Local Model

Local Inputs (1) → spark (S) [deg] → PS22 → Responses (1)

Global Model

Global Inputs (2) → n (N) [rpm], load (L) [ratio], afr (A) [%] → Quadratic → PS22

Data set:

Two-Stage : Data Object

Stages : 2
Records : 270
Factors : 4
Tests : 27

Notes:

Response Models	Local RMSE	Two-Stage RMSE	PRESS R...	Two-Stage T^2	-log L
Two-Stage	0.23489	1.0558	1.3682	2.0716	-476.6914
Two-Stage	0.49907	2.0646	2.4669	1.8374	-334.1863

New Delete Select...

Ready

This example is a two-stage model, the same model used in the Quick Start tutorial. All test plan nodes (one- and two-stage) bring up this view with a block diagram of the test plan. The diagram provides a graphical interface so you can set up inputs and set up models by double-clicking the blocks in the test plan diagram. These functions can also be reached using the TestPlan menu.


The diagram has the following functionality for setting up the stages in hierarchical modeling. At present MBC only supports one- and two-stage models. You can reach these functions via the right-click context menu (on each block) or the menus:

- 1** “Setting Up Inputs” on page 4-4 — Setting the number of inputs for each stage of the model hierarchy.
- 2** “Setting Up Models” on page 4-3 — Setting up the new default models for each stage in the model hierarchy.
- 3** “Designing Experiments” on page 4-69 — Using the Design Editor.

You can access the Design Editor via the right-click menus on the model blocks or the **TestPlan** menu (for a particular model — you must select a model or input block before you can design an experiment). **View > Design Data** also opens the Design Editor where you can investigate the design properties of the data. See .

- 4** “Loading Data from File” on page 3-13 and “New Response Models” on page 4-71

You can attach data to a new test plan by choosing **TestPlan > Select Data** or by double-clicking the Responses block in the diagram, which launches the Data Wizard (if the project already has data loaded). If a test plan already has data attached to it, you can reach the Data Selection views in

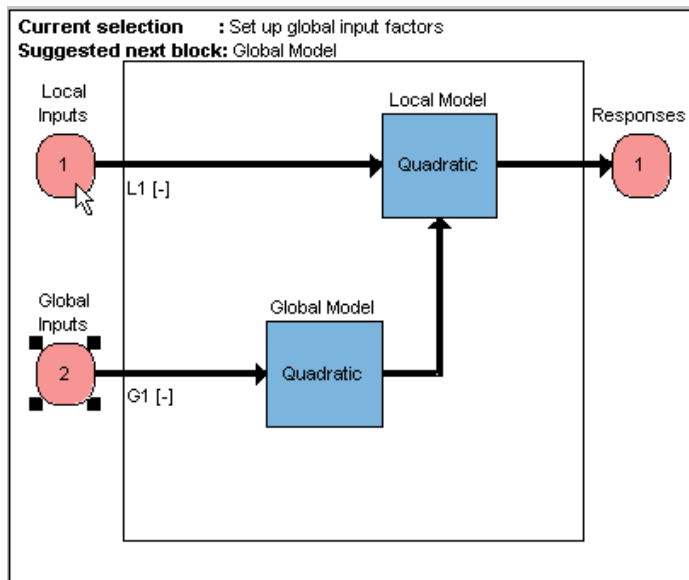
the Data Editor using the **Select Data** toolbar button  or the **TestPlan** menu item. Here you can select data for modeling and match data to a design. For example, after the design changes, new data matching might be necessary. See “Matching Data to Designs” on page 3-41 for details.

- 5** “Viewing Designs” on page 4-69 — Viewing the statistical properties of the collected data.
- 6** “Summary Statistics” on page 1-32 — Right-click on the global model block and select **Summary Statistics** to reach the Summary Statistics dialog. Here you can choose which summary statistics you want displayed to help you evaluate models.

The selected Model block is highlighted in yellow if a Setup dialog is open; otherwise it is indicated by blocks at the corners. The selected Model block indicates the stage of the model hierarchy that is used by the following menu choices:


- Set Up Model
- Design Experiment
- View Design Data
- View Model
- Summary Statistics

The block diagram in the test plan view represents the hierarchical structure of models. Following is an example of a two-stage test plan block diagram.



Other test plan level functionality:

- From the test plan level you can access the Boundary Constraint Modeling functionality. See “Boundary Constraint Modeling” on page 4-153.

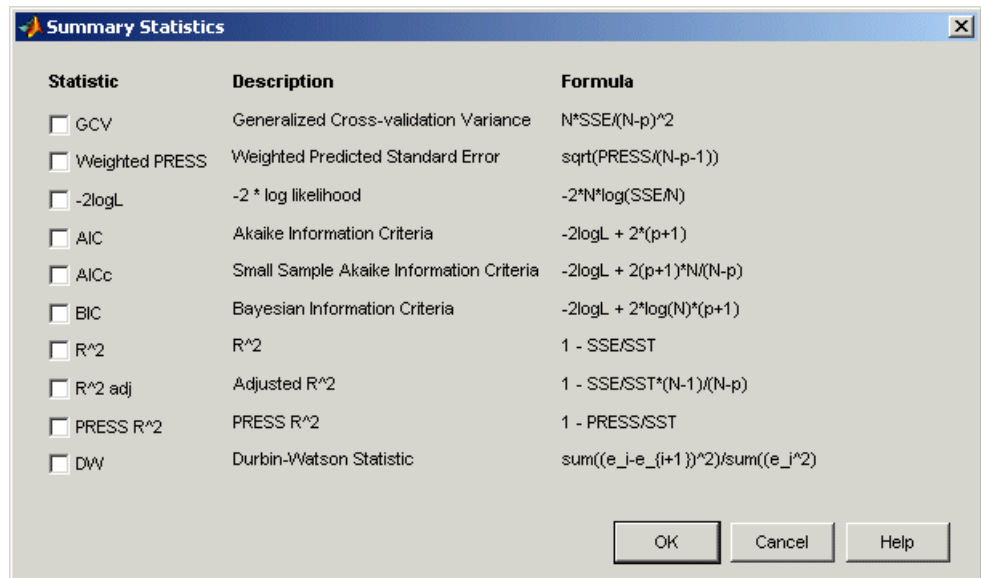
- You can save the current test plan as a template using the **TestPlan > Make Template** command or the toolbar button . This can be useful for speeding up creation of subsequent projects. See “Creating New Test Plan Templates” on page 1-26.

See also “Test Plan Level: Toolbar and Menus” on page 1-34

Summary Statistics

Right-click on the global model block and select **Summary Statistics** to reach the Summary Statistics dialog. Here you can choose which summary statistics you want displayed to help you evaluate models. The standard summary statistics are PRESS RMSE (for linear models only) and RMSE, and these are always displayed. You can choose additional statistics here by selecting the check boxes.

You can also reach the dialog using the **Model** menu from any global model node.



The choices you make in the Summary Statistics dialog determine what summary statistics are displayed in the global model views and the Model Selection window. These choices also determine what criteria you can use for automatically selecting best models in the Local Model Setup dialog for Multiple Models, and the Build Models dialog.

Available summary statistics are:

- GCV, Weighted PRESS, -2LogL , AIC, AICc (small sample), BIC, R^2 , R^2_{adj} , PRESS R^2 , and DW — full names and formulae are given in the dialog.
- The three Information Criteria (AIC, AICc (small sample) and BIC) are only available for comparison if exactly the same data is used for all models (same outliers, same y transform). The values of these will still appear for individual model display, but will disappear from the models list for comparison if, for example, you used a different Box-Cox transform.
 - The differences between two Information Criteria are of interest, not the absolute value. As a rule of thumb a difference of 2 or less means models are roughly the same quality of fit.
 - All the Information Criteria impose a penalty for overfitting, which increases linearly (approximately linearly for AICc) with the number of parameters in the model.
 - As a rough guide, typical models favoured by BIC are less complicated than those chosen by AIC or AICc.
 - AICc (small sample) is designed for cases where parameters/observations < 40 . Use AIC or AICc but not both.

See “Model Selection Guide” on page 4-113 for some guidelines and “Using Information Criteria to Compare Models” on page 4-118 for more statistical background.
- The Durbin-Watson statistic measures the correlation of adjacent observations, usually the residuals of a time series regression. A value of 2.0 for the Durbin-Watson statistic suggests that there is no time correlation.


e_i is the residual at time i .

Summary statistics are inherited from the test plan node or the parent node on creation of a child node. You can select them from individual global models, or select them from the test plan before building models if you want them to apply to all the models within the test plan.

Changes made using the **Model** menu from a global model node are applied immediately to the Summary table and the Models list pane in the global view. Summary statistic selections appear in the Models list pane only when there are child nodes to compare. Resetting a model to the default test plan model also resets the summary statistics.

For definitions of any of the terms in the Summary Statistics formulae, see “Toolbox Terms and Statistics Definitions” on page 4-192. Note that for ordinary least squares cases, ‘p’ is the number of parameters, but for non-ordinary least squares cases (rols and ridge least squares) ‘p’ is the effective number of parameters ($p = N - df$).

Test Plan Level: Toolbar and Menus



The eight buttons on the left (project and node management, plus the **Print** and **Help** buttons) appear in every view level. See “Project Level: Toolbar” on page 1-15 for details.

The right buttons change at different levels.

In the test plan level view, the right buttons are as follows:

- **Design Experiment** opens the Design Editor. Only available when a model or input has been selected in the test plan block diagram. You must specify the stage (local or global) you are designing for. See .
- **Select Data** opens the Data Wizard, or opens the Data Selection views in the Data Editor if data sets have already been selected. See Chapter 3, “Data”.
- **Make Template** opens a dialog to save the current test plan as a template, including any designs and response models. See “Local Level” on page 4-74.

Test Plan Level: Menus

File Menu

Only the **New** (child node) and **Delete** (current node) functions change according to the node level currently selected. Otherwise the **File** menu remains constant. See “File Menu” on page 1-16.

Window Menu. The **Window** menu remains throughout the Model Browser. It allows you to switch windows if there is more than one toolbox window open. See “Window Menu” on page 1-18.

Help Menu. The **Help** menu remains the same throughout the Model Browser. You can always reach the MBC Toolbox Help Roadmap by selecting **Help > MBC Help**. The context help takes you to relevant Help pages, and **Help > About MBC** shows the version notes. See “Help Menu” on page 1-18.

Test Plan Menu.

- **Set Up Inputs** — See “Setting Up Inputs” on page 4-4.
- **Set Up Model** — See “Setting Up Models” on page 4-3.

You can also reach these two functions by double-clicking the blocks in the test plan diagram, and both can only be used when a Model block is first selected in the diagram. You must specify the model to set up, local or global.

- **Design Experiment** — See “The Design Editor” on page 2-3.

This is also available in the toolbar and in the right-click context menu on the blocks in the test plan diagram.

- **Boundary Constraints** — Opens the Constraint Modeling window. See “Boundary Constraint Modeling” on page 4-153.
- **New Data** — Opens the Data Editor to load new data.
- **Select Data** — Opens the Data Selection views of the Data Editor.

You can reach both these functions with the toolbar **Select Data** button. If no data is selected, this button opens the Data Wizard, and if a data set is already selected, it takes you straight to the Data Selection views. See Chapter 3, “Data”.

- **Make Template** — Opens a dialog for saving the current test plan as a new template, with or without designs and response models. Same as the toolbar button. See “Creating New Test Plan Templates” on page 1-26.
- **Export Multimodels** — This provides a smooth interface with the Multimodel Tradeoff in the CAGE browser part of the Model-Based Calibration toolbox. Two global inputs are required (normally speed and load). This item is only enabled if you have set up a two-stage model with the correct number of inputs. This useful application for multiple models allows you to calibrate from local maps. See “Local Model Class: Multiple Models” on page 4-45 for details.

View Menu (Test Plan Level).

- **Design Data** — Opens the Design Editor. The view design facility enables you to investigate the statistical properties of the collected data. This provides access to all the Design Editor and design evaluation utility functions with the current design rather than the prespecified design (after data matching, the data points are used as the new design points). See “The Design Editor” on page 2-3.

For two-stage models, viewing level 1 designs creates a separate design for each test.

- **Model** — Opens a dialog showing the terms in the current model.
- Both of these are only available when a model or input block is selected in the test plan block diagram.

Designs

This section discusses the following topics:

The Design Editor (p. 2-3)	How to open the Design Editor and use the display functionality, plus an overview of reasons to use different design styles.
Creating a Classical Design (p. 2-16)	How to create classical designs. These are good for well-understood systems.
Creating a Space-Filling Design (p. 2-19)	How to create space filling designs. These are best for cases with low system knowledge.
Creating an Optimal Design (p. 2-25)	How to create optimal designs. These are best for cases with high system knowledge.
Adding Design Points (p. 2-35)	How to augment your design. You can add points optimally, randomly, or at specified points.
Merging Designs (p. 2-37)	How to merge the points from two or more designs.
Fixing, Deleting, and Sorting Design Points (p. 2-39)	How to remove points, fix points so they cannot be deleted and sort design points.
Exporting and Importing Designs (p. 2-42)	How to export and import your designs using file, comma separated values or workspace formats.

Applying Constraints (p. 2-43)

How to create constraints for your design. You can create ellipsoid, hyperplane, 1-D lookup table, or 2-D lookup table constraints, and you can import boundary constraints. Designs can have any number of geometric constraints placed upon them.

Prediction Error Variance Viewer (p. 2-54)

How to use the Prediction Error Variance viewer to examine the predictive power of your designs and models.

Design Evaluation Tool (p. 2-62)

Descriptions and definitions of all the information available in the Design Evaluation Tool.

The Design Editor

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.

There is a step-by-step guide to using the Design Editor in “Tutorial: Design of Experiment” in the Getting Started documentation.

Opening the Design Editor

You must first have a test plan before you can open the Design Editor.

- 1** From the startup (project) view of the Model Browser, click **New** and select a one or two-stage test plan. See “Project Level: Startup View” on page 1-11 and “Test Plans” on page 1-25 in the Modeling section.

You can design experiments at both stages, for local models and global models; for most two-stage models the global model is most appropriate for design of experiment.

- 2** Before you design an experiment we recommend that you set up your input variables, by double-clicking the Inputs blocks on the test plan diagram. See “Setting Up Inputs” on page 4-4.

You can choose the number of inputs for your model and set up their names and definitions, then you can design an experiment to collect data. It is much easier to understand your design points if they are labeled with the factor names. Also, if you do not set up model inputs first, then you can only create designs for the default number of variables (one).

- 3** If you want to use optimal designs, then the type of model you are going to use to fit the data is important, and you should choose a model type before opening the Design Editor. Double-click a model block in the test plan diagram to set up model types. 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, so in these cases you should pick your model

type before designing an experiment. See Chapter 4, “Modeling” to find out about model types in the Model-Based Calibration Toolbox.

If you have no idea what model you are going to fit, choose a space-filling design. Model type has no effect on designs that are space-filling or classical, so if you want to create these designs you can leave the model type at the default and open the Design Editor.

You can invoke the Design Editor in several ways from the “Test Plan Level” on page 1-29:

- 1 First you must select the stage (first/local or second/global) for which you want to design an experiment. Click to select the appropriate model block in the test plan diagram.
- 2 Right-click the model block and select **Design Experiment**.

Alternatively, click the **Design Experiment** toolbar icon .

You can also select **TestPlan > Design Experiment**.

For an existing design, **View > Design Data** also launches the Design Editor (also in the right-click menu on each Model block). This shows the selected data as a design.

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.

Classical designs (including full factorial) are very well researched and are suitable for simple regions (hypercube or sphere). See “Creating a Classical Design” on page 2-16.

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 a way as to maximize coverage of the factors’ ranges as quickly as possible. See “Creating a Space-Filling Design” on page 2-19.

Optimal designs are best for cases with high system knowledge, where previous studies have given confidence in the best type of model to be fitted, and the constraints of the system are well understood. See “Creating an Optimal Design” on page 2-25.

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. See “Adding Design Points” on page 2-35.

Design Editor Displays

The following example shows the display after creating an optimal design.

The screenshot shows the Design Editor interface with the following components:

- Design Tree:** Shows a hierarchy with 'Designs' containing 'Radial Basis Function Design', 'Actual Design', and 'AugmentedDesign'.
- Design Table:** A table with 12 rows and 5 columns. The first two columns are highlighted in red for rows 243 and 244. A context menu is open over the table, showing options like 'Current View', 'Viewer Options', 'Display Design Point Numbers', 'Split Vertically', 'Split Horizontally', 'Delete Current View', 'Print to Figure', and a list of projection types (1D, 2D, 3D, 4D, Pairwise, Constraints, Model Description).
- Properties - AugmentedDesign:**

Design Style	Custom
Number of Points	264
Number of Constraints	0
Last Changed	13-Feb-2003, ...
Model	RBF-multiqua...
- Status Bar:** Shows 'Ready'.

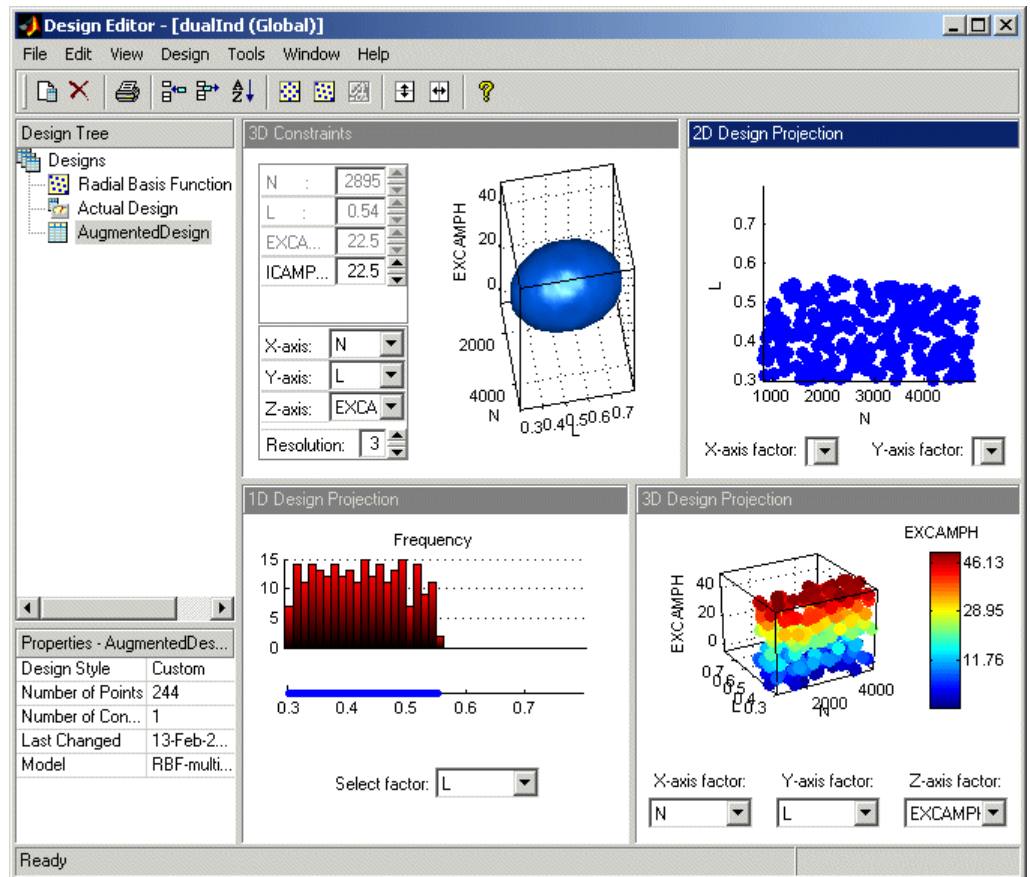
ID	Value 1	Value 2	Value 3	Value 4	Value 5
243	4452.906	0.503			
244	4587.575	0.528	9.88	31.593	
245	1219.158	0.574	6	-2.25	
246	4782.004	0.599	17	14.25	
247	3943.687	0.599	-2.25	41.75	
248	2057.475	0.474	6	-5	
249	2267.054	0.424	47.25	0.5	
250	4572.425	0.499	8.75	6	
251	2686.212	0.624	3.25	6	
252	4153.267	0.424	33.5	28	
253	4572.425	0.299	36.25	8.75	
254	2057.475	0.649	3.25	6	

When you first create or open a design, the main display area shows the default **Design Table** view of the design (see example above). All the views on the right show the design selected in the left tree (see “The Design Tree” on page 2-7). There is a context menu for the views on the right, available by right-clicking the title bars, in which you can change the view of the design to **1-D**, **2-D**, **3-D**, **4-D**, and **Pairwise Projections**, **2-D**, and **3-D Constraint** views, and the **Table** view (also under **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. The split can also be merged again. You can also use the toolbar buttons. 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 current active view. See “Design Editor Toolbar and Menus” on page 2-9 for more information about how to change your display options.

The information pane, bottom left, displays pieces of information for the current design selected in the tree. The amount of information in this pane can change depending on what the design is capable of; for example, only certain models can support the optimal designs and only these can show current optimal values. You can also see this information and more by selecting **File > Properties** or using the context menu in the tree.

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, 3-D, or pairwise plot of the current design points as the main plot.

The following example shows several views in use at once.







The Design Tree


The currently available designs are displayed on the left in a tree structure.

The tree displays three pieces of information:

- The name of the design, which you can edit by clicking it
- The state of the design

- The icon changes from  if it is empty, to the appropriate icon for the design type when it has design points (for example,  optimized, as in the toolbar buttons for Optimal, Classical, and Space-Filling designs).
- The icon changes to  when design points have been added using a different method (for example, augmenting a classical design with optimally chosen points). It becomes a *custom* design style. You can mix and match all design options in this way.
- A padlock appears () if the design is locked. This happens when it has child nodes (to maintain the relationship between designs, so you can retreat back up the design tree to reverse changes).
- The design that is selected as best. This is the design that is used for matching against experimental data. The icon for the selected design is the normal icon turned blue. When you have created more than one design, you must select as best the design to be used in modeling, using the **Edit** menu. Blue icons are also locked designs, and do not acquire padlocks when they have child nodes.
- You can reach a context menu by right-clicking in the design tree pane. Here you can delete or rename designs and add new designs. Choose **Evaluate Design** to open the Design Evaluation window. **Properties** opens the Design Properties dialog, which displays information about the size, constraints, properties (such as optimality values), and modification dates of the selected design.

Design Editor Toolbar and Menus



- **New Design, Delete Design, Print** — See the “File Menu” on page 2-9.
- **Add, Delete and Sort Point** — See the “Edit Menu” on page 2-11.
- **Classical, Space Filling and Optimal Design** — See the “Design Menu” on page 2-14.
- **Split View Horizontally and Vertically**— See the “View Menu” on page 2-12.

File Menu

- **New Design** — Creates a new design node in the tree. This is blank if it is the first design you create, or, if you have an existing design, the child node is a copy of the parent design. Also in the toolbar and the context menu in the design tree. Use the **Design** menu or toolbar buttons to set up designs.
- **Delete Design** — Deletes the currently selected design and its subdesigns. Also in the toolbar and the context menu in the design tree.

- **Rename Design** — Enables you to edit the name of the current design. You can also do this by clicking again on the selected design name, or right-click to use the context menu in the design tree.
- **Properties** — Opens the Design Properties dialog, where you can see information about your current design, such as the number of factors, points, and constraints; the design style; when the design was last modified; optimality values, space filling, and classical settings. Also in the context menu in the design tree.
- **Import Design** — Opens the Import Design dialog, where you can import designs from Design Editor files (*.mvd), comma-separated-values files (*.csv), or from the workspace. You can browse to the required file or specify the source variable if importing from the workspace. If it is not a Design Editor file you can choose to convert the design points from [-1, 1] range. See “Exporting and Importing Designs” on page 2-42.
- **Export Design** — Opens the Export Design dialog, where you can export designs to Design Editor files (*.mvd), comma-separated-values files (*.csv), or from the workspace. You can specify the name of the destination file or variable. For .csv files you can use the check box to choose whether to include factor names, and for the workspace and .csv files you can choose whether to convert design points to [-1, 1] range. See “Exporting and Importing Designs” on page 2-42. Note that you do not have to export your designs to save them — they are saved when you save your project in the Model Browser.
- **Merge Designs** — Opens the Merge Designs dialog, where you can choose which designs to merge and a base design. See “Merging Designs” on page 2-37.
- **Import Constraints** — Opens the Import Constraints dialog. Here you can import any suitable constraints for the currently selected design. You can import any existing constraints in the design tree, or from a design file, or you can import boundary constraints from file or the current project. See “Importing Constraints” on page 2-51.
- **Print** — Prints the current view (plots only). You can also use **Print Preview**. See also **Print to Figure** in the **View** menu. If you want to print information from the Table view you can copy the information to the clipboard by using **Edit > Copy Design Data**, or save the design as a .csv file. Also in the toolbar.

- **Close** — Closes the Design Editor. You return to the Model Browser window. Note that you do not lose your designs, you simply close the Design Editor. The designs reappear when you reopen the Design Editor. When you save your project in the Model Browser your designs remain part of that project.

Edit Menu

- **Copy View** — Copies the current view to the clipboard.
- **Copy Design Data** — Copies the design data to the clipboard. This can be useful if you want to print the contents of the Table view.
- **Clear** — Deletes all points in the current design.
- **Add Point** — Opens the Add Design Points dialog. Here you can choose how many points to add optimally, randomly, or at specified points. Also in the toolbar. See “Adding Design Points” on page 2-35.
- **Delete Point** — Opens the Delete Design Points dialog. Here you can choose the points to delete. Also in the toolbar. See “Fixing, Deleting, and Sorting Design Points” on page 2-39.
- **Sort Points** — Opens the Sort dialog. Here you can choose to sort by any or all of your factors, by custom expression, or at random. Also in the toolbar. See “Fixing, Deleting, and Sorting Design Points” on page 2-39.
- **Fix/Free Points** — Opens the Fix Design Points dialog. You can fix design points so they are not moved by design optimization processes. See “Fixing, Deleting, and Sorting Design Points” on page 2-39.
- **Randomize** — Select this option as a quick way of randomly resorting the points in the current design. This is a shortcut to the same functionality provided by the **Random** option in the Sort dialog.
- **Round Factor** — Opens the Round Design dialog. Here you can select a factor to round, and limit it to a specified number of significant figures or to a fixed interval.
- **Constraints** — Opens the Constraints Manager dialog. Here you can add, edit, duplicate, and delete constraints on your designs. See “Applying Constraints” on page 2-43.

- **Model** — Opens the Global or Local Model Setup dialog (depending on which stage you are designing an experiment for). Here you can change the model for which you are designing an experiment. The model type is important for optimal designs.
- **Select As Best** — Selects the current design as best. This changes the icon in the tree blue. This is the design that is used for matching against experimental data.

View Menu

- **Current View** — Changes the current view to your selection from the submenu:
 - Design Table
 - 1-D Design Projection
 - 2-D Design Projection
 - 3-D Design Projection
 - 4-D Design Projection
 - Pairwise Design Projections
 - 2-D Constraints
 - 3-D Constraints
 - Model Description
- **View Options** — these items depend on the currently selected view:
 - **Plot Properties**— For 1-D, 2-D and 3-D Design Projections. Opens dialogs for configuring details of the current display. You can change basic properties such as color on the projections (1-D, 2-D, 3-D, and 4-D). You can rotate all 3-D views as usual.
 - **Edit Colormap** For the 3-D and 4-D Design Projections. You can also double-click the color bar to edit the colormap.
 - **Graph Size** For the Pairwise Projections, you can choose graph size or to display all graphs.
 - **Value Filter** — For the table view, you can set up a filter to selectively display certain ranges of values.

- **Display Design Point Numbers** — You can select this option to toggle the display of design point numbers in views that support the feature. A design point number is the index of a particular point in the design: this value is permanently displayed in the table view. Views that support the display of design point numbers are
 - 2-D Design Projection
 - 3-D Design Projection
 - 4-D Design Projection
 - Pairwise Design Projections

Because all these views are projections that use a subset of the design's input factors, it is often the case that the resulting view contains points that have been plotted on top of each other. In this case, the design point numbers will stack up in a column above the common point to aid readability. You can use **Display Design Point Count** to see at a glance how many points are overlapping in any stack. You can select point count or point numbers but not both.

Note Displaying multiple views with design point numbers for large designs can significantly slow down the display. You might want to turn off the design point number display in these cases.

- **Display Design Point Count** — If points are plotted on top of each other (in 2-D, 3-D, 4-D, or pairwise plots) this option allows you to see how many points are overlapping in each cluster. A number next to a point indicates that more than one point is plotted there.
- **Split View** — splits the current view and adds your selected new view from the submenu.
- **Split View Vertically** — Splits the current view vertically to produce a new view in addition to the currently selected view. Also in the toolbar and the buttons in the title bar of each view. New plot types are produced for each new view. This is a quick way to produce a variety of different plots. Remember that you can change any existing view type by selecting **View > Current View** or, alternatively, using the context menu to select from the **Current View** submenu.

- **Split View Horizontally** — Splits the current view horizontally to produce new views, as for **Split Vertically**. Also in the toolbar and the buttons in the title bar of each view.
- **Close View** — Deletes the current view.
- **Print to Figure** — This option copies the current view into its own figure, allowing you to use the standard MATLAB plotting tools to annotate and print the display.

Design Menu

- **Optimal** — Opens the Optimal Design dialog. Also in the toolbar. See “Creating an Optimal Design” on page 2-25.
- **Classical** — You can use the submenu here to go directly to the type of classical design you want, or select **Design Browser** to see all the options. See “Creating a Classical Design” on page 2-16. The toolbar button opens the Design Browser.
- **Space Filling** — You can use this submenu to go directly to the type of space filling design you want, or select **Design Browser** to see all the options. See “Creating a Space-Filling Design” on page 2-19. The toolbar button opens the Design Browser.

Tools Menu

- **PEV Viewer** — Opens the Prediction Error Variance Viewer where you can evaluate the predictive power of your designs. See “Prediction Error Variance Viewer” on page 2-54.
- **Evaluate Designs** — Opens the Design Evaluation window where you can examine detailed mathematical properties of your design. Also in the context menu in the design tree. See “Design Evaluation Tool” on page 2-62.



Window Menu

This allows you to switch between the Model Browser and Design Editor windows.

Help Menu

As with everywhere in the toolbox, the **Help** menu provides access to general toolbox help and help specific to the current view. Here you can select **MBC Help** to browse all the toolbox help, **Design Editor Help** to go straight to the Designs documentation, or **About MBC** to see the current version number.

Creating a Classical Design


- 1** Add a new design by clicking the  button in the toolbar or select **File > New**.
- 2** Select the new design node in the tree. An empty Design Table appears if you have not yet chosen a design. Otherwise if this is a new child node the display remains the same, because child nodes inherit all the parent design's properties. 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.
- 3** 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.

- 4** A dialog appears if 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. The default is replacement of the current points with a new design. Click **OK** to proceed, or **Cancel** to change your mind.

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. The options are **Face-center cube**, **Spherical**, **Rotatable**, or **Custom**. If you choose **Custom**, you can then choose a ratio value (α) between the corner points and the face points for each factor and the number of center points to add. Five levels for each factor are used. You can set the ranges for each factor. **Inscribe star points** scales all points within the coded values of 1 and -1 (instead of plus or minus α outside that range). When this box is not selected, the points are circumscribed.

- 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 you should choose 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. You can set the ranges of each factor.

- Full Factorial

Generates an n -dimensional grid of points. You can choose the number of levels for each factor, the number of additional center points to add, and the ranges for each factor.

- Plackett Burman

These are “screening” designs. They are two-level designs that are designed to allow you to work out which factors are contributing any effect to the model while using the minimum number of runs. For example, for a 30-factor problem this can be done with 32 runs. They are constructed from Hadamard matrices and are a class of two-level orthogonal array.



- Regular Simplex

These designs are generated by taking the vertices of a k -dimensional regular simplex (k = number of factors). For two factors a simplex is a triangle; for three it is a tetrahedron. Above that are hyperdimensional simplices. These are economical first-order designs that are a possible alternative to Plackett Burman or full factorials.

You can always toggle coded values by selecting the check box at the top.


Creating a Space-Filling Design

Space-filling designs should be used 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. These designs literally fill out the n -dimensional space with points that are in some way regularly spaced. These designs can be especially useful in conjunction with nonparametric models such as radial basis function (a type of neural network).

- 1 Add a new design by clicking the  button in the toolbar or select **File > New**.
- 2 Select the node in the tree by clicking. An empty Design Table appears if you have not yet chosen a design. Otherwise, if this is a new child node the display remains the same, because child nodes inherit all the parent design's properties.
- 3 Select **Design > Space Filling > Design Browser**, or click the **Space Filling Design** button  on the toolbar.
- 4 A dialog appears if 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. The default is replacement of the current points with a new design. Click **OK** to proceed, or **Cancel** to change your mind.

The Space Filling Design Browser appears.

Note As with the Classical Design Browser, you can select the three types of design you can preview in the Space Filling Design Browser from the **Design > Space Filling** menu in situations when you already know the type of space-filling design you want.



Space-Filling Design Styles

In the **Design** drop-down menu you can see the three design styles available:

- “Latin Hypercube Sampling” on page 2-21
- “Lattice” on page 2-21
- “Stratified Latin Hypercube” on page 2-22

Latin Hypercube Sampling

Latin Hypercube Sampling (LHS) designs are sets of design points that, for an N point design, project onto N different levels in each factor. Here the points are generated randomly. You choose a particular Latin Hypercube by trying several such sets of randomly generated points and choosing the one that best satisfies user-specified criteria.

Lattice


Lattice designs project onto N different levels per factor for N points. The points are not randomly generated but are produced by an algorithm that uses a prime number per factor. If good prime numbers are chosen, the lattice spreads points evenly throughout the design volume. A poor choice of prime numbers results in highly visible lines or planes in the design projections. If all the design points are clustered into one or two planes, it is likely that you cannot estimate all the effects in a more complex model. When design points are projected onto any axes, there are a large number of factor levels.

For a small number of trials (relative to the number of factors) LHS designs are preferred to Lattice designs. This is because of the way Lattice designs are generated. Lattice designs use prime numbers to generate each successive sampling for each factor in a different place. No two factors can have the same generator, because in such cases the lattice points all fall on the main diagonal of that particular pairwise projection, creating the visible lines or planes described above. When the number of points is small relative to the number of factors, the choice of generators is restricted and this can lead to Lattice designs with poor projection properties in some pairwise dimensions, in which the points lie on diagonals or double or triple diagonals. This means that Latin Hypercube designs are a better choice for these cases.

See the illustrations in the following section comparing the properties of good and poor lattices and a hypercube design.


Stratified Latin Hypercube

Stratified Latin Hypercubes separate the normal hypercube into N different levels on user-specified factors. This can be useful for situations where the preferred number of levels for certain factors might be known; more detail might be required to model the behavior of some factors than others. They can also be useful when certain factors can only be run at given levels.



The preceding example shows the different properties of a poor lattice (left) and a good lattice (right), with a similar number of points. The poorly chosen prime number produces highly visible planes and does not cover the space well.

An example of an LHS design of the same size is shown for comparison.



Setting Up a Space-Filling Design

You can toggle coded units for all space-filling designs by selecting the check box at the top.

The default **Design** drop-down menu item is Latin Hypercube Sampling.

For both Latin Hypercube Sampling and Stratified Latin Hypercube, you can choose from several selection criteria available in a drop-down menu:

- Maximize minimum distance (between points). This is the default.
- Minimize maximum distance (between points)
- Minimize discrepancy — Minimizes the deviation from the average point density
- Minimize RMS variation from CDF — Minimizes the Root Mean Square (RMS) variation of the Cumulative Distribution Function (CDF) from the ideal CDF

- Minimize maximum variation from CDF — Minimizes the maximum variation of the CDF from the ideal CDF

The final two (CDF variation) options scale best with the number of points and it is advisable to choose one of these options for large designs.

The same criteria are available for the Stratified Latin Hypercube.



- You can set the number of points using the controls or by typing in the edit box.
- You can set the ranges for each factor.
- If you select the **Enforce Symmetrical Points** check box, you create a design in which every design point has a *mirror* design point on the opposite side of the center of the design volume and an equal distance away. Restricting the design in this way tends to produce better Latin Hypercubes.
- You can use the tabs under the display to view 2-D, 3-D, and 4-D previews.
- Click **OK** to calculate the Latin Hypercube and return to the main Design Editor.

For a Lattice space-filling design:

- You can choose the lattice size by using the buttons or typing in the edit box.
- You can choose the prime number generator by using the up/down buttons on the **Prime number for X** edit box.
- You can choose the range for each factor.
- Click **OK** to calculate the lattice and return to the Design Editor.

Creating an Optimal Design

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. Optimal designs require linear models.

- 1 Click the  button in the toolbar or select **File > New Design**. A new node appears in the design tree. It is named according to the model for which you are designing, for example, Linear Model Design.
- 2 Select the node in the tree by clicking. An empty Design Table appears if you have not yet chosen a design. Otherwise, if this is a new child node the display remains the same, because child nodes inherit all the parent design's properties.
- 3 Set up any constraints at this point. See “Applying Constraints” on page 2-43.
- 4 Choose an Optimal design by clicking the  button in the toolbar, or choose **Design > Optimal**.

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.
- p additional points are added to the design, either optimally or at random. These points are chosen from the candidate set.
- p 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 three tabs that contain the settings for three main aspects of the design:

- **Initial Design** tab: Starting point and number of points in the design
- Candidate Set tab: Candidate set of points from which the design points are chosen
- Algorithm tab: Options for the algorithm that is used to generate the points

Optimal Design: Initial Design Tab

The screenshot shows the 'Optimal Design' dialog box with the 'Initial Design' tab selected. At the top, 'Optimality criteria:' is set to 'V-Optimal'. Below this are three tabs: 'Initial Design', 'Candidate Set', and 'Algorithm'. The 'Initial Design' tab contains two sections. The first section, 'Initial Points From Current Design', has three radio button options: 'Keep current design points', 'Keep current fixed design points', and 'Do not keep any current design points', with the third option selected. The second section, 'Additional Design Points', contains four rows of settings: 'Points from current design:' set to 0, 'Minimum additional points required:' set to 10, 'Optional additional points:' set to 20, and 'Total number of points:' set to 30. At the bottom of the dialog are three buttons: 'OK', 'Cancel', and 'Help'.

The **Initial Design** tab allows you to define the composition of the initial design: how many points to keep from the current design and how many total or additional points to choose from the candidate set.

1 Choose the type of the optimal design, using the **Optimality criteria** drop-down menu:

- D-Optimal designs — Aims to reduce the volume of the confidence ellipsoid to obtain accurate coefficients. This is set up as a maximization problem, so the progress graph should go up with time.

The D-optimality value used is calculated using the formula

$$D_{\text{eff}} = \frac{\log_e(\det(\mathbf{X}\mathbf{X}))}{k}$$

where \mathbf{X} is the regression matrix and k is the number of terms in the regression matrix.

- V-Optimal designs — Minimizes the average prediction error variance, to obtain accurate predictions. This is better for calibration modeling problems. This is a minimization process, so the progress graph should go down with time.

The V-optimality value is calculated using the formula

$$V_{\text{eff}} = \frac{1}{n_c} \sum_j x_j' (\mathbf{X}_c' \mathbf{X}_c)^{-1} x_j$$

where x_j are rows in the regression matrix, \mathbf{X}_c is the regression matrix for all candidate set points, and n_c is the number of candidate set points.

- A-Optimal designs — Minimizes the average variance of the parameters and reduces the asphericity of the confidence ellipsoid. The progress graph also goes down with this style of optimal design.

The A-optimality value is calculated using the formula

$$A_{\text{eff}} = \text{trace}((\mathbf{X}\mathbf{X})^{-1})$$

where \mathbf{X} is the regression matrix.

2 You might already have points in the design (if the new design node is a child node, it inherits all the properties of the parent design). If so, choose from the radio buttons:

- **Replace the current points with a new initial design**
- **Augment the current design with additional points**
- **Keep only the fixed points from the current design**

For information on fixed design points, see “Fixing, Deleting, and Sorting Design Points” on page 2-39.

- 3** You can choose the total number of points and/or the number of additional points to add by clicking the **up/down** buttons or by typing directly into the edit boxes for **Optional additional points** or **Total number of points**.

Optimal Design: Candidate Set Tab

The **Candidate Set** tab allows you to set up a candidate set of points for your optimal design. Candidate sets are a set of potential test points. This typically ranges from a few hundred points to several hundred thousand.

Select variables in this list

Choose algorithm type from this list

The screenshot displays the 'Optimal Design' software window with the 'Candidate Set' tab selected. On the left, there are two projection windows: '3D Projection of Candidate Set' and '2D Projection of Candidate Set'. The 3D window shows a grid of points colored by their 'A' value, with a color scale from 14.84 (blue) to 92.97 (red). Below it are dropdown menus for 'x-axis factor' (N), 'y-axis factor' (L), and 'z-axis factor' (A). The 2D window shows a grid of points with 'N' on the x-axis and 'L' on the y-axis. Below it are dropdown menus for 'x-axis factor' (N) and 'y-axis factor' (L). The main 'Optimal Design' window has 'Optimality Criteria' set to 'V-Optimal' and 'Generation algorithm' set to 'Grid'. Under 'Options', 'Equally spaced levels' is selected, and 'Number of levels for A' is set to 16. The 'Display' section shows icons for different views and a 'Display a maximum of 2500 points' option. At the bottom, there are 'OK', 'Cancel', and 'Help' buttons. Annotations with arrows point to various elements: 'Select variables in this list' points to the variable list in the 3D window; 'Choose algorithm type from this list' points to the 'Generation algorithm' dropdown; 'Open display windows with these buttons' points to the view icons in the 'Display' section; and 'Change the number of levels of the selected variable here' points to the 'Number of levels for A' spinner.

Optimality Criteria: V-Optimal

Start Point Candidate Set Algorithm

Generation algorithm: Grid

Allow replicated points in design

Options

View coded values

N
L
A

Equally spaced levels

Minimum A value: 0

Maximum A value: 100

Number of levels for A: 16

User-specified levels

0:6.6667:100

Display

Display a maximum of 2500 points.

0 constraints will be applied to this candidate set.

OK Cancel Help

Open display windows with these buttons

Change the number of levels of the selected variable here

The set generation schemes are as follows:

- **Grid** — Full factorial grids of points, with fully customizable levels.
- **Lattice** — These have the same definition as the space-filling design lattices, but are typically used with about 10,000 points. The advantage of a lattice is that the number of points does not increase as the number of factors increases; however, you do have to try different prime number generators to achieve a good lattice. See “Creating a Space-Filling Design” on page 2-19.
- **Grid/Lattice** — A hybrid set where the main factors are used to generate a lattice, which is then replicated over a small number of levels for the remaining factors.
- **Stratified Lattice** — Another method of using a lattice when some factors cannot be set to arbitrary values. Stratified lattices ensure that the required number of levels is present for the chosen factor. Note that you cannot set more than one factor to stratify to the same N levels. This is because forcing the same number of levels would also force the factors to have the same generator. As for a lattice space-filling design, no two factors can have the same generator, because in such cases the lattice points all fall on the main diagonal of that particular pairwise projection, creating highly visible planes in the points and poor coverage of the space. For illustrations of this effect, see “Lattice” on page 2-21.
- **User-defined** — Import custom matrices of points from MATLAB or MAT-files.

For each factor you can define the range and number of different levels within that range to select points.

- 1** Choose a type of generation algorithm from the drop-down menu. Note that you could choose different parameters for different factors (within an overall scheme such as **Grid**).
- 2** This tab also has buttons for creating plots of the candidate sets. Try them to preview your candidate set settings. If you have created a custom candidate set you can check it here. The edit box sets the maximum number of points that will be plotted in the preview windows. Candidate sets with many factors can quickly become very large, and attempting to display the entire set will take too long. If the candidate set has more

points than you set as a maximum, only every Nth point is displayed, where N is chosen such that (a) the total displayed is less than the maximum and (b) N is prime. If you think that the candidate set preview is not displaying an adequate representation of your settings, try increasing the maximum number of points displayed.

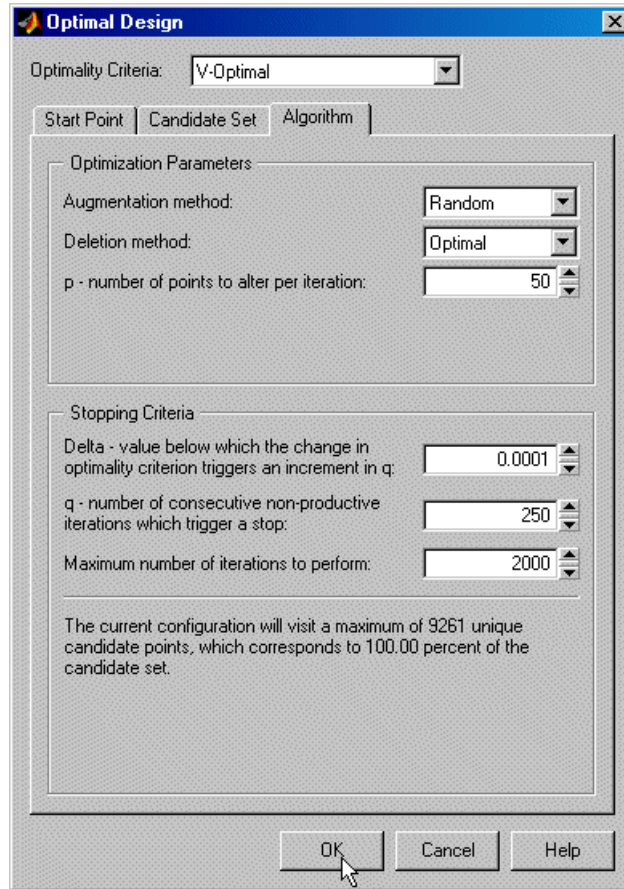
- 3 Notice that you can see 1-D, 2-D, 3-D, and 4-D displays (fourth factor is color) all at the same time as they appear in separate windows (see the example following). Move the display windows (click and drag the title bars) so you can see them while changing the number of levels for the different factors.
- 4 You can change the factor ranges and the number of levels using the edit boxes or buttons.

Optimal Design: Algorithm tab

The **Algorithm** tab has the following algorithm details:

- **Augmentation method** — Random or Optimal— Optimal can be very slow (searches the entire candidate set for points) but converges using fewer iterations. Random is much faster per iteration, but requires a larger number of iterations. The Random setting does also have the ability to lower the optimal criteria further when the Optimal setting has found a local minimum.
- **Deletion method** — Random or Optimal— Optimal deletion is much faster than augmentation, because only the design points are searched.
- **p — number of points to alter per iteration** — The number of points added/removed per iteration. For optimal augmentation this is best kept smaller (~5); for optimal deletion only it is best to set it larger.
- **Delta — value below which the change in optimality criteria triggers an increment in q** — This is the size of change below which changes in the optimality criteria are considered to be not significant.
- **q — number of consecutive non-productive iterations which trigger a stop** — Number of consecutive iterations to allow that do not increase the optimality of the design. This only has an effect if random augmentation or deletion is chosen.


- **Maximum number of iterations to perform** — Overall maximum number of iterations.



- 1** Choose the augmentation and deletion methods from the drop-down menus (or leave at the defaults).
- 2** You can alter the other parameters by using the buttons or typing directly in the edit boxes.
- 3** Click **OK** to start optimizing the design.

When you click the **OK** button on the Optimal Design dialog, another window appears that contains a graph. This window 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.

- 4 You can click **Accept** at any time, but it is most useful to wait until iterations are not producing noticeable improvements; that is, the graph becomes very flat.



You can always return to the Optimal Design dialog (following the same steps) and choose to keep the current points while adding more.

Averaging Optimality Across Multiple Models


The Design Editor can average optimality across several linear models. This is a flexible way to design experiments using optimal designs. If you have no idea what model you are going to fit, you would choose a space-filling design. However, if you have some idea what to expect, but are not sure which model to use, you can specify a number of possible models. The Design Editor can average an optimal design across each model.

For example, if you expect a quadratic and cubic for three factors but are unsure about a third, you can specify several alternative polynomials. You

can change the weighting of each model as you want (for example, 0.5 each for two models you think equally likely). This weighting is then taken into account in the optimization process in the Design Editor. See “Global Model Class: Multiple Linear Models” on page 4-21.

Adding Design Points

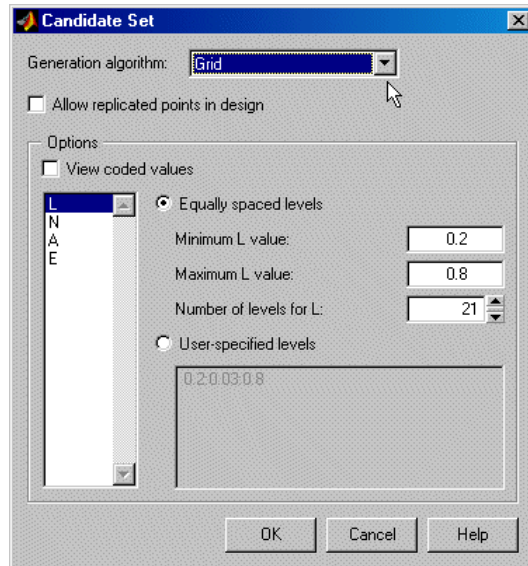
In any design, you can add points using the **Edit** menu. You can specify how many points to add and how to do so: optimally, randomly, or at specified values.

- 1 Select **Edit > Add Point** or click the  button. A dialog appears, as shown following.
- 2 Choose an augmentation method from the drop-down menu: optimal (D,V, or A), random, or user-specified.

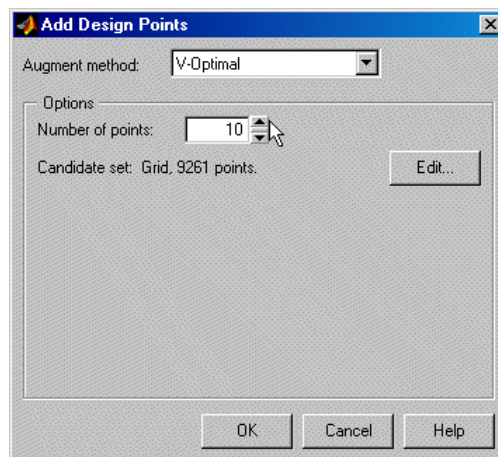
Note You can add points optimally to any design based on a linear or multilinear model, as long as it has the minimum number of points required to fit that model. This means that after adding a constraint you might remove so many points that a subsequent replace operation does not allow optimal addition.

- 3 Choose the number of points to add, using the buttons or typing into the edit box. For user-specified custom points, you also enter the values of each factor for each point you want to add.
- 4 If you choose an optimal augmentation method and click **Edit**, the Candidate Set dialog appears, as shown in the following example. Here you can edit the ranges and levels of each factor and which generation

algorithm to use. These are the same controls you see on the Candidate Set tab of the Optimal Design dialog.




5 Click **OK** to add the points and return to the main display.



Merging Designs

You can merge the points from two or more designs together using the **File** menu. You can merge designs together to form a new design, or merge points into one of the chosen designs. Points that are merged retain their fixed status in the new design.

- 1 Select **File > Merge Designs**. A dialog appears, as shown.



- 2 A list of all the designs from the Design Editor is shown, along with the associated design style and number of points. Select two or more designs from the list by dragging with the mouse or **Ctrl**+clicking.
- 3 When at least two designs have been selected, the options at the bottom of the dialog are enabled. Choose whether you want to create a new design or put the design points into an existing design.
- 4 If you choose to create a new design, you must also choose one of the selected designs to act as a base. Properties such as the model, constraints, and any optimal design setup are copied from this base design. If you

choose to reuse an existing design you must choose one of the selected designs to receive the points from other designs.


- 5 Click **OK** to perform the merging process and return to the main display. If you choose to create a new design, it appears at the end of the design tree.

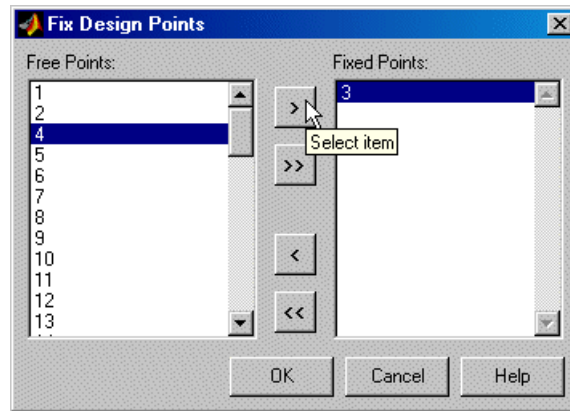
Fixing, Deleting, and Sorting Design Points

You can fix or delete points using the **Edit** menu. You can also sort points or clear all points from the design.

Fixed points become red in the main Design Editor table display. If you have matched data to a design or used experimental data as design points, those points are automatically fixed. You already have the data points, so you do not want them to be changed or deleted. Design points that have been matched to collected data are also fixed. Since these points have already been run they cannot be freed — you will not see them in the Fix Design Points dialog. Once you have fixed points, they are not moved by design optimization processes. This automatic fixing makes it easy for you to optimally augment the fixed design points.

To fix or delete points:

- 1** To delete all points in the current design, select **Edit > Clear**.
- 2** If you want to fix or delete particular points, first change the main Design Editor display pane to the Table view. This gives a numbered list of every point in the design, so you can see where points are in the design.
- 3** Select **Edit > Fix/Free Points** or **Edit > Delete Point** (or click the toolbar button ).
- 4** A dialog appears in which you can choose the points to fix or delete.




The example above shows the dialog for fixing or freeing points. The dialog for deleting points has the same controls for moving points between the **Keep Points** list and the **Delete Points** list.

- 5 Move points from the **Free Points** list to the **Fixed Points** list; or from the **Keep Points** list to the **Delete Points** list, by using the buttons.
- 6 Click **OK** to complete the changes specified in the list boxes, or click **Cancel** to return to the unchanged design.

To sort points:

- 1 Select **Edit > Sort**. A dialog appears (see example following) for sorting the current design — by ascending or descending factor values, randomly, or by a custom expression.



- 2** To sort by custom expression you can use MATLAB expressions (such as $\text{abs}(N)$ for the absolute value of N) using the input symbols. Note that sorts are made using coded units (from -1 to 1) so remember that points in the center of your design space will be near zero in coded units, and those near the edge will tend to 1 or -1.
- 3** Select **Edit > Randomize** as a quick way of randomly resorting the points in the current design. This is a shortcut to the same functionality provided by the **Random** option in the Sort dialog.

Exporting and Importing Designs

Note that you do not need to save your designs separately from the project. When you select **File > Close** you do not lose your designs, you simply close the Design Editor. When you save your project in the Model Browser, your designs remain part of that project. You can also export designs to a file or the workspace.

You can export any design by choosing **File > Export Design**. The selected design *only* is exported.

There are three options:

- Design Editor file (*.mvd) generates a Design Editor file (.mvd).
- Comma separated format file (*.csv) exports the matrix of design points to a CSV (comma-separated-values) file. You can include factor symbols by selecting the check boxes.
- Workspace exports the design matrix to the workspace. You can convert design points to a range of (1, -1) by selecting the check box.

You can choose the destination file or variable by typing in the edit box or using the browse button.

Import designs by selecting **File > Import**. The controls on the dialog are very similar to the Export Design dialog: you can import from Design Editor files, CSV files, or the workspace, and you can convert design points from a (1,-1) range.


Applying Constraints

In many cases designs might not coincide with the operating region of the system to be tested. For example, an automobile engine normally does not operate in a region of low speed (n) and high exhaust gas recirculation (EGR). You cannot run 15% EGR at 1000 RPM. There is no point selecting design points in impractical regions, so you can constrain the candidate set for test point generation.

Designs can have any number of geometric constraints placed upon them. Each constraint can be one of four types: an ellipsoid, a hyperplane, a 1-D lookup table, or a 2-D lookup table.

To add a constraint to a design:

- 1 Select **Edit > Constraints** from the Design Editor menus.
- 2 The Constraints Manager dialog appears.



The example shows the Constraints Manager dialog with two constraints. Here you can add new constraints, and delete, edit, or duplicate existing constraints.

If there are no constraints yet, the Constraints Manager is empty and you can only click **Add** to construct a new constraint.

- 1 Click **Add**.
- 2 The Constraint Editor dialog with available constraints appears. You can select the following from the **Constraint Type** drop-down menu: Linear, Ellipsoid, 1-D Table and 2-Table. See the next section, “Constraint Types” on page 2-44.


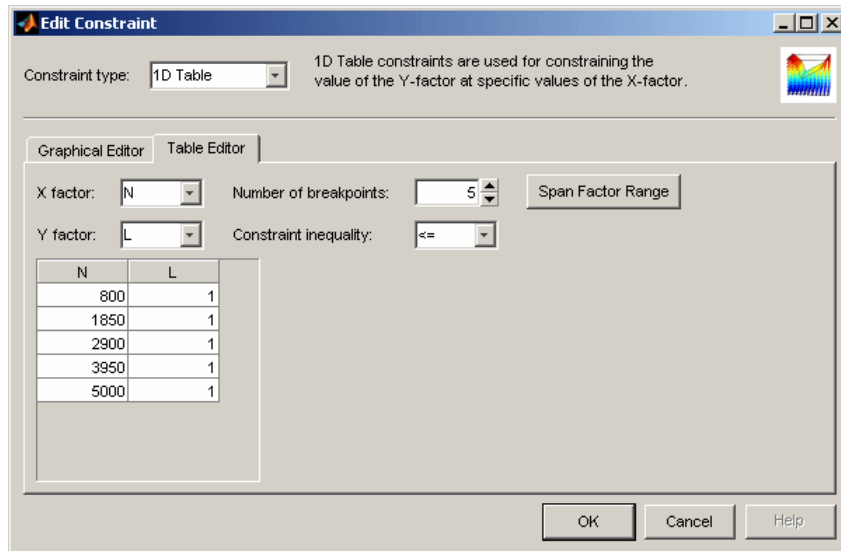
Constraint Types

- “1-D Table Constraints” on page 2-44
- “2-D Table Constraints” on page 2-49
- “Linear Constraints” on page 2-46
- “Ellipsoid Constraints” on page 2-47

Note The Design Editor constraint types are the same as the Linear, Ellipsoid, 1-D Table and 2-D Table constraint types you can apply to optimizations in the CAGE Browser part of the Model-Based Calibration Toolbox.

1-D Table Constraints

1-D table constraints limit the maximum or minimum setting of one factor as a function of another factor. Linear interpolation between user-defined points is used to specify the constraint. You can use the Table Editor tab or the Graphical Editor tab to define the constraint.



- On either tab, select the appropriate factors to use for X and Y, and choose whether to constrain your design above or below the defined boundary using the **Inequality** drop-down menu.

- On the **Table Editor** tab you can enter the number of breakpoints, click the button to **Span Factor Range** to space your breakpoints evenly, and you can enter breakpoint values into the table.
- On the **Graphical Editor** tab you can add and remove points, and click and drag the points to define the boundary. You can also enter values in the edit boxes for the selected point.

Linear Constraints

You specify the coefficients of the equation for an (N-1) dimensional hyperplane in the N-factor space. The form of the equation is $A \cdot x = b$ where A is your defined coefficient vector, x is the vector of factor settings, and b is a scalar. The equation is applied by substituting design point settings (in coded values) for x . For example,

In two dimensions: $A=(1, 2)$, $x=(L, A)$, $b=3$

Then $A \cdot x = b$ expands to

$$1 \cdot L + 2 \cdot A = 3$$

Rearranging this, you can write it as

$$A = -L/2 + 3/2$$

which corresponds to the traditional equation of a 2-D straight line, $y = mx + c$, with $m = -1/2$ and $c = 3/2$. $A \cdot x = b$ is thus the higher dimensional extension of this equation.

The linear constraints work by selecting the region below the defined plane (that is, $A \cdot x \leq b$). To select a region above the plane, multiply all your values by -1: $A \rightarrow -A$, $b \rightarrow -b$.

For example, to select a simple plane where $N < 0.8$ as a constraint boundary, enter 8 under N and set all the other factors to 0.

Ellipsoid Constraints

The ellipsoid constraint allows you to define an N-dimensional ellipsoid. You can specify the center of the ellipsoid, the length of each axis, and the rotation of the ellipsoid.

Ellipsoid center. You can specify the center of the ellipsoid by entering values in the **Center point** columns. These are the values, in natural units, that mark where you want the ellipsoid to be centered in each of the factor dimensions. The defaults are the midpoint of each factor range.


Axis length. You specify the size of the ellipsoid by entering values along the diagonal of the matrix. The default values create an ellipsoid that touches the edge of the design space in each of the factor dimensions. In general, for an entry value X in the diagonal, the ellipsoid size in that factor is $1/\sqrt{X}$.

If you want a radius of r in a factor, enter $1/(r^2)$. For example, if you want to restrict N to a radius of 2000 from the center point, enter $1/2000^2 = 2.5e-7$, as shown in the example below.

Enter a zero in the diagonal to not constrain with respect to that factor.

Rotation. The matrix entries that are not on the main diagonal control rotation of the ellipsoid.


The following example shows a defined ellipsoid constraint.



You must enter values in the table to define the ellipsoid. If you leave the values at the defaults, the candidate set is a sphere.

In this example, entering 2.5e-7 in the ENGSPEED diagonal restricts that axis to $1/\sqrt{2.5e-7} = 2000$. Entering zero in the INTCAM diagonal leaves INTCAM unconstrained (that is, the constraint is a cylinder extending to the ends of the INTCAM factor range). The ellipse is not rotated as the non-diagonal matrix entries are all zero.

A 3-D display of this constraint in the Design Editor can be seen below.




2-D Table Constraints


2-D table constraints are an extension of the 1-D table. Constraint boundary values for a factor are specified over a 2-D grid of two other factors.

- You can specify these grid locations by entering values in the top row and left column, while the matrix of values for the third factor is entered in the rest of the edit boxes. To specify grid values, you can enter values directly or just choose the number of breakpoints for your grid and space them over the factors' ranges, using the controls described below.
- You can specify the number of breakpoints for the X and Y factors.
- You can click **Span Range** to space your breakpoints evenly over the range of X or Y. This is useful if you add some breakpoints, as new points are often all at the maximum value for that factor. It is much quicker to use the **Span Range** button than to change points manually.
- You can specify to keep the region below (\leq) or above (\geq) the constraint boundary, as for the 1-D table, using the **Inequality** drop-down menu for the Z factor.
- You can switch to coded values using the check box. See the example.

The constraint boundary between the defined grid points is calculated using bilinear interpolation.



- After defining any constraint, click **OK**. Your new constraint appears in the Constraint Manager list box. Click **OK** to return to the Design Editor, or **Add** to define more constraints.

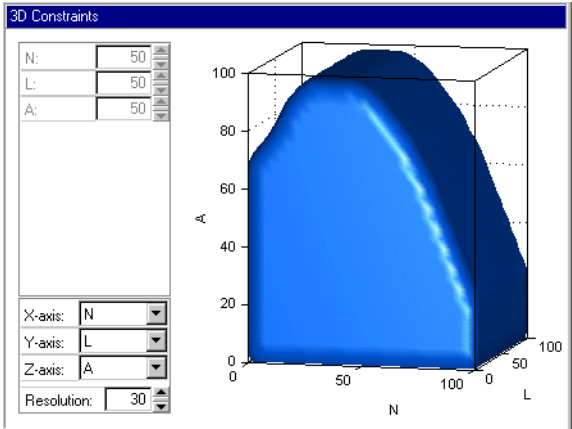


A dialog appears if there are points in the design that fall outside your newly constrained candidate set. You can simply continue (delete them) or cancel the constraint. Fixed points are *not* removed by this process. For optimal designs you can also replace them with new random points within the new candidate set, as shown in the preceding example dialog.

Note You only get the **Replace** points option for optimal designs. If you want to replace points removed by constraints from other designs, you can always use **Edit > Add Point** to add points optimally, randomly, or at chosen places. However, if so many points have been removed by a constraint that there are not enough left to fit the current model, optimal addition is not possible. See “Adding Design Points” on page 2-35.

To view constraints in the Design Editor,


- 1 Right-click the Design Editor display pane to reach the context menu.
- 2 Select **Current > View > 3D Constraints**. (You can also select 2-D constraints view). An example is shown.



These views are intended to give some idea of the region of space that is currently available within the constraint boundaries.

Importing Constraints

Select **File > Import Constraints**. The Import Constraints dialog appears, as shown in the following example.



Here you can import constraints for the currently selected design from

- Current designs — import any existing constraints in the design tree
- Design editor file (*.mvd) — extract constraints from a design file
- Boundary constraints (current project) — import boundary constraints from the Model Browser project
- Boundary constraints (*.mat file) — extract boundary constraints from file

Note You can only import design constraints from designs that have the same number of factors and have the same coded range for each factor. For designs of N factors you can import boundary constraints with N or less active factors.

Select constraints in the list by clicking, or **Ctrl**+click to select multiple constraints.

If importing from a file you can type the filename in the edit box or use the **browse** button to locate the file.


Click **OK** to import. If importing boundary constraints a dialog appears (once for each constraint) where you can match up factor names.

Prediction Error Variance Viewer

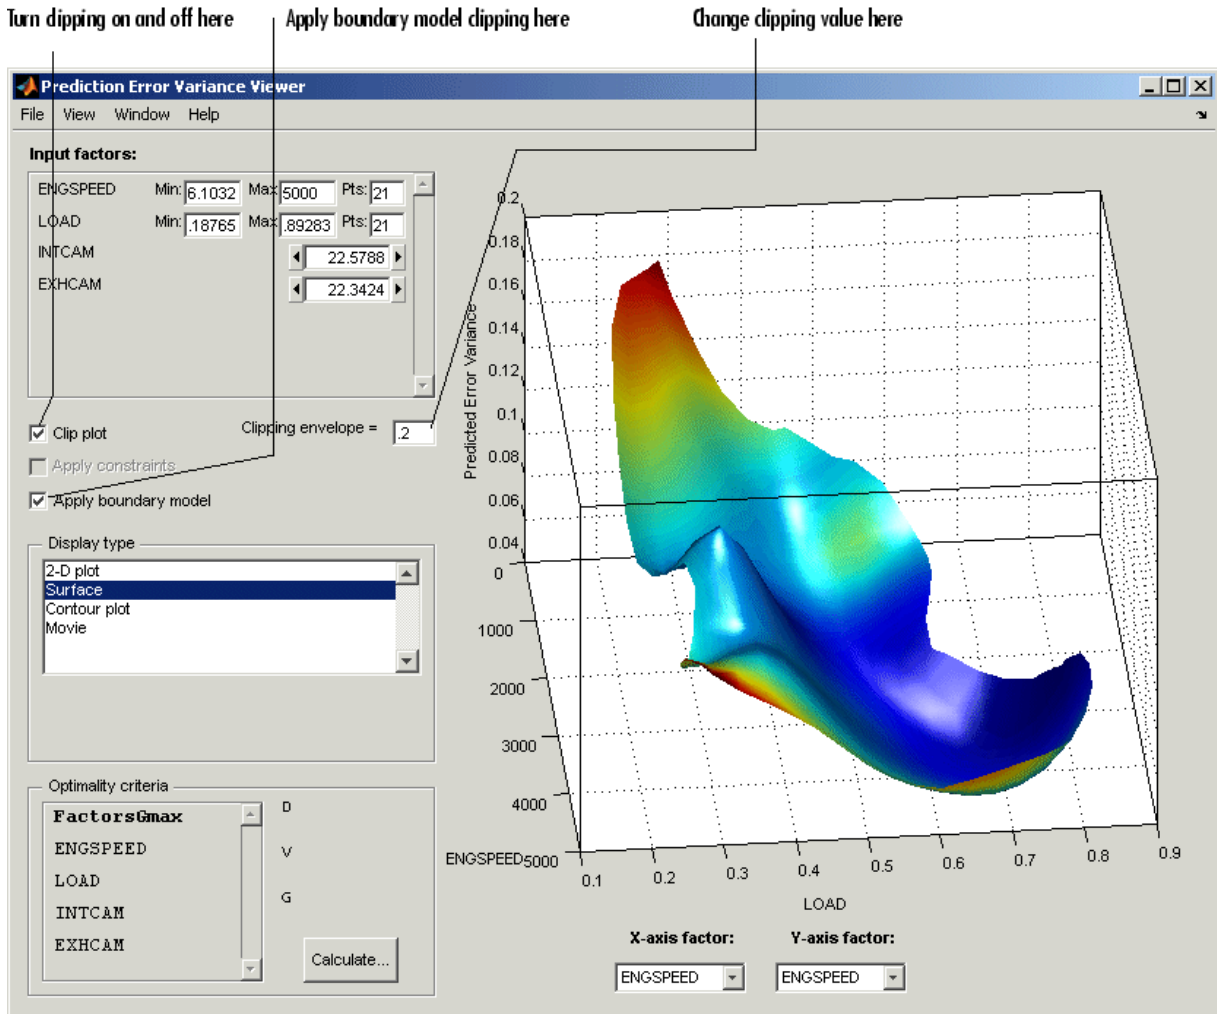
You can use the Prediction Error Variance (PEV) viewer to examine the quality of the model predictions. You can examine the properties of designs or global models. When you open it from the Design Editor, you can see how well the underlying model predicts over the design region. When you open it from a global model, you can view how well the current global model predicts. A low PEV (tending to zero) means that good predictions are obtained at that point.

The PEV viewer is only available for linear models and radial basis functions.

When designs are rank deficient, the PEV Viewer appears but is empty; that is, the PEV values cannot be evaluated because there are not enough points to fit the model.

- From the Design Editor, select **Tools > PEV Viewer**.
- From the global level of the Model Browser, if the selected global model is linear or a radial basis function,
 - Click the  toolbar button to open the PEV viewer.
 - Alternatively, select **Model > Utilities > Prediction Error Surface**.

If a model has child nodes you can only select the Prediction Error Surface from the child models.



The default view is a 3-D plot of the PEV surface.

The plot shows where the model predictions are best. This example shows an MBT model response feature. The model predicts well where the PEV values are lowest.

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 drop-down 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 or edit boxes in the frame, 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. See the contour plot section (in "Response Surface View" on page 4-127) for a description of the controls.
 - You can select the **Clip Plot** check box, as shown in the preceding example. Areas that move above the value in the **Clipping envelope** edit box are removed. You can enter the value for the clipping envelope; a white line appears on the plot where the PEV values pass through the clipping value.
 - You can also clip with the boundary model or design constraints if available.
- When you use the PEV viewer to see design properties, optimality values for the design appear in the **Optimality criteria** frame.

Note that you can choose Prediction Error shading in the Response Feature view (in Model Selection or Model Evaluation). This shades the model surface according to Prediction Error values ($\sqrt{\text{PEV}}$). This is not the same as the PEV viewer, which shows the shape of a surface defined by the PEV values. See "Response Surface View" on page 4-127.

Optimality Criteria

No optimality values appear in the **Optimality criteria** frame until you click **Calculate**. Clicking **Calculate** opens the Optimality Calculations dialog. Here iterations of the optimization process are displayed.

In the **Optimality criteria** frame in the PEV viewer are listed the D, V, G and A optimality criteria values, and the values of the input factors at the point of maximum PEV (**Location of G value**). This is the point where the model has its maximum PEV value, which is the G-optimality criteria. The

D, V and A values are functions of the entire design space and do not have a corresponding point.

For statistical information about how PEV is calculated, see the next section “Prediction Error Variance” on page 2-57.

Prediction Error Variance

Prediction Error Variance (PEV) is a very useful way to investigate the predictive capability of your model. It gives a measure of the precision of a model’s predictions.

You can examine PEV for designs and for models. It is useful to remember that:

$$\text{PEV (model)} = \text{PEV (design)} * \text{MSE}$$

So the accuracy of your model’s predictions is dependent on the design PEV and the mean square errors in the data. You should try to make PEV for your design as low as possible, as it is multiplied by the error on your model to give the overall PEV for your model. A low PEV (close to zero) means that good predictions are obtained at that point.

You can think of the design PEV as multiplying the errors in the data. If the design PEV < 1, then the errors are reduced by the model fitting process. If design PEV > 1, then any errors in the data measurements are multiplied. Overall the predictive power of the model will be more accurate if PEV is closer to zero.

You start with the regression (or design) matrix, for example, for a quadratic in N (engine speed) and L (load or relative air charge):

$$\mathbf{X} = \begin{bmatrix} 1 & L_1 & N_1 & L_1^2 & L_1 N_1 & N_1^2 \\ 1 & L_2 & N_2 & L_2^2 & L_2 N_2 & N_2^2 \\ \dots & \dots & \dots & \dots & \dots & \dots \\ 1 & L_n & N_n & L_n^2 & L_n N_n & N_n^2 \end{bmatrix}$$

If you knew the actual model, you would know the actual model coefficients β . In this case the observations would be:

$$y = X\beta + \varepsilon$$

where ε is the measurement error with variance

$$\text{var}(\varepsilon) = \text{MSE}$$

However you can only ever know the predicted coefficients:

$$\hat{\beta} = (X^T X)^{-1} X^T y$$

which have variance

$$\text{var}\hat{\beta} = (X^T X)^{-1} \text{MSE}$$

Let x be the regression matrix for some new point where you want to evaluate the model, for example:

$$x = \begin{bmatrix} 1 & L_{new} & N_{new} & L_{new}^2 & L_{new}N_{new} & N_{new}^2 \end{bmatrix}$$

Then the model prediction for this point is:

$$\hat{y} = x\hat{\beta} = x(X^T X)^{-1} X^T y$$

Now you can calculate PEV as follows:

$$\text{PEV}(x) = \text{var}(\hat{y}) = (x(X^T X)^{-1} X^T)(X(X^T X)^{-1} X^T) \text{MSE}$$

$$\text{PEV}(x) = x(X^T X)^{-1} x^T \text{MSE}$$

Note the only dependence on the observed values is in the variance (MSE) of the measurement error. You can look at the PEV(x) for a design (without MSE, as you don't yet have any observations) and see what effect it will have on the measurement error - if it is greater than 1 it will magnify the error, and the closer it is to 0 the more it will reduce the error.

You can examine PEV for designs or global models using the Prediction Error Variance viewer. When you open it from the Design Editor, you can see how well the underlying model predicts over the design region. When you open it from a global model, you can view how well the current global model predicts. A low PEV (tending to zero) means that good predictions are obtained at that point. See “Prediction Error Variance Viewer” on page 2-54.

For information on the calculation of PEV for two-stage models, see “Prediction Error Variance for Two-Stage Models” on page 2-59.

Prediction Error Variance for Two-Stage Models

It is very useful to evaluate a measure of the precision of the model’s predictions. You can do this by looking at Prediction Error Variance (PEV). Prediction error variance will tend to grow rapidly in areas outside the original design space. The following section describes how PEV is calculated for two-stage models.

For linear global models applying the variance operator to Equation 4-15 yields:

$$\text{Var}(\beta) = (Z^T W^{-1} Z)^{-1} Z^T W^{-1} \text{Var}(P) W^{-1} Z (Z^T W^{-1} Z)^{-1} \text{ so}$$

(2-1)

$$\text{Var}(\beta) = (Z^T W^{-1} Z)^{-1} \tag{2-1}$$

since $\text{Var}(P) = W$. Assume that it is required to calculate both the response features and their associated prediction error variance for the i^{th} test. the predicted response features are given by:

(2-2)

$$\hat{p}_i = z_i \hat{\beta} \tag{2-2}$$

where z_i is an appropriate global covariate matrix. Applying the variance operator to Equation 2-2 yields:

(2-3)

$$\text{Var}(\hat{\mathbf{p}}_i) = \mathbf{z}_i \text{Var}(\hat{\boldsymbol{\beta}}) \mathbf{z}_i^T = \mathbf{z}_i (\mathbf{Z}^T \mathbf{W}^{-1} \mathbf{Z})^{-1} \mathbf{z}_i^T \quad (2-3)$$

In general, the response features are non-linear functions of the local fit coefficients. Let \mathbf{g} denote the non-linear function mapping $\boldsymbol{\theta}_i$ onto \mathbf{P}_i . Similarly let \mathbf{h} denote the inverse mapping.

(2-4)

$$\hat{\boldsymbol{\theta}}_i = \mathbf{h}(\hat{\mathbf{p}}_i) \quad (2-4)$$

Approximating \mathbf{h} using a first order Taylor series expanded about \mathbf{P}_i (the true and unknown fixed population value) and after applying the variance operator to the result:

(2-5)

$$\text{Var}(\hat{\boldsymbol{\theta}}_i) = \dot{\mathbf{h}} \text{Var}(\hat{\mathbf{p}}_i) \dot{\mathbf{h}}^T \quad (2-5)$$

where the dot notation denotes the Jacobian matrix with respect to the response features, \mathbf{P}_i . This implies that $\dot{\mathbf{h}}$ is of dimension $(p \times p)$. Finally the predicted response values are calculated from:

(2-6)

$$\hat{\mathbf{y}}_i = \mathbf{f}(\hat{\boldsymbol{\theta}}_i) \quad (2-6)$$

Again, after approximating \mathbf{f} by a first order Taylor series and applying the variance operator to the result:

(2-7)

$$\text{Var}(\hat{\mathbf{y}}_i) = \left[\frac{\partial \mathbf{f}}{\partial \boldsymbol{\theta}} \right]_{\hat{\boldsymbol{\theta}}} \dot{\mathbf{h}} \text{Var}(\hat{\mathbf{p}}_i) \dot{\mathbf{h}}^T \left[\frac{\partial \mathbf{f}}{\partial \boldsymbol{\theta}} \right]_{\hat{\boldsymbol{\theta}}}^T \quad (2-7)$$

After substituting Equation 2-3 into Equation 2-7 the desired result is obtained:

(2-8)


$$\text{Var}(\hat{\mathbf{y}}_i) = \left[\frac{\partial \mathbf{f}}{\partial \boldsymbol{\theta}} \right]_{\hat{\boldsymbol{\theta}}}^T \mathbf{h}_{z_i} (\mathbf{Z}^T \mathbf{W}^{-1} \mathbf{Z})^{-1} \mathbf{z}_i^T \mathbf{h} \left[\frac{\partial \mathbf{f}}{\partial \boldsymbol{\theta}} \right]_{\hat{\boldsymbol{\theta}}} \quad (2-8)$$

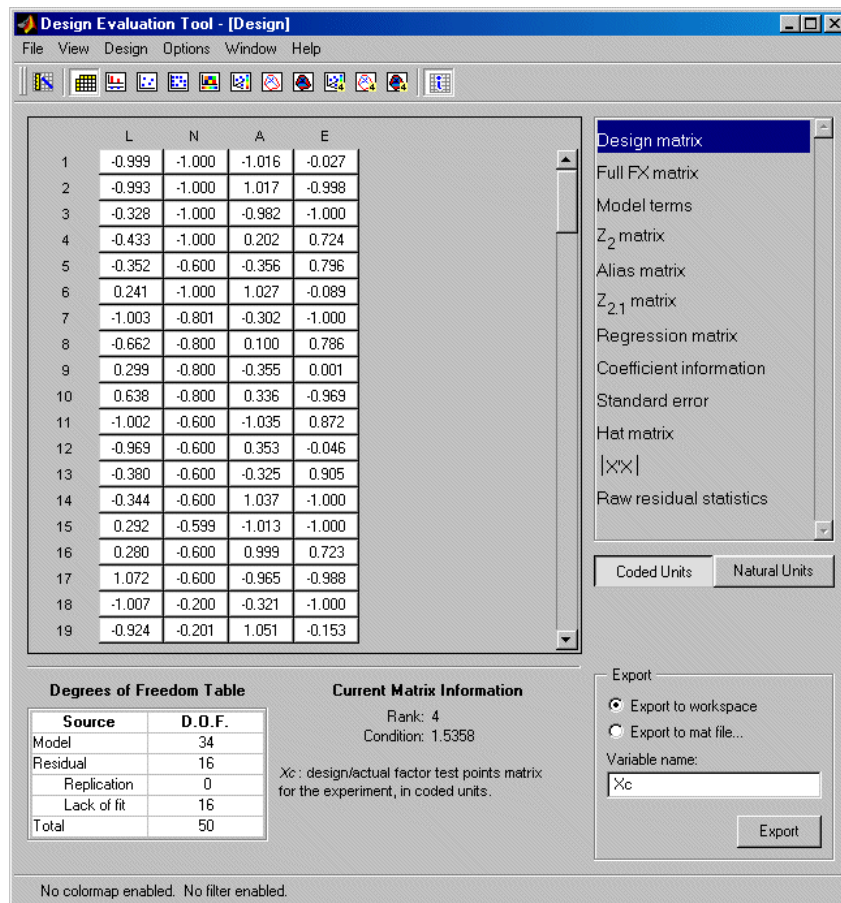
This equation gives the value of Prediction Error Variance.

Design Evaluation Tool

The Design Evaluation tool is only available for linear models.

You can open the Design Evaluation tool from the Design Editor or from the Model Browser windows. From the Design Editor select **Tools > Evaluate Designs** and choose the design you want to evaluate. From the Model

Browser global view, you can click the  button.



Design Evaluation Tool - [Design]

File View Design Options Window Help

	L	N	A	E
1	-0.999	-1.000	-1.016	-0.027
2	-0.993	-1.000	1.017	-0.998
3	-0.328	-1.000	-0.982	-1.000
4	-0.433	-1.000	0.202	0.724
5	-0.352	-0.600	-0.356	0.796
6	0.241	-1.000	1.027	-0.089
7	-1.003	-0.801	-0.302	-1.000
8	-0.662	-0.800	0.100	0.786
9	0.299	-0.800	-0.355	0.001
10	0.638	-0.800	0.336	-0.969
11	-1.002	-0.600	-1.035	0.872
12	-0.969	-0.600	0.353	-0.046
13	-0.380	-0.600	-0.325	0.905
14	-0.344	-0.600	1.037	-1.000
15	0.292	-0.599	-1.013	-1.000
16	0.280	-0.600	0.999	0.723
17	1.072	-0.600	-0.965	-0.988
18	-1.007	-0.200	-0.321	-1.000
19	-0.924	-0.201	1.051	-0.153

Design matrix
Full FX matrix
Model terms
 Z_2 matrix
Alias matrix
 $Z_{2,1}$ matrix
Regression matrix
Coefficient information
Standard error
Hat matrix
 $|X'X|$
Raw residual statistics

Coded Units Natural Units

Export
 Export to workspace
 Export to mat file...
 Variable name:

 Export

Degrees of Freedom Table

Source	D.O.F.
Model	34
Residual	16
Replication	0
Lack of fit	16
Total	50

Current Matrix Information
 Rank: 4
 Condition: 1.5358
 Xc: design/actual factor test points matrix for the experiment, in coded units.

No colormap enabled. No filter enabled.

In the Design Evaluation tool you can view all the information on correlations, covariance, confounding, and variance inflation factors (VIFs). You can investigate the effects of including or excluding model terms aided by this information (you must remove them in the Stepwise window). Interpretation is aided by color-coded values based on the magnitude of the numbers. You can specify changes to these criteria.

When you open the Design Evaluation tool, the default view is a table, as shown in the preceding example. You choose the elements to display from the list on the right. Click any of the items in the list described below to change the display. Some of the items have a choice of buttons that appear underneath the list box.


To see information about each display, click the  toolbar button or select **View > Matrix Information**.

Table Options

You can apply color maps and filters to any items in a table view, and change the precision of the display.

To apply a color map or edit an existing one:

- 1 Select **Options > Table > Colors**. The Table Colors dialog appears.
- 2 Select the check box **Use a colormap for rendering matrix values**.
- 3 Click the **Define colormap** button. The Colormap Editor dialog appears, where you can choose how many levels to color map, and the colors and values to use to define the levels. Some tables have default color maps to aid analysis of the information, described below.

You can also use the **Options > Table** menu to change the precision (number of significant figures displayed) and to apply filters that remove specific values or values above or below a specific value from the display.

The status bar at bottom left displays whether color maps and filters are active in the current view.

When evaluating several designs, you can switch between them with the **Next design** toolbar button or the **Design** menu.

Design Matrix

X_n/X_c: design/actual factor test points matrix for the experiment, in natural or coded units. You can toggle between natural and coded units with the buttons on the right.

Full FX Matrix

Full model matrix, showing all possible terms in the model. You can include and exclude terms from the model here, by clicking on the green column headings. When you click one to remove a term, the column heading becomes red and the whole column is grayed.

The Full FX matrix is the same as the Jacobian for linear models if all terms included. Jacobian only includes 'in' terms. In the general case, jacobian is expressed:

$$(i, j) = df/dp_j (x_i)$$

In the case of linear models and RBFs this simplifies to:

$$(i, j) = j^{\text{th}} \text{ term evaluated at } i^{\text{th}} \text{ data point} = \text{Jacobian matrix.}$$

Model Terms

You can select terms for inclusion in or exclusion from the model here by clicking. You can toggle the button for each term by clicking. This changes the button from in (green) to out (red) and vice versa. You can then view the effect of these changes in the other displays.

Note Removal of model terms only affects displays within the Design Evaluation tool. If you decide the proposed changes would be beneficial to your model, you must return to the Stepwise window and make the changes there to fit the new model.

Z₂ Matrix

Z₂: Matrix of terms that have been removed from the model. If you haven't removed any terms, the main display is blank apart from the message "All terms are currently included in the model."

Alias Matrix

Like the Z₂ matrix, the alias matrix also displays terms that are not included in the model (and is therefore not available if all terms are included in the model). The purpose of the alias matrix is to show the pattern of confounding in the design.

A zero in a box indicates that the row term (currently included in the model) is not confounded with the column term (currently not in the model). A complete row of zeros indicates that the term in the model is not confounded with any of the terms excluded from the model. A column of zeros also indicates that the column term (currently not in the model) could be included (but at the cost of a reduction in the residual degrees of freedom).

A: the alias matrix is defined by the expression

$$A = (X'X)^{-1}X'Z_2$$

Z_{2.1} Matrix

As this matrix also uses the terms not included in the model, it is not available if all terms are included.

Z_{2.1}: Matrix defined by the expression $Z_{2.1} = Z_2 - XA$

Regression Matrix

Regression matrix. Consists of terms included in the model. $n \times p$ matrix where n is the number of test points in the design and p is the number of terms in the model.

Coefficient Information

When you select **Coefficient information**, six buttons appear below the list box. Covariance is displayed by default; click the buttons to select any of the others for display.

Covariance

Cov(b): variance-covariance matrix for the regression coefficient vector b .

$$\text{Cov}(b) = (X'X)^{-1}$$

Correlation

Corr(b): correlation matrix for the regression coefficient vector b .

$$\text{Corr}(b)_{ij} = \frac{\text{Cov}(b)_{ij}}{\sqrt{\text{Cov}(b)_{ii}}\sqrt{\text{Cov}(b)_{jj}}}$$

By default **Correlation** has an active color map to aid analysis. Values below -0.9 are red, -0.9 to -0.7 are orange, -0.7 to 0.707 are black, 0.707 to 0.9 are orange, and greater than 0.9 are red. You can view and edit the color map using **Options > Table > Colors**.

Partial VIFs

Variance Inflation Factors (VIFs) are a measure of the nonorthogonality of the design with respect to the selected model. A fully orthogonal design has all VIFs equal to unity.

The Partial VIFs are calculated from the off-diagonal elements of Corr(b) as

$$\text{VIF}_{ij} = \frac{1}{(1 - \text{Corr}(b)_{ij}^2)} \text{ for } p \geq i > j > 1$$

Partial VIFs also has a default color map active (<1.2 black, >1.2<1.4 orange, >1.4 red). A filter is also applied, removing all values within 0.1 of 1. In regular designs such as Box-Behnken, many of the elements are exactly 1 and so need not be displayed; this plus the color coding makes it easier for

you to see the important VIF values. You can always edit or remove color maps and filters.

Multiple VIFs

Measure of the nonorthogonality of the design. The Multiple VIFs are defined as the diagonal elements of $\text{Corr}(b)$:

$$VIF_i = [\text{Corr}(b)^{-1}]_{ii}$$

Multiple VIFs also has a default color map active (<8 black, 8><10 orange, >10 red). A filter is also applied, removing all values within 0.1 of 1. Once again this makes it easier to see values of interest.

2 Column Corr.

$\text{Corr}(X)$; correlation for two columns of X.

$$w_{ij} = \frac{x_{ij} - \bar{x}_j}{\sqrt{\sum_{i=1}^N (x_{ij} - \bar{x}_j)^2}}$$

Let W denote the matrix of w_{ij} values. Then the correlation matrix for the columns of X (excluding column 1) is $\text{Corr}(X)$, defined as

$$\text{Corr}(X) = W'W$$

2 Column Correlation has the same default color map active as **Correlation**.

Single Term VIFs

Measure of the nonorthogonality of the design. The Single Term VIFs are defined as

$$VIF_{ij} = \frac{1}{(1 - \text{Corr}(X)_{ij}^2)} \text{ for } p \geq i > j > 1$$

Single term VIFs have a default color map active (<2 black, 2>red) and values within 0.1 of 1 are filtered out, to highlight values of interest.

Standard Error

σ_j : Standard error of the j^{th} coefficient relative to the RMSE.

Hat Matrix

Full Hat matrix

H: The Hat matrix.

$$H = QQ'$$

where Q results from a QR decomposition of X. Q is an $n \times n$ orthonormal matrix and R is an $n \times p$ matrix.

Leverage values

The leverage values are the terms on the leading diagonal of H (the Hat matrix). Leverage values have a color map active (<0.8 black, 0.8>orange<0.9, >0.9 red).

|X'X|

D; determinant of X'X.

D can be calculated from the QR decomposition of X as follows:

$$D = \left[\prod_{i=1}^p (R_{1ii}) \right]^2$$

where p is the number of terms in the currently selected model.

This can be displayed in three forms:

$$|X'X|$$

$$\log(|X'X|)$$

$$|X'X|^{(1/p)}$$

Raw Residual Statistic

Covariance

Cov(e): Variance-covariance matrix for the residuals.


$$\text{Cov}(e) = (I-H)$$

Correlation

Corr(e): Correlation matrix for the residuals.

$$\text{Corr}(e)_{ij} = \frac{\text{Cov}(e)_{ij}}{\sqrt{\text{Cov}(e)_{ii}}\sqrt{\text{Cov}(e)_{jj}}}$$

Degrees of Freedom Table

To see the Degrees of Freedom table (and the information about each display), click the  toolbar button or select **View > Matrix Information**.

Source	D.O.F.
Model	5
Residual	14
Replication	0
Lack of fit	14
Total	19

Source	D.F.
Model	p
Residual	n-p
Replication	by calculation
Lack of fit	by calculation
Total	n

Replication is defined as follows:

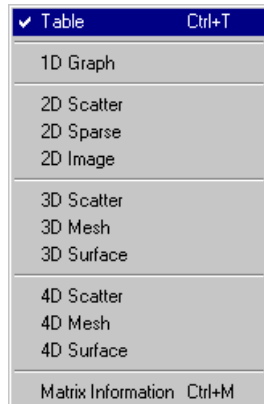
Let there be n_j (>1) replications at the j^{th} replicated point. Then the degrees of freedom for replication are

$$\sum_j (n_j - 1)$$

and Lack of fit is given by $n - p -$ degrees of freedom for replication.

Note: replication exists where two rows of X are identical. In regular designs the factor levels are clearly spaced and the concept of replication is unambiguous. However, in some situations spacing can be less clear, so a tolerance is imposed of 0.005 (coded units) in all factors. Points must fall within this tolerance to be considered replicated.

Design Evaluation Graphical Displays



The Design Evaluation tool has options for 1-D, 2-D, 3-D, and 4-D displays. You can switch to these by clicking the toolbar buttons or using the **View** menu.

Which displays are available depends on the information category selected in the list box. For the Design matrix, (with sufficient inputs) all options are available. For the Model terms, there are no display options.

You can edit the properties of all displays using the **Options** menu. You can configure the grid lines and background colors. In the 2-D image display you can click points in the image to see their values. All 3-D displays can be rotated as usual. You can edit all color map bars by double-clicking.

Export of Design Evaluation Information

All information displayed in the Design Evaluation tool can be exported to the workspace or to a .mat file using the radio buttons and **Export** button at the bottom right. You can enter a variable name in the edit box.

Data

This section discusses the following topics:

Data (p. 3-3)	Overview providing links to all information about using data in the Model-Based Calibration Toolbox.
The Data Editor (p. 3-6)	How to open the Data Editor and use the toolbar and menus to view and manipulate your data.
Data Loading and Merging (p. 3-13)	How to load and merge data from files and the workspace.
Data Editor Views (p. 3-19)	This section describes the functionality available in the different types of views available in the Data Editor. There are plots, tables, and list views for examining your data and design points.
The Variable Editor (p. 3-25)	How to use the Variable Editor to define new variables.
The Filter Editors (p. 3-27)	How to use the Filter Editors to filter your data by record or test, and add test notes to mark every test that fulfills set criteria.
Storage (p. 3-29)	How to use the Storage dialog to store plot preferences, user-defined variables, filters, and test notes so they can be applied to other data sets and other sessions.

Test Groupings (p. 3-33)

How to use the Test Groupings dialog to define groupings for your data. Here you can define rules to collect records of the current data set into groups.

Data Wizard (p. 3-36)

How to use the Data Wizard to select data for modeling. You can also set up matching data to designs by setting tolerances and automatically opening the cluster plots views within the Data Editor.

Matching Data to Designs (p. 3-41)

How to use the cluster plots to match design points to data and select data for modeling.

Data Loading Application
Programming Interface (p. 3-49)

An example to show you how to use the data loading API (application programming interface) to write your own data loading function, plug these functions into the toolbox, and subsequently use data loaded by these functions within the toolbox.

Data

This section describes all aspects of loading, manipulating, and selecting data in the Model Browser. The Data Editor provides a powerful graphical interface for dealing with data:

- You use the Data Editor for loading, filtering, grouping, and editing data, and you can define new variables. You can match data to designs. You can reach the Data Editor from every node in the model tree, so you can also examine and export your modeling data. The Data Editor contains various graphical interfaces for these tasks:
 - Data Import Wizard for loading and merging data
 - Variable Editor, Filter Editor, Test Filter Editor, and Test Notes Editor are dialogs for creating and editing new variables and data filters.
 - Storage dialog for storing new variables, data filters, and plot settings
 - Test Groupings dialog can be used for plotting and manipulating data groups.
 - Within the Data Editor there are 2-D, 3-D, and multiple data plots for viewing data, and cluster plots for viewing data and design points.
 - You use the cluster plots views within the Data Editor for matching data to experimental designs. You can set tolerances for automatic selection of the nearest data points to the specified design points, or select data points manually.
- You use the Data Wizard to select data for modeling. You can also set up matching data to designs by setting tolerances and automatically opening the cluster plots views within the Data Editor. You reach the Data Wizard from test plan level.

You can load and merge data from the following:

- From files (Excel, Concerto, MATLAB)
- From the workspace

See “Data Loading and Merging” on page 3-13.

You can also write your own data-loading functions. See “Data Loading Application Programming Interface” on page 3-49.

Within “The Data Editor” on page 3-6, you can do the following:

- View plots, edit and add data records. See “Data Editor Views” on page 3-19.
- Define new variables. See “New Variables” on page 3-25.
- Apply filters to remove unwanted records, filter by test, and apply notes to tests that fulfill set criteria. See “New Filters” on page 3-27.
- Store and retrieve user-defined variables and filters. See “Storage” on page 3-29.
- Define test groupings to collect data into groups.
- Match data to experimental designs using the cluster plots views.
- Export data and modeling output to file and to the workspace.

You use the Data Wizard to do the following:

- Select the data set and design to use for modeling.
- Select the data signals to use for model input factors (one-stage, or local and global for two-stage).
- Select matching tolerances (if matching data to a design).
- Select data signals for response model input factors.

There is a tutorial to guide you through using the Data Editor. See “Tutorial: Data Editor” in the Getting Started documentation.

Tip for Fast Data Editor Access

When you open the Data Editor from the test plan level, a Cluster Plot is displayed. If you want to open the Data Editor frequently from the test plan level, it will save time if you previously had a Cluster Plot view open when you last used the Data Editor.

The Data Editor retains memory of the types of views open when you last used the Data Editor. If the previous views did not include a Cluster Plot view, then

when you open the Data Editor from the test plan level the existing views are saved and replaced with a Cluster Plot view and a Cluster Information list view. The views revert to the previous views when you close the Data Editor. This process takes longer to display than if you had an existing Cluster Plot view (the Cluster Information list view is not required).

The Data Editor

To reach the Data Editor,

- From the test plan node, choose **TestPlan > Select Data**.
- Alternatively, from the project node, do one of the following:
 - Choose **Data > New Data, Copy Data, or Edit Data**.
 - Select any of the equivalent toolbar buttons.
 - Double-click a data set in the **Data Sets** pane.
- From any modeling node, click the View Modeling Data toolbar button.

As can be seen in the following example, you can split the view to display several plots at once. As in the Design Editor, you 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, 3D plots, multiple data plots, data tables, and list views of filters, variables, test filters, and test notes.

Data Editor - Data Object

File View Tools Window Help

266 / 270 records
7 + 0 variables
27 / 27 tests

Source files: D:\sandbox\MBC\toolbox\mbc\mbc

Tests: 537, 586, 587, 588, 589, 590, 591, 592, 593, 594, 595, 596, 597, 598, 600, 601, 602, 603, 604, 605, 606, 607, 608, 609, 610, 611, 612

Legend: Data remaining (green), Data removed (red), Variables added (blue)

2D Data Plot

Tests: 537, 586, 587, 588, 589, 590, 591

Y-axis sel...: tq [ft lbf]

X-axis sel...: spark [deg]

Data Table

Scroll to test number: 589

test \ name	afr	egr
589	11.05	0
	11.05	0
	11.07	0
	11.09	0
	11.09	0
	11.1	0
	11.09	0

2D Data Plot

Tests: 586, 587, 589

Y-axis sel...: n [rpm]


X-axis sel...: spark [deg]

Filter List

Filter Expression	Results
✓ spark < 50	Filter successfully applied : 4 records excluded

Data has 266/270 Records, 7 + 0 Variables, and 27 Tests.

See “Data Editor Views” on page 3-19 for details on the functions available in each type of view.



The list box at the top right contains the source file information for the data, and other information is displayed on the left: the numbers of records, variables, and tests it contains. See the preceding example. The bars and figures show the proportion of records removed by any filters applied, and the number of user-defined variables is shown. For this example with two user-defined variables added to a data set originally containing seven variables, you see '7 + 2 variables.'

The **Tests** list pane on the left is constant. Tests selected here apply to 3-D plots, multiple data plots, and tables. It does not apply to 2-D plots because they have their own independent test controls. If you are viewing read-only local modeling data, the selected test is shown in the **Tests** pane and remains synchronized if you change test in the Model Browser.


See “Data Editor Toolbar and Menus” on page 3-9 for other controls.

By default new data sets are called Data Object. You can change the names of data sets at the project node by select-clicking a data set in the **Data Sets** list, or by pressing **F2** (as when selecting to rename in Windows Explorer).

To edit data plot properties (2-D, 3-D, or multiple), you can right-click the plot and select **Viewer Options > Properties**. Here you can choose to show the legend and the grid. You can choose the line style if you want to connect the data points and the data marker point style, if any. **Reorder X Data** (2-D plots) redraws the line joining the points in order from left to right. The line might not make sense when drawn in the order of the records.

Note Dismissing the Data Editor automatically brings up the Data Wizard if you entered it from the test plan level.

Data Editor Toolbar and Menus



- Import Data — See “Loading Data from File” on page 3-13; also in the **File** menu.
- Edit Variables — Opens the Variable Editor.
- Edit Filters — Opens the Filter Editor.
- Change Test Grouping — Opens the Test Groupings dialog; also in the **Tools** menu.
- Split View Horizontally — Divides the currently selected plot into two plots side by side. Also in the **View** menu.
- Split View Vertically — Divides the currently selected plot into two plots one above the other. Also in the **View** menu.
- Open Storage Window — Opens “Storage” on page 3-29Storage where you can store and retrieve variables, filters, and plot settings.

The **Window** and **Help** menus are the same as everywhere in the Model Browser. See “Window Menu” on page 1-18 and “Help Menu” on page 1-18.

File Menu

- **Import**
 - **File** — See “Loading Data from File” on page 3-13.
 - **Workspace** — See “Loading Data from the Workspace” on page 3-15.
 - **Excel** — Loads Excel files.
- **Export**
 - **Workspace** — Exports data to the workspace.
 - **Excel** — Exports data to an Excel file.
- **Print** — Prints your currently selected plot.
- **Close** — Closes the Data Editor.

View Menu

- **Current View**
 - Here you can choose from **2-D Data Plot**, **3-D Data Plot**, **Multiple Data Plots**, **Data Table**, **Variable Definitions**, **Filter Definitions**, **Test Filter Definitions**, and **Test Note Definitions**, **Notes View**, **Cluster View**, **Cluster Information**. A tick shows the type of view currently selected.
- **Viewer Options** — These change depending on the currently selected view, and are duplicated in the right-click menu in each view.
 - **Properties** — Only for data plots (2-D, 3-D, or multiple), not table or list views. Opens a properties dialog for the current plot. Here you can change settings for grids, lines, markers, legends, and axes. For 2-D and multiple data plots you can choose whether to plot bad data (that is, outliers you have removed with the **Tools** menu).
 - Multiple data plot views have several specific options. You can select **Add Plot** to add another subplot to the current view, or **Remove All Plots** to clear the current view. For the selected subplot (right-click the plot, or click first and use the menu) you can select **Plot Variables** to set up axes or **Remove Plot** to delete.

- **Split View** — Divides the currently selected view and adds the selected option. You can choose from the same options available under **Current View**.
- **Split View Horizontally** — Divides the currently selected plot into two plots side by side. The new plot is a 1-D plot by default; you can select other views using the right-click context menu. Also in the toolbar.
- **Split View Vertically** — Divides the currently selected plot into two plots one above the other. Also in the toolbar.
- **Delete Current View** — Deletes the currently selected view.
- **Print to Figure** — Copies the currently selected view to a Figure window. See Toolbar above.

Tools Menu

Note that if you are viewing modeling data in the Data Editor (from a modeling node) you cannot change the data, so editing and user-defined variables and filtering are not enabled; you cannot use the **Tools** menu to alter your modeling output.

- **Variables**
 - **Add** opens the Variable Editor. Here you can add new variables to the data set by typing into the edit box.
 - **Edit** is only enabled if you already have user-defined variables. This also opens the Variable Editor with an additional list on the left containing your defined variables. Select these from the list and edit in the edit box. You can also add new variables here by clicking the button to add a new item to the list.

The toolbar button Edit Variables performs the same functions.

- **Filters**
 - **Add** and **Edit** both open the Filter Editor, in a similar way to the variable items. Here you can remove records by imposing constraints on the data. Also in the toolbar.
 - **Remove Outliers** — Removes points you have selected in plots (outlined in red). Use the check box to **Show Bad Data** (in the Plot Properties dialog) to see removed points on plots.

- **Restore Outliers** — Opens a dialog where you can choose any or all removed points to restore.
- **Test Filters** — **Add** and **Edit** both open the Test Filter Editor.
- **Test Notes** — **Add** and **Edit** both open the Test Note Editor.
- **Change Test Groupings** — Opens the Test Groupings dialog; also in the toolbar.
- **Tolerances** — Opens the Tolerance Editor. These settings are for matching data to designs using the cluster views. See “The Tolerance Editor” on page 3-45.
- **Open Storage**— Opens Storage where you can store and retrieve variables, filters, and plot settings.

Data Loading and Merging


Data can be loaded and merged from files, from the workspace, and from tailor-made Excel sheets. A test plan can only use a single data set, so you need to use the merge functions to combine data variables from more than one source file in order to incorporate desired variables into one model. See

- “Loading Data from File” on page 3-13
- “Loading Data from the Workspace” on page 3-15
- “Tailor-Made Excel Sheets” on page 3-17
-

Note that you can also write your own data-loading functions. See “Data Loading Application Programming Interface” on page 3-49. This includes a description of the correct structure for importing data from the workspace. This is the same structure of data as when you export to the workspace.

Loading Data from File

From the project node choose **Data > New Data**. The Data Editor appears.

- Click the Open File button in the toolbar  to load data from a file.
- Alternatively, choose **File > Import > File**.

The Data Import Wizard appears, to help you select a file.

Alternatively, from the test plan node,

- 1 Select **TestPlan > Select Data**. The Data Wizard appears.
- 2 Click **Load new data set** and the Data Import Wizard appears.

Note If you already have some data loaded, the Data Import Wizard can merge new data into the existing data set.

The Data Import Wizard

- 1** To import data from a file, enter the file pathname in the edit box, or use the Browse button to find and select the data file. Double-click to load the file.

The drop-down menu contains the file types recognized by the Model Browser (Excel, Delimited text files, MATLAB). The default Auto tries to determine the type of file by looking at the file extension. For text files, see “Text File Formats” on page 3-15.

- 2** For Excel files with multiple sheets, you must next select the sheet you want to use and click **OK**.
- 3** The Import Wizard now 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. You can edit variable names and units in this list. Click **Finish** to accept the data, unless you have data to merge. See “Merging Data” on page 3-14.

The Data Import Wizard disappears and the view returns to the Data Editor (or the Data Wizard if you are selecting data from a test plan — see “Data Wizard” on page 3-36).

Merging Data

- 1** If you already have some data loaded, you cannot click **Finish** but must click **Next** instead. This brings you to the data merging screen of the wizard.
- 2** Here you must choose one of the radio buttons:
 - Merge extra channels (more variables)
 - Merge extra records (more data, same variables)
 - Overwrite old data (use only the new data)
- 3** Click **Next**. The wizard summary screen appears, showing the results of the merge. You can click **<Prev** to go back and change merging options.
- 4** Click **Finish** to accept the data and return to the Data Editor.

Note Obviously the merge might not succeed if the data sets do not contain the same variables. A message appears if the merge is impossible when you click **Next** and you must make another choice.

Text File Formats. Select Delimited Text File from the **Open As** list to read delimited text files into MBC. These files can be delimited by tabs, |, commas or spaces. They may optionally contain a variable name header and units header line.

Select Delimited Text Sweep File from the **Open As** list to read delimited text files with multiple blocks of data into MBC. These files can be delimited by tabs, |, commas or spaces. Each block of data may optionally contain a variable name header and units header line. The first block's names and headers are used, subsequent rows of non-numeric data are stripped out. The data can be split into multiple blocks where each block contains data for a separate sweep. Within each block, a column that consists of a value in the first row and empty values in the remaining rows for initial columns will be assumed as meaning that for these columns each row should take its value from the first row in the block. Note that the data is imported as a single array and you must use the Test Groupings window to decide how to sort the data into sweeps. See "Test Groupings" on page 3-33.

Loading Data from the Workspace

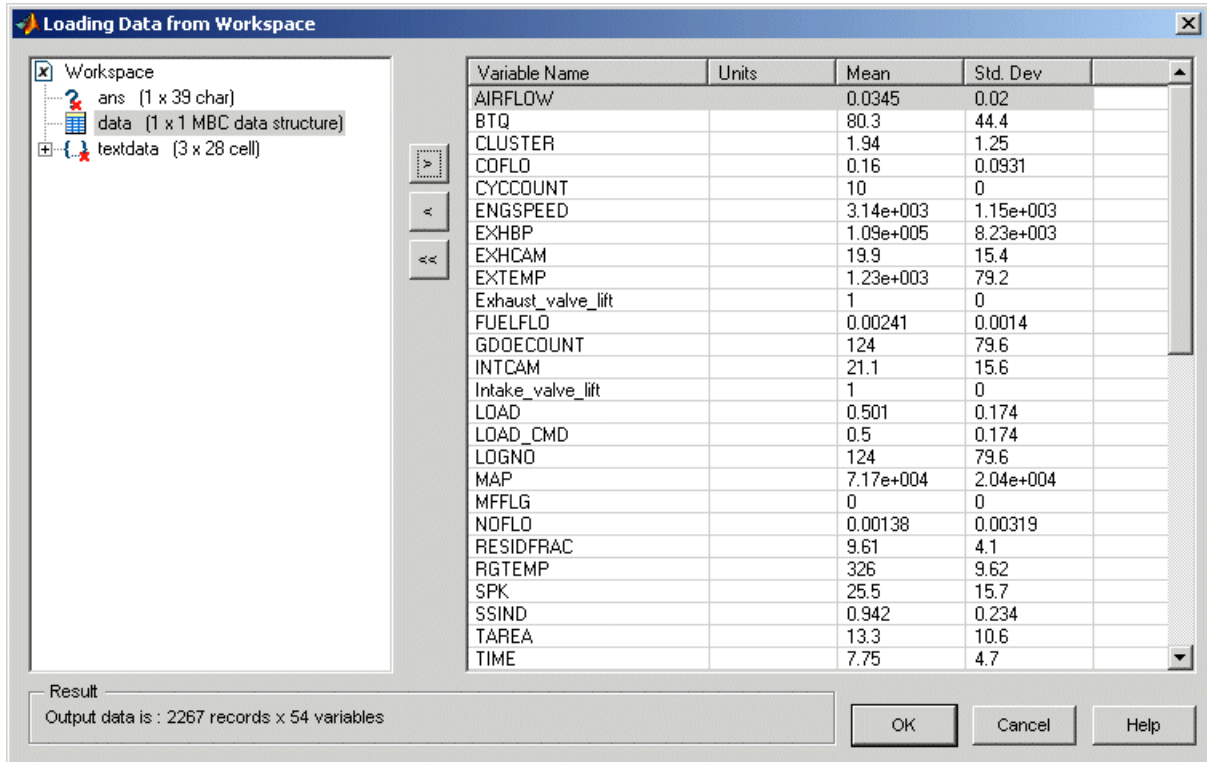
- 1 From the project node choose **Data > New Data**.

The Data Editor appears.

You can import variables from the MATLAB® workspace by choosing **File > Import > Workspace**.

See "Data Editor Toolbar and Menus" on page 3-9.

The Loading Data from Workspace dialog appears.



Variables in the workspace are displayed in hierarchical form in the top left pane. Select a variable here, and information about that variable is displayed in the pane on the right.

- 1 Select a variable to import in the tree at top left.
- 2 Click the **Add** button.

The number of records and variables appears in the **Output Data** pane. You can add variables (one at a time) as many times as you like (as long as there are no name conflicts).

You can double click to edit variable names and units.

- 3 You can use the buttons to remove selected variables one at a time from the right list, or remove all variables at once.
- 4 Click **OK** to accept the data to import and return to the Data Editor.

Data Merging

If you already have data loaded, the Data Merging Wizard appears, where you must choose one of three radio buttons:

- Merge extra channels (more variables)
- Merge extra records (more data, same variables)
- Overwrite old data (use only the new data)

Click **Next**. The wizard summary screen appears, showing the results of the merge. You can click **<Prev** to go back and change merging options.

Note Obviously the merge might not succeed if the data sets do not contain the same variables. A message appears if the merge is impossible when you click **Next**, and you must make another choice.

Click **Finish** to accept the data and return to the Data Editor.

Tailor-Made Excel Sheets

The Data Editor can create a tailor-made Excel sheet for you to fill with data and then import. This sheet will be in the format the Data Editor expects to import data.

- 1 Select **File > Import Excel**.

Excel is opened with a new sheet created, containing the labeled rows Name, Unit, and Data.

- 2 Copy your data and variable names into this sheet, then click **Next** in the wizard to import into the Data Editor.

If the data has been entered in the columns in a way that the Data Editor expects, a summary screen shows you information about the numbers, ranges, means, units, and standard deviations of the records and variables you can import.

3 Click **Finish** to import the data.

Data Editor Views


You can open multiple different views at the same time in the Data Editor, as shown in the following example. You can display 2-D plots, 3-D plots, multiple data plots, cluster plots, data tables, and list views of your filters, variables, and test notes. Use the toolbar buttons, the **View** menu, or right-click a view title bar to split views and change view types.

Functions of different views are described in the following sections:

- “2-D Plots” on page 3-20
- “3-D Plots” on page 3-21
- “Multiple Data Plots” on page 3-21
- “Cluster Plots” on page 3-22
- “List Views — Variables, Filters, Test Filters, Test Notes, and Clusters” on page 3-23
- “Table View” on page 3-23
- “Notes View” on page 3-24

2-D Plots

In the 2-D plot view you can select combinations of variables and tests to plot from the list boxes on the left. Multiple selection of tests and y -axis variables is possible — as show in the following figure, multiple tests are selected to view several tests simultaneously.



See “Data Editor Toolbar and Menus” on page 3-9. Note that you use the 2-D plot view controls for selecting the tests to display. This is unlike the other plots (3-D and multiple data plots), which you control by using the **Tests** list box, always visible on the left of the Data Editor.

To edit data plot properties, you can right-click the plot and select **Viewer Options > Properties**. Here you can choose to show the legend and the grid. You can choose the line style if you want to connect the data points and the data marker point style, if any. **Reorder X Data** redraws the line joining the points in order from left to right. The line might not make sense when drawn in the order of the records. You can use the **Show Bad Data** check box to plot bad data (that is, outliers you have removed).

Click points to select outliers, and remove them by selecting **Tools > Filters > Remove Outliers** (or use the keyboard shortcut **Ctrl+A**). Select **Tools > Filters > Restore Outliers** (or use the keyboard shortcut **Ctrl+Z**) to open a dialog where you can choose to restore any or all removed points.

3-D Plots

In 3-D data plot views you can select the variables for each axis from the drop-down menus, and rotate the plot by clicking and dragging. To edit data plot properties, you can right-click the plot and select **Viewer Options > Properties**. Here you can choose the color and style of the axes, whether to show the grid in each axis, and perspective or orthographic axes projection.

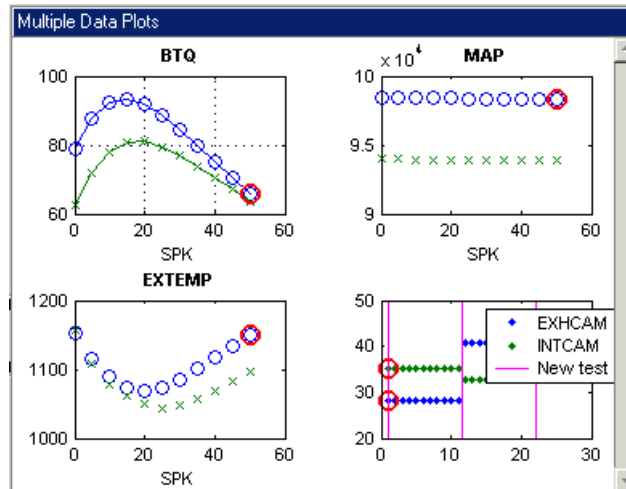
Multiple Data Plots

Here you can add as many 2-D plots as desired to the same view, plotting the same selection of tests in a variety of different plots. Use the right-click context menu to add and remove plots, select plot variables, and edit plot properties. You can select single or multiple Y variables to plot against a single X variable (or no X variable) in the Plot Variables dialog. Select tests to display in the list on the left of the Data Editor, as for 2-D and 3-D plots and the table view.

Note that you can select different plot properties and variables for each plot within the Multiple Data Plots view, as shown in the example following.

Click to select a plot (or right-click) before selecting **Viewer Options > Plot Variables** or **Plot Properties**. For each plot you can use the same plot properties options as for the single 2-D data plots. You can choose to show the legend and the grid. You can choose the line style if you want to connect the data points and the data marker point style, if any. **Reorder X Data** redraws the line joining the points in order from left to right. The line might not make sense when drawn in the order of the records. You can use the **Show Bad Data** check box to plot bad data (that is, outliers you have removed).

Click points to select outliers, and remove them by selecting **Tools > Filters > Remove Outliers** (or use the keyboard shortcut **Ctrl+A**). Selected outliers are outlined in red, as shown in the following example.



Select **Tools > Filters > Restore Outliers** (or use the keyboard shortcut **Ctrl+Z**) to open a dialog where you can choose to restore any or all removed points.

Cluster Plots

You use these views for matching data to design points. Use the cluster information list view with the cluster view to examine your data and design. Click points in the cluster view to select them across the Data Editor — that

is, the selected points are displayed in the table view and other data plot views (except 2-D plots, which have separate controls).

See “Matching Data to Designs” on page 3-41 for detailed information on cluster plot controls.

List Views – Variables, Filters, Test Filters, Test Notes, and Clusters

These views show lists and information such as variable and filter definitions, the notes applied to filtered tests, and the data and design points in selected clusters. Variable and Filter List views show the definitions of each variable or filter. Double-click to select particular filters or variables to edit. See “The Variable Editor” on page 3-25 and “The Filter Editors” on page 3-27.

The Test Notes List view shows the rules used to define notes on the data, along with the actual note, the color of the note, and the number of tests to which that note applies. The specified rule is applied to each test in turn to decide if that test should be noted: e.g., $\text{mean}(\text{TQ}) > 0$ evaluates the mean torque for each test and notes those tests where the value is greater than zero. The list of individual tests and notes appears in the “Notes View” on page 3-24.

In the Cluster List view you can choose to include or exclude design and data points by using the check boxes in cluster list views. See “Matching Data to Designs” on page 3-41 for more information on using these.

Table View

In the Table view, you can view your data, edit, and add records.

Points you have selected in plots (by clicking) are outlined in red in the table. Points you have removed as outliers are light red in the table with a filter icon next to the row number. Edited cells become blue.

Select **View > Viewer Options** or right-click the title bar to reach these options.

- **Duplicate Selected Records** — First select one or more records, then use this option to duplicate them. Each duplicate appears directly underneath

the parent record. Edit duplicates to create new records. You must first select **Allow Editing** to enable this option.

- **Undo Edits in Selected Region** — You can click and drag to highlight a region, then use this option to reverse any edits in the highlighted area. You must first select **Allow Editing** to enable this option.
- **Allow Editing** — Toggles editing, and, as a side effect, causes all records to be shown, including those which are filtered out. Records which are filtered out appear light red in the table, with a filter icon next to the row number. You can alter records by clicking a cell and then typing a new value. Edited cells become blue. Editing the value of a cell may cause that row to be filtered out. If so, the background colour of the row will change after the cell has been edited.
- **Select Columns To Display** — Opens a dialog where you can use the check boxes to select the columns to display in the table. Select a column, then press **Ctrl+A** to select all columns, and then you can select or clear all check boxes with one click. You can click and drag column headers in the table view to rearrange columns.

Notes View


This view shows a list of all tests in the data, along with those notes that apply to particular tests. Each note that applies to one or more tests has its own column in the view. You can sort the columns in the list by clicking the column header (once to sort ascending and once more to sort descending) allowing you to quickly identify the tests that satisfy a given note definition — this facility could be used to screen incoming data for problem tests.

You can click a particular test in the Notes view list to select that test across the Data Editor. 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.

The Variable Editor

New variables you create in the Variable Editor can be viewed in the Data Editor by using the right-click menu in any view to select **Current View > Variable List**. Alternatively, select the menu item **View > Split View > Variable List** to add this new view to the currently selected plot.

You create variables by doing the following:

- Clicking the  toolbar button
- Selecting **Tools > Variables > Add**
- Alternatively, by selecting an existing **Variable List** view by clicking in it, then pressing **Insert**

The Variable Editor appears. See below for instructions.

You can also load user-defined variables from storage.

You can edit existing variables

- Directly, by select-clicking in the **Variable Expression** pane or pressing **F2** (as in renaming in Windows Explorer)
- By double-clicking, which opens the Variable Editor
- By choosing **Tools > Variables > Edit**
- You can delete variables by selecting them in the **Variable Expression** pane and pressing **Delete**

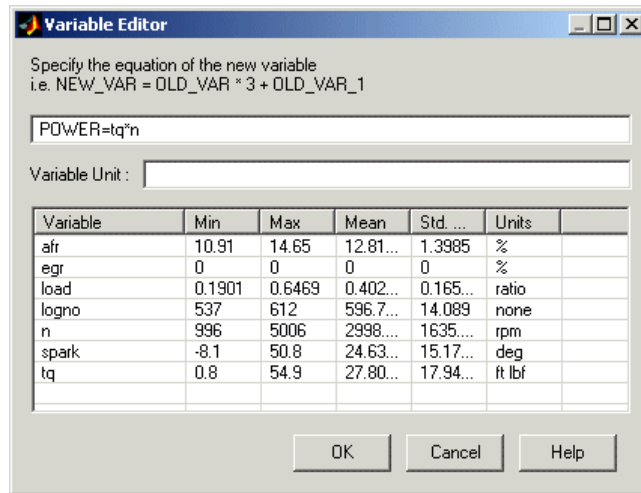
New Variables

Variable Editor

You can define new variables in terms of existing variables:

- Define the new variable by writing an equation in the edit box at the top of the Variable Editor dialog.

You can type directly in the edit box or add variable names by double-clicking them. In the case of variable names especially, this latter method avoids typing mistakes. Variable names are case sensitive.



The preceding example shows a definition for a new variable called POWER that is defined as the product of two existing variables, tq and n, by entering $\text{POWER} = \text{tq} \times \text{n}$.

- Click **OK** to add the new variable to the current data set.


Note that the Variable Editor looks different depending on whether you opened it to create a new variable or edit an existing one. The example above shows the editor when adding a new variable. If you open the editor to edit a variable there is an additional list on the left. You can choose which of your existing variables to edit from this list, or click the button to add a new item to the list if you want to add a new variable.

The Filter Editors

New Filters

New filters you create in the Filter Editor can be viewed in the Data Editor by using the right-click menu in any view to select **Current View > Filter List**. Alternatively, select the menu item **View > Split View > Filter List** to add this new view to the currently selected plot. Their effects are shown graphically in the bars at the top of the Data Editor — removed data is shown in red.

You can create filters

- By clicking the toolbar button 
- By selecting **Tools > Filters > Add**
- Alternatively, by selecting an existing **Filter List** view by clicking in it, then pressing **Insert**

The Filter Editor appears. See below for instructions.

You can also load user-defined filters from Storage.

After you create them, filters can be edited in the same way as variables,

- Directly, after you select-click them in a **Filters List** pane, or by pressing **F2**
- Using **Tools > Filters > Edit**, which opens the Filter Editor
- By double-clicking, which also opens the Filter Editor
- Delete filters by selecting them and pressing **Delete**.

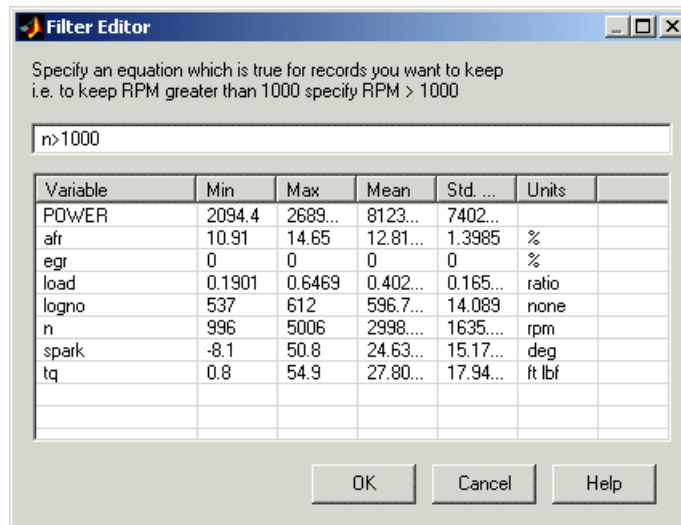
Test Filters and Test Notes

Similarly you can add test filters (to filter out entire tests, instead of individual observations) and test notes (to mark every test that fulfills set criteria). The Test Filter Editor and Test Notes Editor can be reached from the **Tools** menu. You can view these filters in the Data Editor in the same way as the other filters by selecting **Test Filter List** or **Test Notes List** as

the type of view. You define, edit, store, and delete these filters in the same way, and they can also be loaded from storage.

Filter Editor

A filter is a constraint on the data set used to exclude some records. You define the filter using logical operators on the existing variables.



In the preceding example, $n > 1000$, the effect of this filter is to keep all records with speed (n) greater than 1000.

Click **OK** to impose new filters on the current data set.

Note that the Filter Editor looks different depending on whether you opened it to create a new filter or edit an existing one. The example above shows the editor when adding a new filter. If you open the editor to edit a filter there is an additional list on the left. You can choose which of your existing filters to edit from this list, or click the button to add a new item to the list if you want to add a new filter.

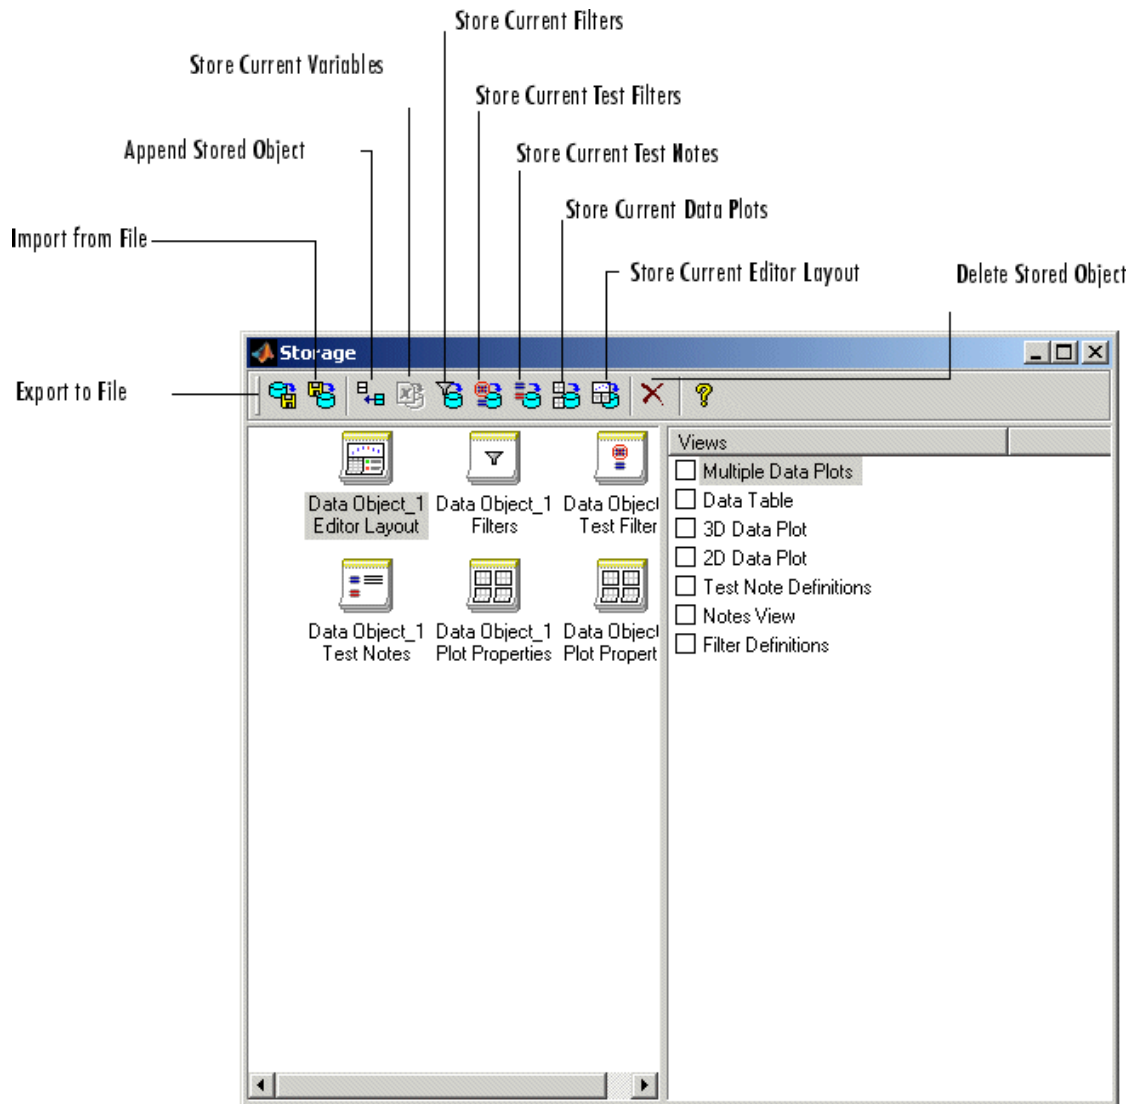
Storage

Storage allows you to store plot preferences, user-defined variables, filters, and test notes so that they can be applied to other data sets loaded later in the session.

You can open the Storage window from the Data Editor window in either of these ways:

- Using the menu **Tools > Open Storage**
- Using the toolbar button 

The toolbar buttons **Store Current Variables** and **Store Current Filters, Test Filters**, or **Test Notes** allow you to put all user-defined variables and filters from the current session into storage. They appear in the Storage window.



All stored user-defined variables and filters appear here regardless of which project is open — once created and brought into storage, they remain there. If you do not delete them, they are there indefinitely. You can also store plot

settings for multiple data plot views with the toolbar button Get Current Data Plots. All the toolbar functions are described below.

You can select Export to File to send the stored objects to a file. You might do this to move the objects to a different user or machine. Select Import from File to bring such variables and filters into storage, and use Append Stored Object to add items from storage to your current project.

You can edit the names of stored objects, by select-clicking as in Windows Explorer, or by clicking once and pressing **F2**.

Note that if you enter the data editor from a modeling node you cannot append variables or filters, as modeling data is read only.

The toolbar buttons in Storage allow the following functions:

- Export to File sends Storage objects to file so they can be sent to another user or machine. Objects remain in storage indefinitely unless deleted; export is only for transporting elsewhere.
- Import from File loads storage objects from file; it does not overwrite current objects.
- Append Stored Object appends the currently selected Storage object to those already in the Data Editor window. Double-clicking an object in storage also appends it to the current session (except Data Editor Layout, which replaces all the current views with the stored layout). If you double-click an object in the left pane the entire object is appended. The pane on the right displays the details of the current object selected in the left pane (variables, filters, or plot properties). In the right pane you can double-click to add these elements individually. For example, you could add a single plot out of several, or a single variable. You can also hold down **Shift** to make multiple selections here.
 - Variables and filters will appear in variable and filter list views. Note that the views are not a requirement for appending variables or filters.
 - Stored Plot Properties (plotted variables, grid style, legend, line style, and so on) apply to multiple data plot views. If you have more than one multiple data plot view open, the stored plot properties are appended to the currently selected view (with the highlighted title bar), or to all views if none are currently selected. The plots stored in a Plot Properties

object are added to multiple data plot views as additional plots — they do not overwrite existing plots. You cannot append stored plot properties if you do not have a multiple data plot view open.


- **Store Current Variables** creates a Storage object containing definitions of all current user-defined variables.
- **Store Current Filters** creates a Storage object containing definitions of all current filters. Similarly, **Store Current Test Filters** and **Store Current Test Notes** create storage objects containing test filters and test notes.
- **Store Current Data Plots** creates a storage object containing the current settings for multiple data plot views (variables plotted and styles for grids, lines, and legends). The settings are taken from the currently selected multiple data plot view. If there are multiple views and none are selected the settings from all the plots are concatenated into a single Plot Properties object.
- **Store Current Data Editor Layout** creates a storage object containing the current settings for all the views. Saving a set of views and then double-clicking this object in Storage can be a very quick way to achieve a variety of plots and list views in a new session.

Delete Stored Object deletes the currently selected Storage object.

Test Groupings


The Define Test Groupings dialog collects records of the current data set into groups; these groups are referred to as *tests*. Test groupings are used to define hierarchical structure in the data for two-stage modeling.

You access the dialog from the Data Editor by doing one of the following:

- Using the menu **Tools > Change Test Groupings**
- Using the toolbar button 

When you enter the dialog, no plot is displayed.

1 Select a variable in the list box to use in defining groups within the data.

2 The **Add Variable** button () adds the currently selected variable in the **Variables** list to the list view on the left. Alternatively, double-click variables to add them.


You can now use this variable to define groups in the data.

In the following example, the variable *n* is being used to define groups. The maximum and minimum values of *n* are displayed.


The **Tolerance** is used to define groups: on reading through the data, when the value of *n* changes by more than the tolerance, a new group is defined. You can change the tolerance by typing directly in the edit box.

You can define additional groups by selecting another variable and choosing a tolerance. Data records are then grouped by *n* or by this additional variable changing outside their tolerances.

You can plot variables without using them to define groups by clearing the **Group By** check box.

You can remove variables from consideration by selecting the unwanted variable in the list view (the selection is highlighted in blue) and clicking the Remove Variable button .

The following plot shows the scaled values of all variables in the list view (the color of the **Tolerance** text corresponds to the color of data points in the plot). Vertical pink bars show the tests (groups). You can zoom the plot by **Shift**-click-dragging or middle-click-dragging the mouse.



Test number variable contains a drop-down menu showing all the variables in the current data set. You can select any of these to number the tests (for example, lognumber could be useful (instead of 1,2,3...) if the data was taken in numbered tests and you want access to that information during modeling).

Every record in a test must share the same test number to identify it, so when you are using a variable to number tests, the value of that variable is taken in the first record in each test.

Test numbers must be unique, so if any values in the chosen variable are the same, they are assigned new test numbers for the purposes of modeling. (This does not change the underlying data, which retains the correct lognumber or other variable.)

Reorder records allows you to reorder records in the data set. This sorts records before grouping. Otherwise, the groups are defined using the order of records in the original data set.

Show original displays the original test groupings if any were defined.

One test/record defines one test per record, regardless of any other grouping. This is required if the data is to be used in creating one-stage models.

Clicking **OK** accepts the test groupings defined and dismisses the dialog.

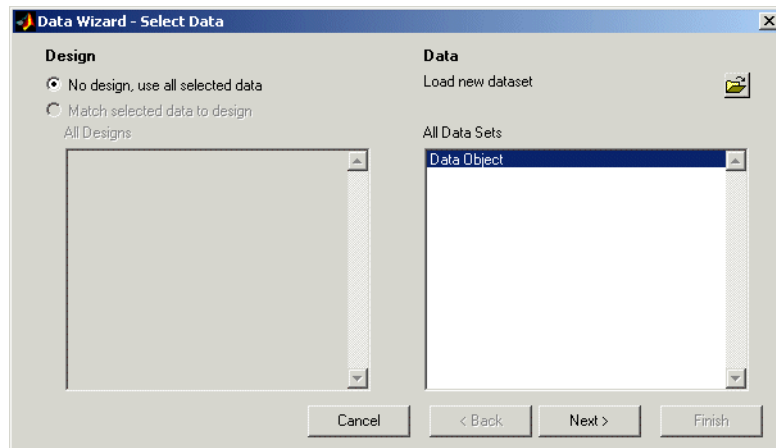
Data Wizard

After setting up a new test plan, *if no data has been selected in that test plan*, then either of the following brings up the Data Wizard.

- Choosing **Select Data** from the test plan node (toolbar button or **TestPlan** menu item)
- Double-clicking the Responses block in the test plan diagram

Step 1: Select Data Set

Use the first screen of the wizard to select the data set to build models from. You can also select whether to use all the data set or to match the data to a design, if any designs are in use in the test plan. Designs appear in the left list box.

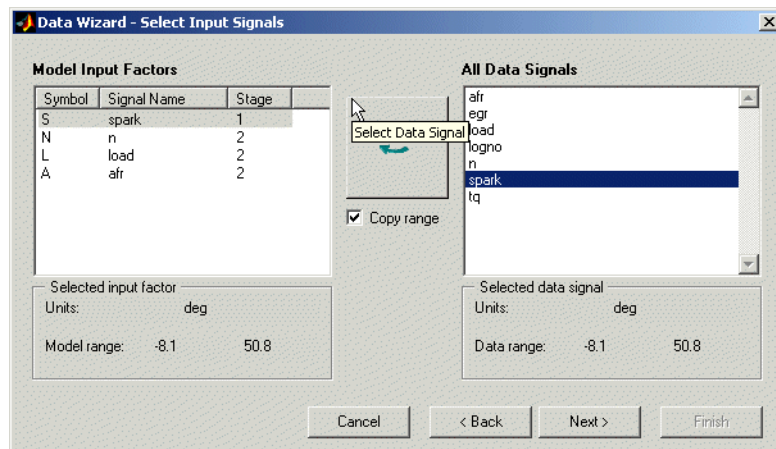


Step 2: Select Input Signals

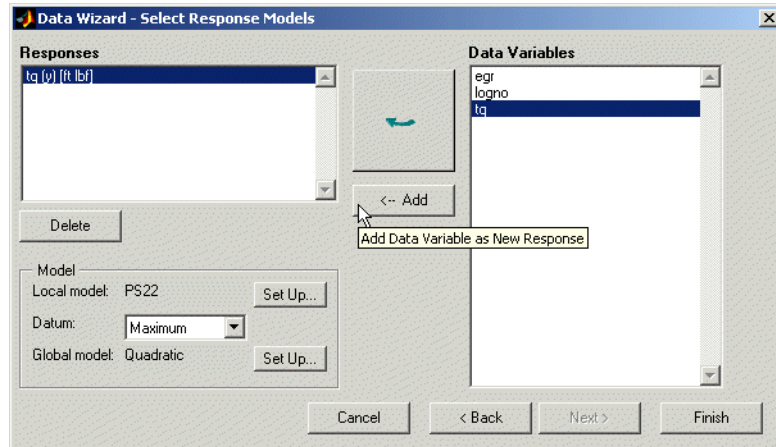
Select the input signals for the model (on all stages of the hierarchical model) from the list box of data signals on the right, and match them to the correct model variables using the big button with the arrow. Double-clicking an item in the data signals list selects that signal for the currently selected input factor (with the range if the **Copy range** check box is selected) and then moves to the next input.

If you entered the correct signal name at the model setup stage, the appropriate signal is automatically selected as each model input factor is selected. This can be time-saving if there are many data signals in a data set. If the signal name is not correct, you must select the correct variable for each input by clicking.

Select the check box **Copy range** if you want to use the range of the selected data signal for the model input range. Ranges are not automatically copied, although stored templates have the ranges that were set when the template was saved.



Step 3: Select Response Models



Starting from scratch (with an empty **Responses** list box), select the desired response in the **Data Variables** list and click **Add**.

In the preceding example, the test plan template specified torque as the response model, so it already appears in the **Responses** list box. If you want to change the response, select another variable and click the large button with the arrow. This replaces the current selected response. The previous response appears in brackets to inform you what has changed.

When there is already a response in the list box, clicking **Add** does not replace the selection but increases the number of responses. The replace button (with the arrow) is not available when the **Responses** box is empty.

You can use **Delete** to remove selected responses. You can select datum models (if the local model supports them), and you can use the **Set Up** buttons to change the local and global models. See “Global Model Setup” on page 4-6, “Local Model Setup” on page 4-26, and “Datum Models” on page 4-72 for details.

Step 4: Set Tolerances

Setting tolerances is only relevant if you are matching data to a design. You can only match data to designs for global models.

You can also edit tolerances later using the **Tools** menu in the Data Editor window. See “The Tolerance Editor” on page 3-45 for definitions of clusters.

Set **Input factor tolerances** for each variable to determine the size of the tolerance in each dimension. This is used for selecting data. The tolerance in each dimension determines the size of “clusters” centered on each design point. Data points that lie within tolerance of a design point are included in that cluster. Data points that fall inside the tolerance of more than one design point form a single cluster containing all those design and data points. If no data points fall within tolerance of a particular design point, they remain unmatched. Default tolerance values are related to variable ranges.

The choices you make in the **Default data selection** options determine how the cluster algorithm is first run to select matching data and design points. This only affects the status of the check boxes for data and design points when you first see the Cluster Plot view in the Data Editor. You can always alter the results of this later in the Cluster Information list view, where you can manually select the data and design points you want to use.

When you exit the Data Editor these selections determine what data is used for modeling and how design points are augmented and replaced. Selected

data is used for modeling and added to the design. Data you have decided to exclude is not used for modeling or added to the design.

- **Unmatched data** — Use or Do not use

Data that does not lie within tolerance of any design point is unmatched. You can decide what to do with these. If you select Use, this data is selected for modeling and added to your design.

If you select Do not use, then unmatched data is not used for modeling or added to the design; it is excluded data.

- **Clusters with more data** — Use all data or Use closest match only

This refers to clusters containing more data points than design points. If you choose Use all data, all the data points in these clusters are selected for modeling and added to the design, replacing the design points in those clusters.

If you choose Use closest match only, then a one-to-one match of the data point closest to each design point is selected, and these are the only points that are selected for modeling and added to the design (replacing a design point each).

- **Cluster with more design** — Do not replace design points or Replace design with closest

All data from these clusters is selected for modeling. The setting here only affects selections for the design.

Where clusters contain more design points than data points, you can choose to leave the design unchanged by selecting Do not replace design points.

If you choose Replace design with closest, this replaces the design points where possible with the closest data point and leaves the rest of the design points unchanged.

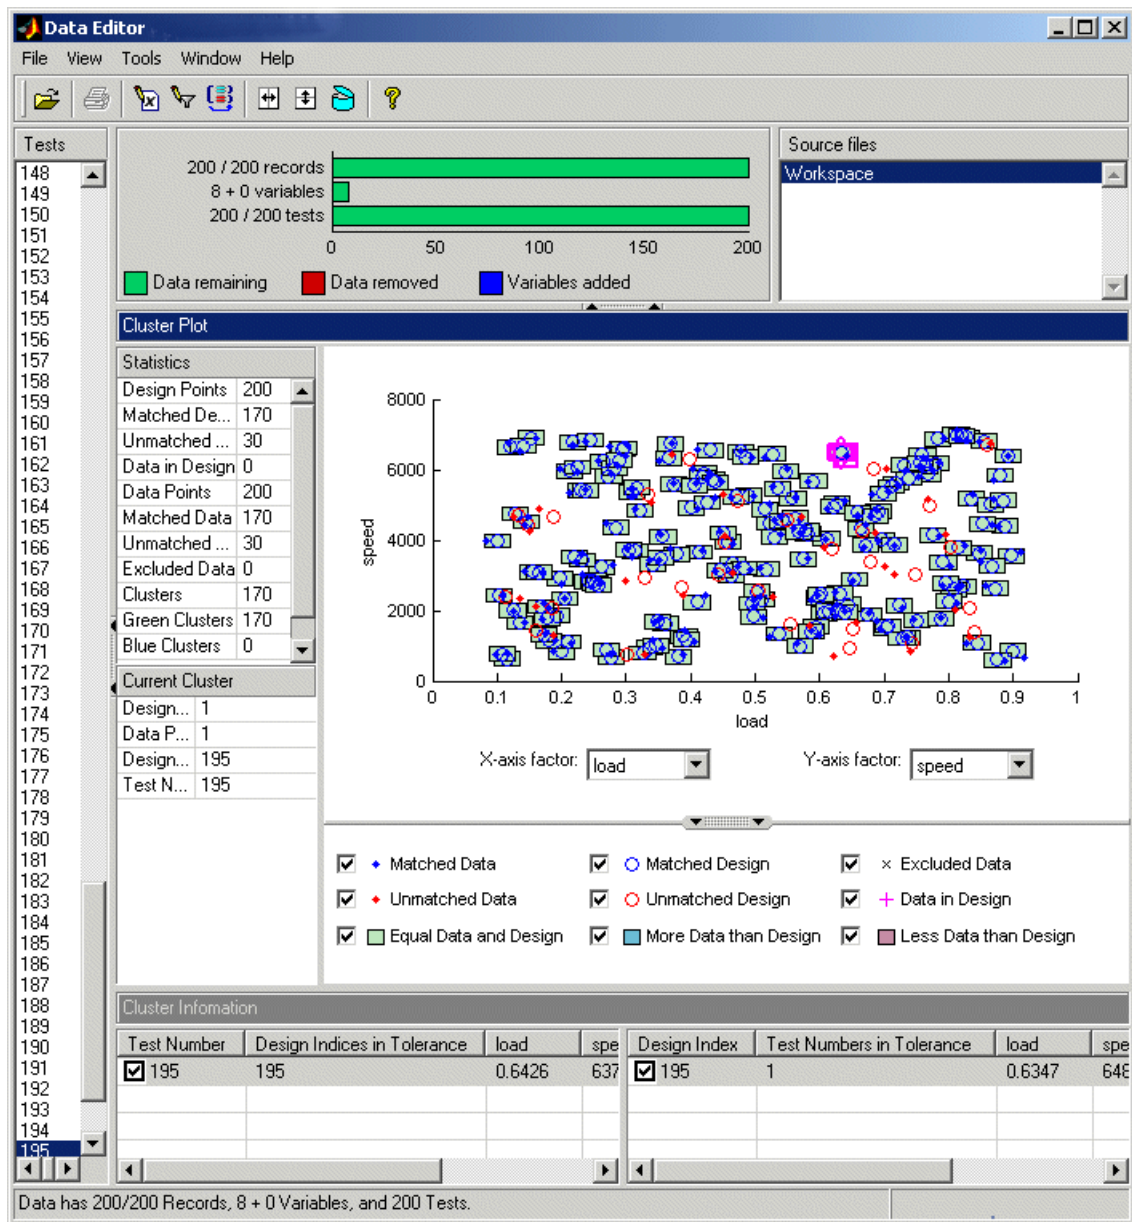
Remember you can override any of these selections manually in the Data Editor; changes are only applied when you close the Data Editor after matching. The cluster plots view in the Data Editor window appears by default when you close the Data Wizard while you are matching data to designs. See the next section, “Matching Data to Designs” on page 3-41.

Matching Data to Designs

The Cluster Plot view is designed for matching data to experimental designs for global models. Here you can select data for modeling. You can use an iterative process: make a design, collect some data, match that data with your design points, modify your design accordingly, then collect more data, and so on. You can use this process to optimize your data collection process in order to obtain the most robust models possible with the minimum amount of data.

Use the Cluster Plot view and associated Cluster Information List view to select data for modeling. All data you select is also added to a new design called Actual Design. You can use the matching process to produce an Actual Design that accurately reflects your current data. You can then use this new design to decide the best points to use if you want to augment your current design in order to collect more data.

The following example shows the Data Editor displaying a Cluster Plot view and Cluster Information List view. You can right-click to change and add views as for the rest of the Data Editor displays.

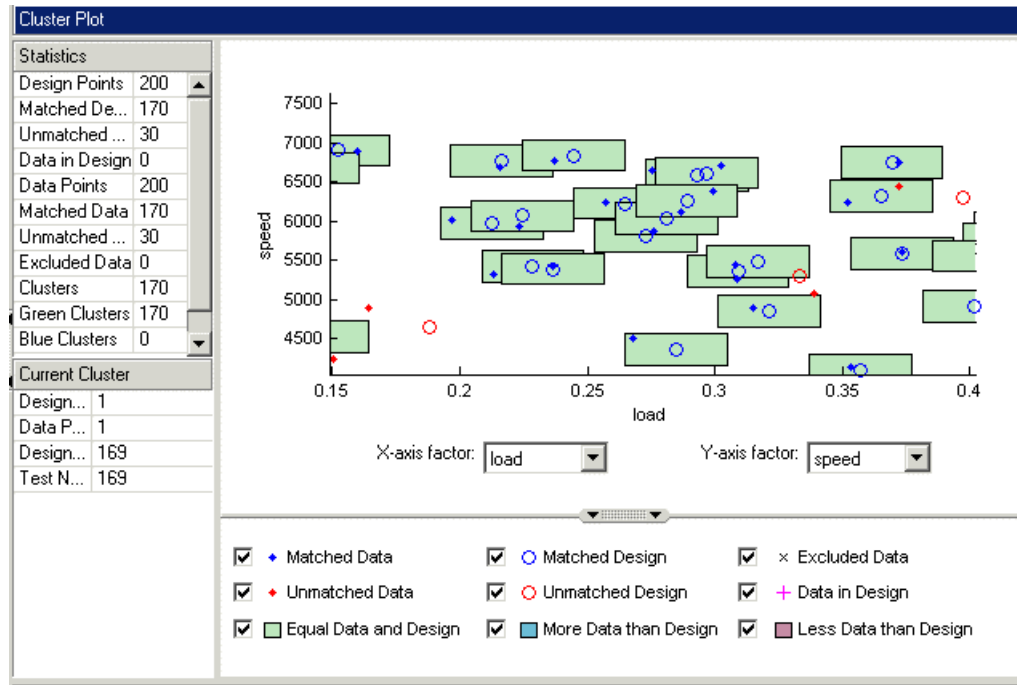


For instructions, see the following section, “How to Use the Cluster Plot View” on page 3-43.

How to Use the Cluster Plot View

For a step-by-step guide to matching data to a design using an example project, see “Matching Data to Designs” on page 3-41 in the Getting Started documentation.

You can **Shift**+click (or center+click) and drag to zoom in on clusters of interest, as shown in the following example. Double-click the plot to return to full size.



We suggest the following sequence as a guideline for matching data to designs using the Cluster Plot:

- 1** It is unlikely that you will get the tolerances right immediately. Open the Tolerance Editor using the **Tools** menu and try different values for different variables. These values determine the size of clusters centered on each design point. Data points that lie within tolerance of any design point in a cluster are matched to that cluster. See “The Tolerance Editor” on page 3-45 for cluster definitions.
- 2** For matching data to designs, you might want to clear the check box in the Cluster Plot for green clusters (with equal data and design points). These clusters are matched; you are more likely to be interested in unmatched points and clusters with uneven numbers of data and design points. Removing the green clusters allows you to focus on these points of interest. If you want your new Actual Design to accurately reflect your current data, your aim is to get as many data points matched to design points as possible, that is, as few red clusters as possible. See “Red Clusters” on page 3-45.
- 3** You can see the values of variables at different points by clicking and holding. Selected points have a pink border. Once points are selected, you can change the plot variables using the X- and Y-axis factor drop-down menus to track those points through the different dimensions.

This can give you a good idea of which tolerances to change in order to match points. Remember that points that do not form a cluster can appear to be perfectly matched when viewed in one pair of dimensions; you must view them in other dimensions to find out where they are separated beyond the tolerance value. Use this tracking process to decide whether you want particular pairs of points to be matched, and then change the tolerances until they form part of a cluster.

Remember that points you select in the cluster view are selected across the Data Editor, so if you have other data plots or a table view open you can investigate the same points in different views.

- 4** Once you have found useful values for the tolerances by trial and error, you can make selections of points within clusters that have uneven numbers of data and design points. These clusters are blue (more data than design) or red (more design than data). Select any cluster by clicking it. The details of every data and design point contained in the selected cluster appear in the Cluster Information List view. Choose the points you want to keep or discard by selecting or clearing the check boxes next to each point in the Cluster Information List view. Notice that your selections can cause

clusters to change color as you adjust the numbers of data and design points within them.

- 5 You can also select unmatched points by right-clicking and selecting **Viewer Options > Select Unmatched Data**. All unmatched points then appear in the Cluster Information List view. You can decide whether to include or exclude them in the same way as points within clusters, by using the check boxes in the Cluster Information List view. If you decide to exclude data points (within clusters or not) they appear on the plot as black crosses (if the **Excluded Data** check box is selected for display).

Note that it is a single fast operation to multiple-select points before selecting or clearing a check box, rather than selecting points individually. To do this, use **Shift**+click to select multiple points and hold the **Shift** key when clicking one of the check boxes.

You can right-click and select **Viewer Options > Show Labels** to see design and data point numbers on the plot (also in the View menu).

Continue this process of altering tolerances and making selections of points until you are satisfied that you have selected all the data you want for modeling. All selected data is also added to your new Actual Design, except that in red clusters.

Red Clusters

These contain more design points than data points. These data points are not added to your design, because the algorithm cannot choose the design points to replace, so you must manually make selections to deal with red clusters if you want to use these data points in your design. If you don't care about the Actual Design (for example, if you do not intend to collect more data) and you are just selecting data for modeling, then you can ignore red clusters. The data points in red clusters are selected for modeling. For information about the effects of your selections, see "What Will Happen to My Data and Design?" on page 3-47.

The Tolerance Editor

You can open the Tolerance Editor by selecting **Tools > Tolerances**.

Here you can edit the tolerance for selecting data points. You can choose values for each variable to determine the size of tolerance in each dimension.

- Data points within the tolerance of a design point are included in that cluster.
- Data points that fall inside the tolerance of more than one design point form a single cluster containing all those design and data points.
- Excluded data (shown as black crosses) that lies within tolerance appears in the Cluster Information list when that cluster is selected. You can then choose whether to use it or continue to exclude it.
- Data in Design (pink crosses) is the only type of data that is not included in clusters.

Note Tolerances are set for global variables. Data used for matching uses test means of global variables, not individual records, unlike other Data Editor views. Click points to inspect values of global variables.

Using the Tolerance Editor is the same process as setting tolerances within the Data Wizard. In the Data Wizard you can also choose in advance what to do with unmatched data and clusters with uneven numbers of data and design points. These choices affect how the cluster algorithm is first run; you can always change selections later in the Data Editor. See “Step 4: Set Tolerances” on page 3-38.

Note If you modify the data in any way while the cluster plot view is open (e.g., by applying a filter) the cluster algorithm will be rerun. You might lose your design point selections.

See the next section, “What Will Happen to My Data and Design?” on page 3-47, for information about what happens to your data set and design when you close the Data Editor after data selection in the cluster plot view.

What Will Happen to My Data and Design?

As with everywhere else in the Data Editor, the changes you make are only applied to the data set when you exit. When you close the Data Editor, your choices in the cluster plots are applied to the data set and a new design called Actual Design is created. All the changes are determined by your check box selections for data and design points.

Note All data points with a selected check box are selected for modeling. All data points with a cleared check box are excluded from the data set.

All data points with a selected check box are put into the new Actual Design *except* those in red clusters. See below.

When you close the Data Editor, these changes are applied:

- Green clusters — equal number of data and design points

The design points are replaced by the equal number of data points. These points become fixed design points (red in the Design Editor table) and appear as Data in Design (pink crosses) when you reopen the Data Editor.

This means that these points are not included in clusters when matching again. These fixed points are also not changed in the Design Editor when you add points, although you can unlock fixed points if you want. This can be very useful if you want to optimally augment a design, taking into account the data you have already obtained.

- Blue clusters — more data than design points

The design points are replaced by all the data points.

Note Design points with selected check boxes in green or blue clusters are the points that will be replaced by your selected data points. You may have cleared the check boxes of other design points in these clusters, and these points will be left unchanged.

- Red Clusters — more design than data points

Red clusters indicate that you should make a decision if you want your new Actual Design to reflect your most current data. The algorithm cannot choose the design points to replace with the data points, so no action is taken. Red clusters do not make any changes to the design when you close the data editor. The existing design points remain in the design. The data points are included or excluded from the data set depending on your selections in the Cluster Information list, but they are not added to the design.

- Unmatched Design Points

These remain in the design.

- Unmatched Data Points

If you have selected the check boxes for unmatched data, they become new fixed design points, which are red in the Design Editor. When you reopen the Data Editor these points are Data in Design, which appear as pink crosses. Note that in the Data Wizard you could choose Use to select all these initially, or you could choose Do not use, which clears all their check boxes. See “Step 4: Set Tolerances” on page 3-38.

- Data in Design

These remain in the design.

- Excluded Data

These data points are removed from the data set and are not displayed in any other views. If you want to return them to the data set you can only do so by selecting them in the Cluster Plot view.

Data Loading Application Programming Interface

Data Loading API Specification

You can use the data loading API (application programming interface) to write your own data loading function, plug these functions into the toolbox, and subsequently use data loaded by these functions within the toolbox. To allow this, there are several stages that need to be followed as described below. For an example, see `xregReadConcerto.m` (in the `mbctools` directory).

Data Function Prototype

A function to successfully load data has the following prototype:

```
[OK, msg, out] = dataLoadingFcn(filename, protoOut)
```

Input Arguments

`filename` is the full path to the file to be loaded.

`protoOut` is an empty structure with the fields expected in the return argument `out`. This allows the data loading API to be easily extended without the need for data loading functions to change when MBC changes.

Output Arguments

The first return argument, `OK`, allows the function to signal that it has successfully loaded the data file. A value of 1 signals success, and 0 failure. If the function fails, it can return a message, `msg`, to indicate the reason for the failure. This message is displayed in a warning dialog box if the function fails. If the function is successful, the return argument `out` contains the data necessary for MBC.

`out.varNames` is a cell array of strings that hold the names of the variables in the data ($1 \times n$ or $n \times 1$).

`out.varUnits` is a cell array of strings that hold the units associated with the variables in `varNames` ($1 \times n$ or $n \times 1$). This array can be empty, in which case no units are defined.

`out.data` is an array that holds the values of the variables ($m \times n$).

`out.comment` is an optional string holding comment information about the data.

Data Function Check In

Once you have written the function, you need to check it into the MBC Toolbox, using the `xregCheckinDataLoadingFcn` function. This function has the following prototype:

```
OK= xregCheckinDataLoadingFcn(fcn, filterSpec, fileType)
```

`fcn` is either a function handle or a string that calls the data loading function via `feval`. This function must be on the MATLAB path.

`filterSpec` is a 1×2 element cell array that contains the extensions that this function loads and the descriptions of those files. This cell array is used in the `uigetfile` function, for example, `{ '*.m;*.fig;*.mat;*.mdl', 'All MATLAB Files' }`. MBC attempts to decide automatically which sort of file is loaded, based on the extension. In the case of duplicate extensions, the first in the list is selected; however, it is always possible to override the automatic selection with a user selection. You will see a warning message if there is any ambiguity.

`fileType` is a string that describes the file type, for example, 'MATLAB file' or 'Excel file'.

Modeling

This section discusses the following topics:

Setting Up Models (p. 4-3)	How to set up all the model types available in the toolbox, including descriptions and illustrations of model shapes.
New Response Models (p. 4-71)	How to add new response models.
Local Level (p. 4-74)	How to examine models at the local level.
Global Level (p. 4-93)	How to examine models at the global level.
Selecting Models (p. 4-111)	How to use the Model Selection window to compare models.
MLE (p. 4-141)	How to use Maximum Likelihood Estimation in modeling.
Response Level (p. 4-145)	How to examine models at the response level.
Model Evaluation Window (p. 4-148)	How to use the Model Evaluation Window to examine model fits and verify models against data.
Boundary Constraint Modeling (p. 4-153)	How to construct and use boundary models.
Stepwise (p. 4-166)	How to use the Stepwise window to improve your models.

Box-Cox Transformation (p. 4-176)	How to use Box-Cox transformation and decide if this improves your models.
Maximum Likelihood Estimation for Nonlinear Repeated Measures (p. 4-180)	This section contains an overview of the mathematics of Maximum Likelihood Estimation two-stage models and links to model definitions.
Two-Stage Models for Engines (p. 4-181)	An explanation of the reasons why using two-stage models is an effective way to model data containing repeated measures.
Linear Regression (p. 4-191)	Introducing linear regression in the Model-Based Calibration Toolbox.
Toolbox Terms and Statistics Definitions (p. 4-192)	A table defining terms and statistics used in the toolbox.
Exporting Models (p. 4-195)	How to export models.

Setting Up Models

The following sections cover setting up inputs and models in detail:

- “Setting Up Inputs” on page 4-4
- “Global Model Setup” on page 4-6
 - “Global Linear Models: Polynomials and Hybrid Splines” on page 4-8
 - “Global Model Class: Radial Basis Function” on page 4-14
 - “Global Model Class: Hybrid RBF” on page 4-20
 - “Global Model Class: Multiple Linear Models” on page 4-21
 - “Global Model Class: Free Knot Spline” on page 4-22
 - “Global Model Class: Neural Network” on page 4-24
 - “Global Model Class: User-Defined and Transient Models” on page 4-26
- “Local Model Setup” on page 4-26
 - “Local Model Class: Polynomials and Polynomial Splines” on page 4-28
 - “Local Model Class: Truncated Power Series” on page 4-32
 - “Local Model Class: Free Knot Spline” on page 4-34
 - “Local Model Class: Growth Models” on page 4-36
 - “Local Model Class: Linear Models” on page 4-42
 - “Local Model Class: Average Fit” on page 4-44
 - “Local Model Class: Multiple Models” on page 4-45
 - “Local Model Class: User-Defined Models” on page 4-49
 - “Local Model Class: Transient Models” on page 4-55
 - “Covariance Modeling” on page 4-66
 - “Correlation Models” on page 4-67
 - “Transforms” on page 4-68
- “Designing Experiments” on page 4-69 tells you how to proceed to experimental design after setting up models.

What Models Are Available?

*Multiple Models gives you access to global model types for your local model.

** Polynomial and polynomial spline are two special case linear models for local models with one input factor. You can use polynomial and polynomial spline models (with more settings) for local models with more factors by choosing Linear Models.

User defined and transient models must be checked into the toolbox before you can use them. They will be available as local models for the number of factors specified. There is an example user-defined model for a single factor preregistered with the toolbox. The example transient model provided must have exactly two factors, one of which must be time. See “Local Model Class: User-Defined Models” on page 4-49 and “Local Model Class: Transient Models” on page 4-55 for details.

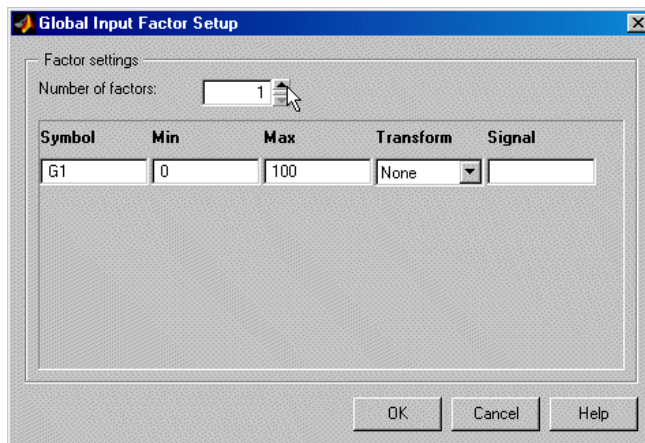
Setting Up Inputs

You set up the number and definition of model input factors for each stage by double-clicking or right-clicking the inports of the block diagram.

You can create a block diagram by clicking **New** in the **Test Plans** pane at project level (the view where you begin). See “Creating a New Test Plan” on page 1-25. Then double-click the blocks labeled inputs.

See “Test Plan Level” on page 1-29.

The following example shows the input setup dialog for the global model. The dialog for the local model contains exactly the same controls.



Number of Factors

You can change the number of input factors using the buttons at the top.

Symbol

The input symbol is used as a shortened version of the signal name throughout the application. The symbol should contain a maximum of three characters.

Min and Max Model Range

This setting is important before you design experiments. The default range is [0.100]. There is usually some knowledge about realistic ranges for variables. If you are not designing an experiment you can use the data range as the model range later, in the data selection stage. In some cases you might not want to use the data range (for example, if the data covers too wide a range, or not wide enough) if you are interested in modeling a particular region. In that case you can set the range of interest here.

Transform

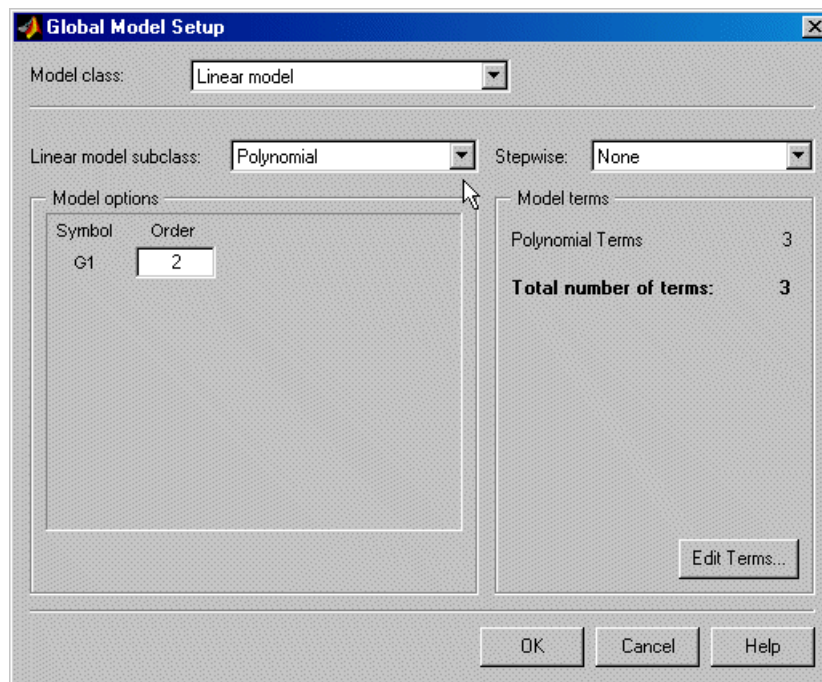
You can use input transformations to change the input factor for designing experiments. The available input transformations are $1/x$, \sqrt{x} , $\log_{10}(x)$, x^2 , $\log(x)$.

Signal


You can set up the signal name in the input factor setup dialog. It is not necessary to set this parameter at this stage, as it can be defined later at the data selection stage (as with the range). However, setting the signal name in this dialog simplifies the data selection procedures, as the Model Browser looks for matching signal names in loaded data sets. When the number of data set variables is large this can save time.

Global Model Setup

The following example shows the default settings for the global model.




Global Model Classes



The default global model is linear polynomial. There are a number of different global model classes available, covered in the following sections:

- “Global Linear Models: Polynomials and Hybrid Splines” on page 4-8
- “Global Model Class: Radial Basis Function” on page 4-14
- “Global Model Class: Hybrid RBF” on page 4-20
- “Global Model Class: Multiple Linear Models” on page 4-21
- “Global Model Class: Free Knot Spline” on page 4-22

Note that some other models are also available for one-stage models:



Of these, neural network models are only available for one-stage models. The single preregistered user-defined model is only available for a single input. You can check in your own user-defined models and transient models with as many factors as you choose; these only appear as options when the appropriate number of inputs are present. User-defined and transient models can also be used for one-stage models.

- “Global Model Class: Neural Network” on page 4-24
- “Global Model Class: User-Defined and Transient Models” on page 4-26

“Global Level” on page 4-93 has a description of the statistics displayed when a global model is analyzed.

See “Setting Up Models” on page 4-3 for a list of all information about setting up local and global models.

Global Linear Models: Polynomials and Hybrid Splines

Global linear models can be

- Polynomial
- Hybrid Spline

Polynomials

Polynomials of order n are of the form

$$\beta_0 + \beta_1x + \beta_2x^2 + \beta_3x^3 \dots \beta_nx^n$$

You can choose any order you like for each input factor.



Quadratic curve




Cubic curve

As shown, a quadratic polynomial $y = ax^2 + bx + c$ can have a single turning point, and a cubic curve $y = ax^3 + bx^2 + cx + d$ can have two. As the order of a polynomial increases, it is possible to fit more and more turning points. The curves produced can have up to $(n-1)$ bends for polynomials of order n .

Term Editor

Click the **Edit Terms** button to see the terms in the model. This opens the Term Editor dialog. Here you can remove any of the terms as shown in the following example.




See also “Interaction” on page 4-12.

Hybrid Splines

You can use the Hybrid Spline model to fit a spline to one factor and polynomials to all other factors.


A spline is a piecewise polynomial function, where different sections of polynomials are fitted smoothly together. The locations of the breaks are called knots. You can choose the required number of knots (up to a maximum of 50) and their positions. In this case all the pieces of curves between the knots are formed from polynomials of the same order. You can choose the order (up to 3).

The following example illustrates the shape of a spline curve with one knot and third-order basis functions. The knot position is marked on the N axis.



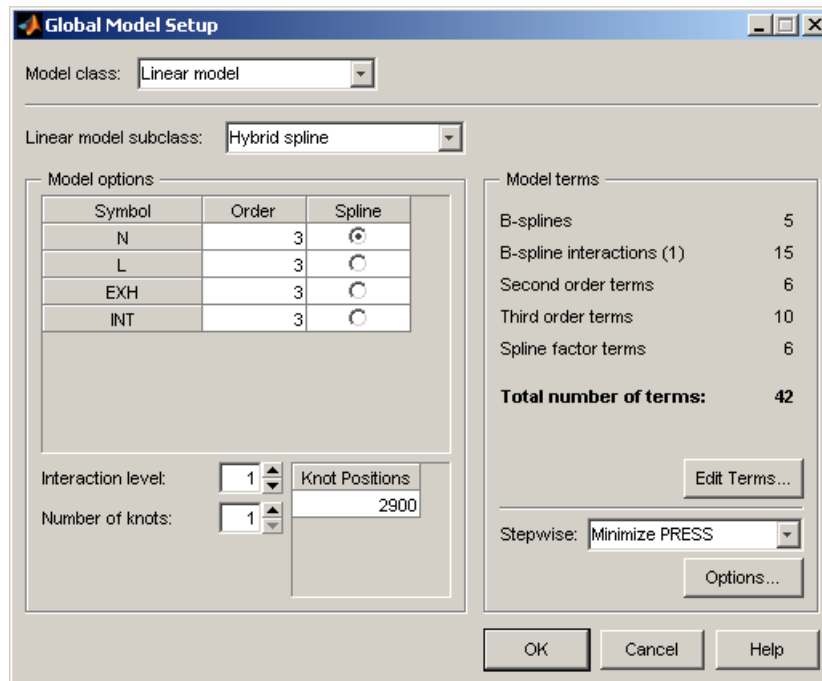
You can fit more complicated curves using splines, so they can be useful for the factor you expect to behave in the most complex way. This allows you to model detailed fluctuations in the response for one factor, while simpler models are sufficient to describe the other factors.

The following example clearly shows that the response (Blow_2 in this case) is quadratic in the Load (L) axis and much more complex in the RPM (N) axis.



You can choose the order of the polynomial for each factor and the factor to fit the spline to. The maximum order for each factor is cubic.

The following example shows the options available for the Hybrid Spline linear model subclass.



This example only shows one input factor. When there are more, you use the radio buttons to select which factor is modeled with a spline. Select the order for each factor in the edit boxes.

Interaction

You can choose the interaction level on both linear model subclasses (polynomial and hybrid spline). The maximum interaction level is the same as the polynomial order (for example, 3 for cubics).

The interaction level determines the allowed order of cross-terms included.

You can use the Term Editor to see the effects of changing the interaction level. Click the **Edit Terms** button. The number of constant, linear, second- and third-order (and above) terms can be seen in the **Model Terms** frame.

For polynomials. With an interaction level of 1, there are no terms in the model involving more than one factor. For example, for a four-factor cubic, for factor L, you see the terms for L, L^2 , and L^3 , but no terms involving L *and* other factors. In other words, there are no cross-terms included.


If you increase the interaction level to 2, under second-order terms you see L^2 and also L multiplied by each of the other factors; that is, second-order cross-terms (for example, LN, LA, and LS).

Increase the interaction to 3, and under third-order terms you see L^3 multiplied by each of the other factors (L^2N , L^2A , L^2S), L multiplied by other pairs of factors (LNA, LNS, LAS), and L multiplied by each of the other factors squared (N^2L , A^2L , S^2L). Interaction level 3 includes all third-order cross-terms.

The preceding also applies to all four factors in the model, not just L.

For hybrid splines. The interaction function is different: it refers only to cross-term interactions between the spline term and the other variables. For example, at interaction order 0, raw spline terms only; interaction 1, raw terms and the spline terms x the first-order terms; interaction 2, includes spline terms x the second-order terms; and so on.

Stepwise



Take care not to overfit the data; that is, you do not want to use unnecessarily complex models that “chase points” in an attempt to model random effects.

The **Stepwise** feature can help. Stepwise will select the terms using various criteria. Stepwise generally means that terms are removed in steps (one at a time). The stepwise algorithms are Minimize Press, Forward Selection, Backwards Selection, and Prune. The most commonly used stepwise

algorithm is Minimize PRESS, where at each step the term that will improve the PRESS statistic the most is removed or included. Minimize PRESS throws away terms in the model to improve its predictive quality, removing those terms that reduce the PRESS of the model. Forward and Backwards Selection uses statistical significance at the alpha % level.

Predicted sum of squares error (PRESS) is a measure of the predictive quality of a model. See “Two-Stage Models for Engines” on page 4-181 for an explanation of PRESS and why it is useful as a diagnostic statistic.

Prune is one of the alternative algorithms for stepwise. The order of the terms matter, and the terms will only be removed from the end, provided they improve the quality of the fit (measured by various criteria: PRESS, GCV etc.). The other stepwise algorithms do not have this restriction - they can remove any term in any order. Removing terms only from the end is valid when there is ordering in the terms e.g. polynomials (from low order terms to high order terms) or rbfs where the fitting algorithms select the most important terms first.


Click **Options** to open a dialog containing further settings for the selected **Stepwise** option. You should choose from the list a criteria for removing terms (PRESS, RMSE, AIC, BIC etc.). For the Prune settings, see “Global Model Class: Radial Basis Function” on page 4-14. For a guide to all the settings in the Stepwise window (which explains the other Stepwise settings available here), see “Stepwise” on page 4-166. Note you can also use the Stepwise window after model fitting to try other Stepwise settings, and replace excluded model terms if you want.

Global Model Class: Radial Basis Function

A variety of radial basis functions (RBFs) are available in MBC. They are all radially symmetrical functions that can be visualized as mapping a flexible surface across a selection of hills or bowls, which can be circular or elliptical.

Networks of RBFs can model a wide variety of surfaces. You can optimize on the number of centers and their position, height and width. You can have different widths of centers in different factors. RBFs can be useful for investigating the shapes of surfaces when system knowledge is low. Combining several RBFs allows complicated surfaces to be modeled with relatively few parameters.

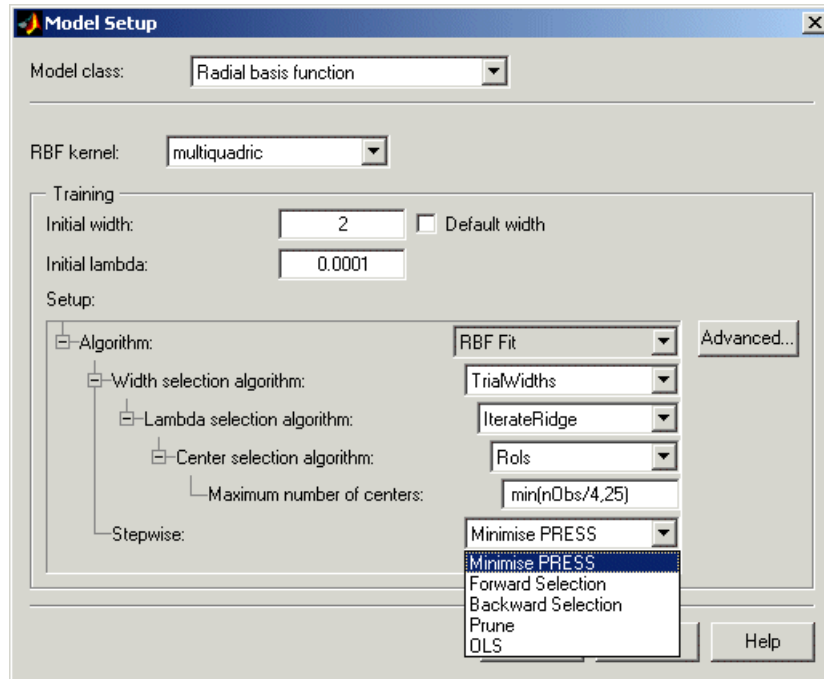
The following example shows a surface of an RBF model.



There is a detailed user guide for modeling using RBFs in . See especially “Tips for Modeling with Radial Basis Functions” on page 5-34 for guidelines and “Types of Radial Basis Functions” on page 5-5 for illustrations of different types.

The statistical basis for each setting in the RBF global models is explained in detail in “Guide to Radial Basis Functions for Model Building” on page 5-3.

The following example illustrates the basic RBF settings available.



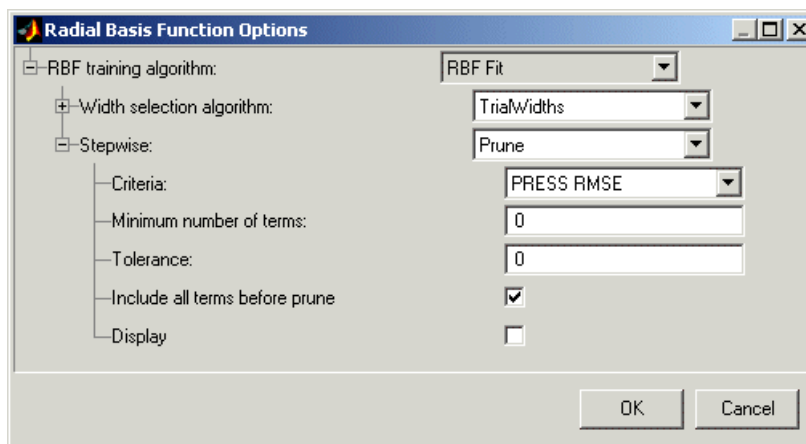
You can use the drop-down menus to set RBF kernel type, initial width and lambda, width, lambda, and center selection algorithm and maximum number of centers. After you have fitted a model once to get some idea of what to expect, you can try different maximum numbers of centers as a useful method for homing in on a good fit. There are more options for fine tuning in the Advanced options dialog, but you can use the main controls from here to narrow down the search for the best model. See below.

For most algorithms the **Initial width** is a single value. However, for WidPerDim (available in the **Width selection algorithm** pull down), you can specify a vector of widths to use as starting widths. WidPerDim produces elliptical basis functions that can have a different width in each factor direction. If supplying a vector of widths, there should be the same number as the number of global variables, and they must be in the same order as specified in the test plan. If you provide a single width, then all dimensions

start off from the same initial width, but are likely to move from there to a vector of widths during model fitting.

You can use the last drop-down menu to choose to automatically run **Stepwise** at the end of the center/width selection algorithm to remove less useful model terms, Ordinary Least Squares (OLS) or the Prune algorithm to home in on the best number of centers (using your choice of the Summary Statistics as selection criteria).

Note if you choose Prune here, there are further settings you need which can be found by clicking **Advanced**. This opens the Radial Basis Function Options dialog, as shown.




All the settings under Width selection algorithm are for fine tuning the RBF model. See “Tips for Modeling with Radial Basis Functions” on page 5-34 for guidelines, and details on specific parameters for different algorithms can be found in Chapter 5, “Radial Basis Functions”.

The options in the **Stepwise** drop-down menu are the same as the main Model Setup dialog — Min. PRESS, Forward, Backward, Prune and OLS (Ordinary Least Squares). If you choose Prune the further options appear below as shown in this example. You need to choose one of the Summary Statistics as selection criteria for the Prune algorithm. All the Summary Statistics options are available as criteria, and do not depend on your choices

of these statistics in the Summary Statistics dialog. See “Summary Statistics” on page 1-32 for more information.

We recommend you select the check box to **Include all terms before prune** (otherwise the current number of terms is used at the start). You can choose a **Minimum number of terms**, and the **Tolerance** you set determines how far from this number of terms the algorithm can go — within the limits of the tolerance, the algorithm searches for fewer terms that reduce the value of your selection criteria.

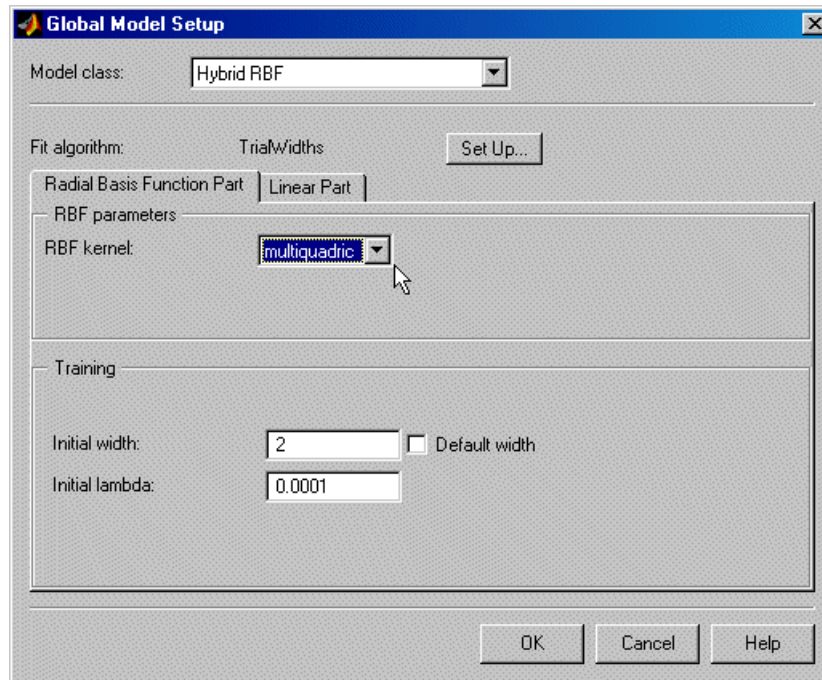
If you select the **Display** check box a figure appears illustrating the Prune process, like the example shown following, plotting the number of parameters against the selection criteria, in this case PRESS RMSE. You can use this information to see if you should change the minimum number of terms and the tolerance and refit, in case you need to avoid a local minimum.



Note Once you have a global model you can use the RBF template in the “Build Models Dialog” on page 4-106 to automatically build a variety of radial basis function models with varying maximum numbers of centers and/or different kernels.

Global Model Class: Hybrid RBF

This option combines an RBF model with a linear model.



The **RBF kernel** drop-down menu offers the same options as for normal RBF.

The **Linear Part** tab contains the same options as the other global linear models; see “Global Linear Models: Polynomials and Hybrid Splines” on page 4-8.


See “Hybrid Radial Basis Functions” on page 5-32.

Click **Set Up** to reach the Hybrid RBF Options dialog where you can change all the settings for the RBF part of the model. Here you can also choose to automatically run Stepwise, OLS or Prune. These settings are the same as the Radial Basis Functions Options dialog, see “Global Model Class: Radial Basis Function” on page 4-14 for details.

See also Chapter 5, “Radial Basis Functions” for a detailed guide to the use of all the available RBFs in modeling.

Global Model Class: Multiple Linear Models

The following example shows the defaults for multiple linear models. You can add linear models (as seen in the single linear model settings).



This is primarily for designing experiments using optimal designs. If you have no idea what model you are going to fit, you would choose a space-filling design. However, if you have some idea what to expect, but are not sure exactly which model to use, you can specify a number of possible models here. The Design Editor can average optimality across each model.

For example, if you expect a quadratic and cubic for two factors but are unsure about a third, you can enter several alternative polynomials here. You can change the weighting of each model as you want (for example, 0.5 each for two

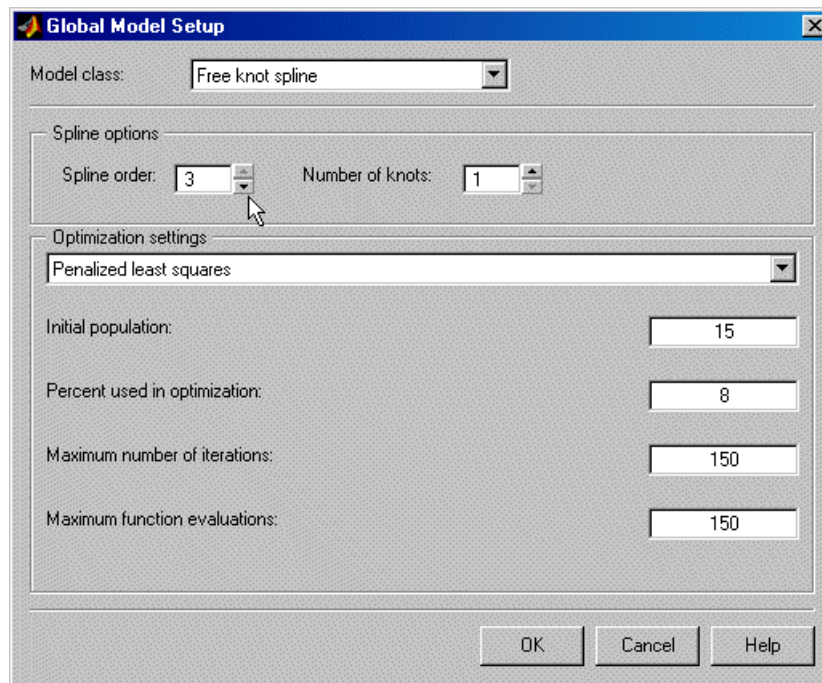
models you think equally likely). This weighting is then taken into account in the optimization process in the Design Editor.

The model that appears in the model tree is the one you select, listed as **Primary model**. Click the model in the list, then click **Use Selected**. The **Primary model** changes to the desired model. If you do not select a primary model, the default is the first in the list.

When the model has been fitted, you can view the primary model at the global node. To compare the fit of all the alternatives, click **Build Models** in the toolbar, select **Multiple Linear Models** in the dialog, and click **OK**. One of each model is then built as a selection of child nodes.

See also “Polynomials” on page 4-8, and “Local Model Class: Multiple Models” on page 4-45.

Global Model Class: Free Knot Spline




This option is only available for global models with only one input factor. See also “Hybrid Splines” on page 4-9 for a description of splines. The major difference is that you choose the position of the knots for hybrid splines; here the optimal knot positions are calculated as part of the fitting routine.

You can set the number of knots and the spline order can be between one and three.

There are three different algorithms under **Optimization settings**: Penalized least squares, Genetic algorithm, and Constrained least squares.

For all three methods, you can set the **Initial population**. This is the number of initial guesses at the knot positions. The other settings put limits on how long the optimization takes.

The following example shows a free knot spline model with three knots. The position of the knots is marked on the N axis.



See also the local models involving splines:

- “Polynomial Spline” on page 4-29
- “Local Model Class: Truncated Power Series” on page 4-32
- “Local Model Class: Free Knot Spline” on page 4-34

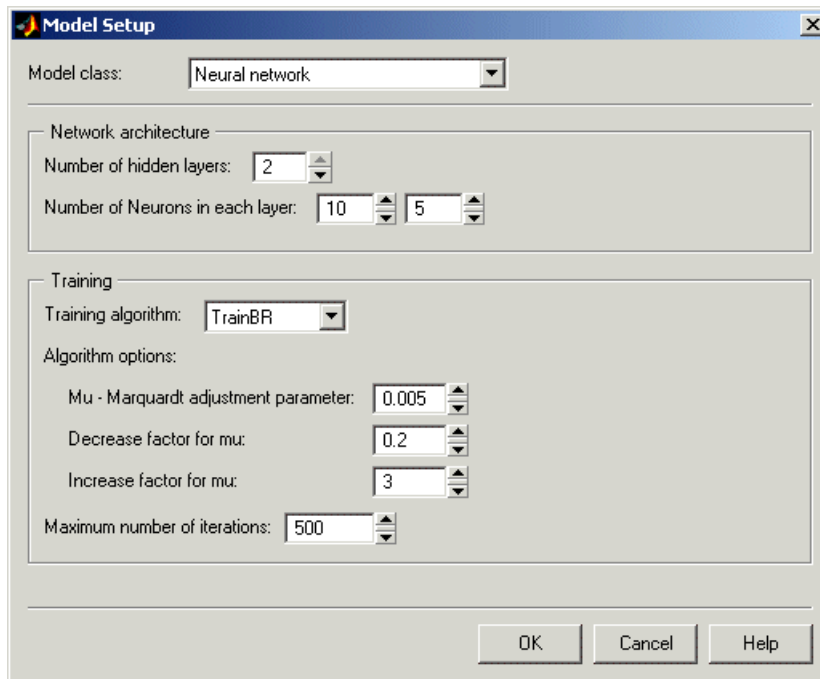
Global Model Class: Neural Network

Neural networks are only available for one-stage models.

Neural nets contain no preconceptions of what the model shape will be, so they are ideal for cases with low system knowledge. They are useful for functional prediction and system modeling where the physical processes are not understood or are highly complex.

The disadvantage of neural nets is that they require a lot of data to give good confidence in the results, so they are not suitable for small data sets. Also, with higher numbers of inputs, the number of connections and hence the complexity increase rapidly.

MBC provides an interface to some of the neural network capability of the Neural Network Toolbox. Therefore these functions are only available if the Neural Network Toolbox is installed.



For help on the neural net models implemented in the Model-Based Calibration Toolbox, see the documentation in the Neural Nets Toolbox . At the MATLAB command line, enter

```
>>doc nnet
```

The training algorithms available in the Model-Based Calibration Toolbox are `traingdm`, `trainlm`, `trainbr`.

These algorithms are a subset of the ones available in the neural nets toolbox. (The names indicate the type: gradient with momentum, named after the two authors, and bayesian reduction). Neural networks are inspired by biology, and attempt to emulate learning processes in the brain.

Neural nets contain no preconceptions of what the model shape will be, so they are ideal for cases with low system knowledge. They are useful for

functional prediction and system modeling where the physical processes are not understood or are highly complex.

The disadvantage of neural nets is that they require a lot of data to give good confidence in the results, so they are not suitable for small data sets. Also, with higher numbers of inputs, the number of connections and hence the complexity increase rapidly.

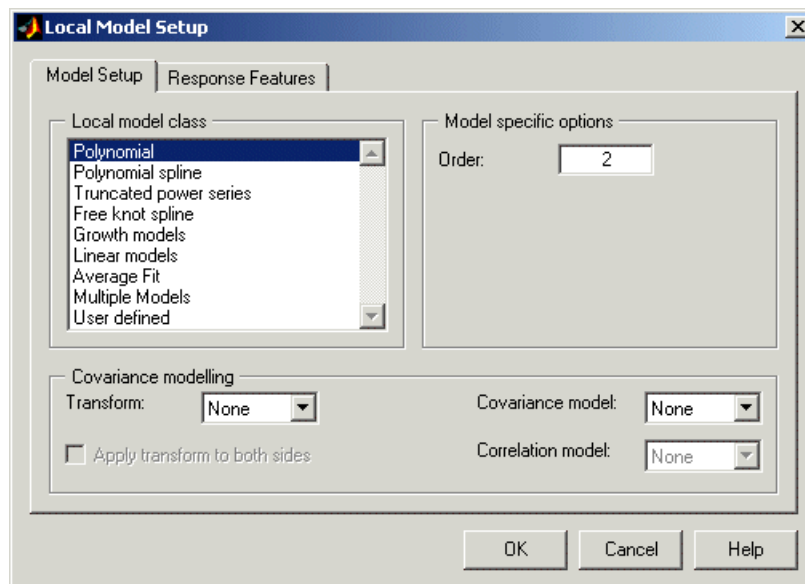
See the Neural Network Toolbox documentation for more help.

Global Model Class: User-Defined and Transient Models

These models can be one-stage models, or local model types for two-stage modeling. For set up information see “Local Model Class: User-Defined Models” on page 4-49 and “Local Model Class: Transient Models” on page 4-55.

Local Model Setup

The following example shows the default settings for the local model.



Local models can be of the following types:

- Polynomial
- Polynomial Spline
- Truncated Power Series
- Free Knot Spline
- Growth Models
- Linear Models
- Average Fit
- Multiple Models
- User-Defined Models

You can choose additional response features at this stage using the **Response Features** tab. These can also be added later. The Model Browser automatically chooses sufficient response features for the current model.

See also

- “Covariance Modeling” on page 4-66
- “Correlation Models” on page 4-67
- “Transforms” on page 4-68

See each local model type for statistical details and available response features.

See “Setting Up Models” on page 4-3 for a list of all information about setting up local and global models.

Local Model Class: Polynomials and Polynomial Splines

Polynomials

At the local level, if you have one input factor, you can choose Polynomial directly from the list of local model classes. Here you can choose the order of polynomials used, and you can define a datum model for this kind of local model (see below).

If there is more than one input factor, you can only choose Linear Models, Average Fit, Multiple Models or Transient Models from the Local Model Class list.

Under Linear Models you can choose Polynomial or Hybrid Spline. This is a different polynomial model where you can change more settings such as Stepwise, the Term Editor (where you can remove any model terms) and you can choose different orders for different factors (as with the global level polynomial models). See “Local Model Class: Linear Models” on page 4-42.

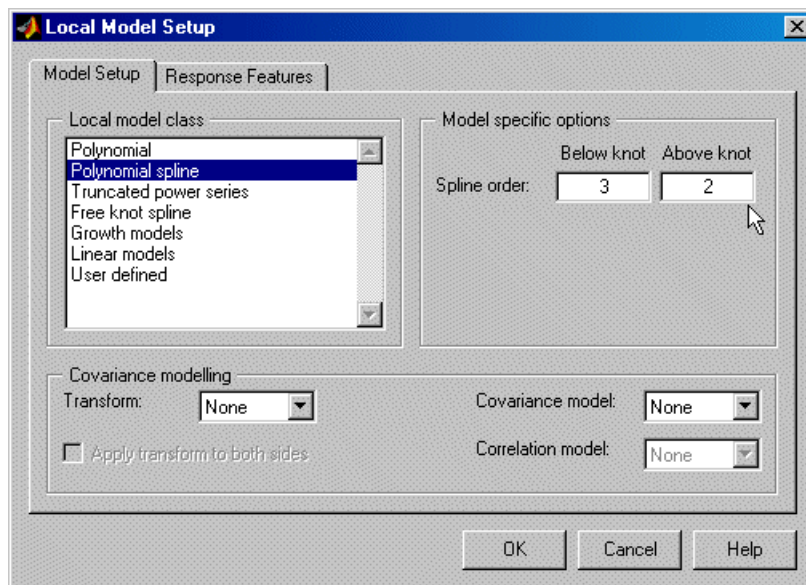
Different response features are available for this polynomial model and the Linear Models: Polynomial choice. You can view these by clicking the **Response Features** tab on the Local Model Setup dialog. Single input polynomials can have a datum model, and you can define response features relative to the datum. See “Datum Models” on page 4-72.

The following response features are permitted for the polynomial model class:

- Location of the maximum or minimum value (when using datum models; note that the datum model is not used in reconstructing).
- Value of the fit function at a user-specified x-ordinate. When datum models are used, the value is relative to the datum (for example, $\text{mbt} - x$).
- The n^{th} derivative at a user-specified x-ordinate, for $n = 1, 2, \dots, d$ where d is the degree of the polynomial.

See the global model section “Polynomials” on page 4-8 for a general description of polynomial models.


Polynomial Spline



A spline is a piecewise polynomial, where different sections of polynomial are fitted smoothly together. The location of each break is called a knot. Polynomial splines are essential for modeling torque/spark curves.

This model has only one knot. You can choose the orders of the polynomials above and below the knot. See also “Hybrid Splines” on page 4-9. These global models also use splines, but use the same order polynomial throughout.

Polynomial splines are only available for single input factors. The following example shows a typical torque/spark curve, which requires a spline to fit properly. The knot is shown as a red spot at the maximum, and the curvature above and below the knot is different. In this case there is a cubic basis function below the knot and a quadratic above.



To model responses that are characterized in appearance by a single and well defined stationary point with asymmetrical curvature either side of the local minimum or maximum, we define the following spline class,

$$y_{ij} = \beta_o + \sum_{a=2}^c \beta_{Low_a} (x_j - k)_-^a + \sum_{b=2}^h \beta_{High_b} (x_j - k)_+^b$$

where k is the knot location, β denotes a regression coefficient,

$$(x_j - k)_- = \min\{0, (x_j - k)\}, (x_j - k)_+ = \max\{0, (x_j - k)\}$$

where c is the user-specified degree for the left polynomial, h is the user-specified degree for the right polynomial, and the subscripts Low and High denote to the left (below) and right of (above) the knot, respectively.

Note that by excluding terms in $(x_j - k)_-$ and $(x_j - k)_+$ we ensure that the first derivative at the knot position is continuous. In addition, by definition the constant β_0 must be equal to the value of the fit function at the knot, that is, the value at the stationary point.

For this model class, response features can be chosen as

- Fit constants $\{\beta_0, \beta_{Low_2}, \dots, \beta_{Low_p}, \beta_{High_2}, \beta_{High_q}\}$
- Knot position $\{k\}$
- Value of the fit function at a user-specified delta

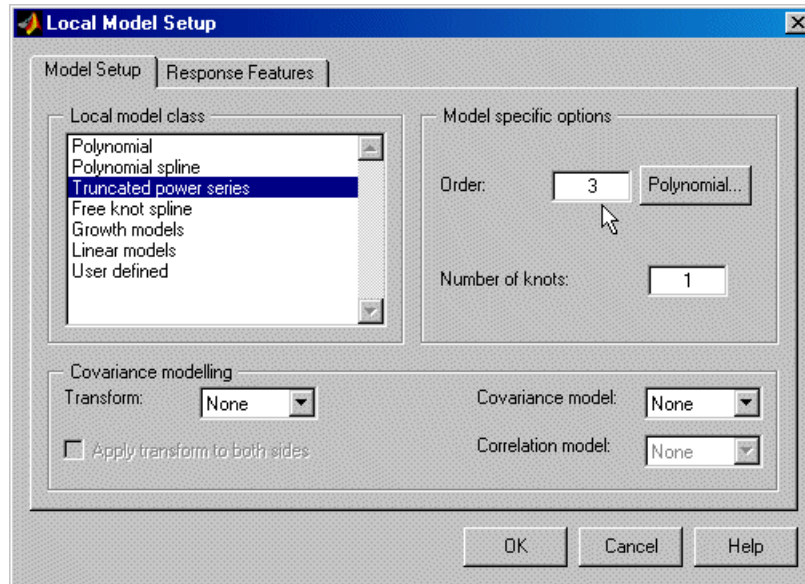
$$\{\Delta a_j = x_j - k\}$$

from the knot position $\{f(\pm \Delta a_j)\}$ if the datum is defined, otherwise the value is absolute.

- Difference between the value of the fit function at a user-specified delta from the knot position and the value of the fit function at the knot

$$\{f(\pm \Delta a_j) - f(k)\}$$

Local Model Class: Truncated Power Series



This is only available for a single input factor.

You can choose the order of the polynomial basis function and the number of knots for Truncated Power Series Basis Splines. A spline is a piecewise polynomial, where different sections of polynomial are fitted smoothly together. The point of each break is called a knot. The truncated power series changes the coefficient for the highest power when the input passes above the knot value.

It is *truncated* because the power series is an approximation to an infinite sum of polynomial terms. You can use infinite sums to approximate arbitrary functions, but clearly it is not feasible to fit all the coefficients.

Click **Polynomial** to see (and remove, if you want) the polynomial terms. One use of the remove polynomial term function is to make the function linear until the knot, and then quadratic above the knot. In this case we remove the quadratic coefficient.

See also

- “Polynomial Spline” on page 4-29, where you can choose different order basis functions either side of the knot
- “Hybrid Splines” on page 4-9, a global model where you can choose a spline order for one of the factors (the rest have polynomials fitted)
- “Local Model Class: Free Knot Spline” on page 4-34, free knot splines, where you can choose the number of knots and the order of the basis functions

Truncated Power Series Basis (TPSBS) Splines

A very general class of spline functions with arbitrary (but strictly increasing) knot sequence:

$$\mathbf{k} = \{k_1, k_2, \dots, k_k\}^T ;$$

$$f(x) = \sum_{i=0}^m \beta_i x^i + \sum_{i=0}^k \beta_{m+i} (x - k_i)_+^m$$

This defines a spline of order m with knot sequence $\mathbf{k} = \{k_1, k_2, \dots, k_k\}^T$

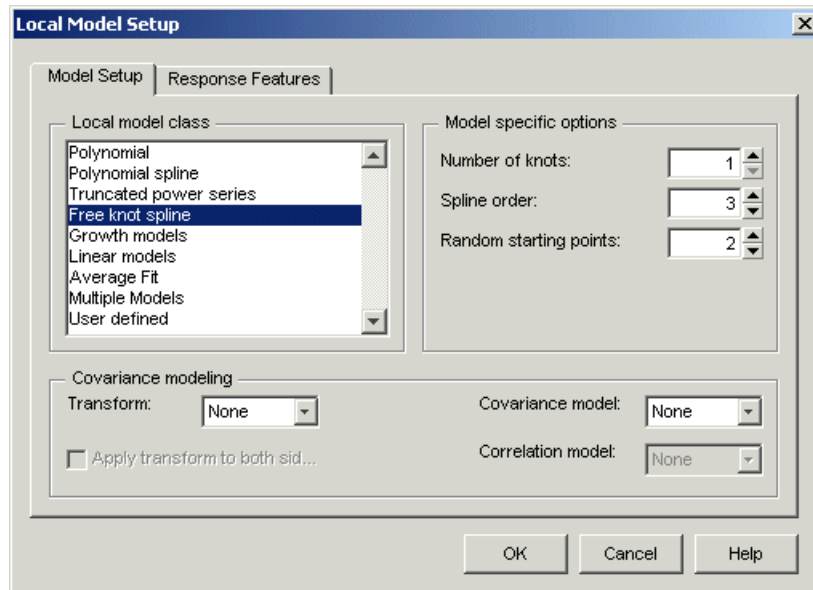
For this model class, response features can be chosen as

- Fit constants $\{\beta_0, \beta_1, \dots, \beta_{m-1+k}\}$
- Knot position vector $\{\mathbf{k}\}$.
- Value of the fit function $\left\{f(a_j)\right\}$ at a user-specified value $\{a_j\}$
- Value of the n^{th} derivative of the fit function with respect to x_j $\left\{f(a_j)^n\right\}$ at a user-specified value $\{a_j\}$, with $n = 1, 2, \dots, m-2$

Any of the polynomial terms can be removed from the model.

Local Model Class: Free Knot Spline

These are the same as the “Global Model Class: Free Knot Spline” on page 4-22 (which is also only available for one input factor). See the global free knot splines for an example curve shape.



A spline is a piecewise polynomial, where different sections of polynomial are fitted smoothly together. The point of the join is called the knot.

You can choose the number of knots. You can choose the order of polynomial fitted (in all curve sections) from 1 to 3. The default is cubic.

You can set the number of **Random starting points**. These are the number of initial guesses at the knot positions.

See also

- “Polynomial Spline” on page 4-29, where you can choose different order basis functions for either side of the knot
- “Local Model Class: Truncated Power Series” on page 4-32, where you can choose the order of the basis function
- “Hybrid Splines” on page 4-9, a global model where you can choose a spline order for one of the factors (the rest have polynomials fitted)

Free Knot Splines

The $(x_j - k_i)_+^{m-1}$ basis is not the best suited for the purposes of estimation and evaluation, as the design matrix might be poorly conditioned. In addition, the number of arithmetic operations required to evaluate the spline function depends on the location of x_i relative to the knots. These properties can lead to numeric inaccuracies, especially when the number of knots is large. You can reduce such problems by employing *B-splines*.

The most important aspect is that for moderate m the design matrix expressed in terms of B-splines is relatively well conditioned.

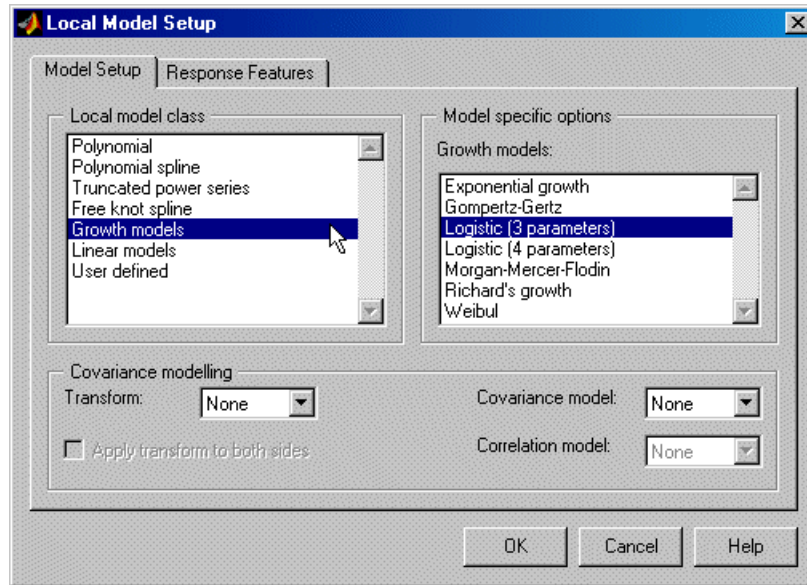
For this model class, response features can be chosen as

- Fit constants $\{\beta_{-(m-1)}, \beta_{-m}, \dots, \beta_k\}$
- Knot position vector $\{k\}$.


Value of the fit function $\{f(a_j)\}$ at a user specified value $\{a_j\}$

Local Model Class: Growth Models

Growth models have a number of varieties available, as shown. They are only available for single input factors.



These are all varieties of sigmoidal curves between two asymptotes, like the following example.



Growth models are often the most appropriate curve shape for air charge engine modeling.

See the following sections for mathematical details on the differences between these growth models:

- “Three Parameter Logistic Model” on page 4-37
- “Morgan-Mercer-Flodin Model” on page 4-38
- “Four-Parameter Logistic Curve” on page 4-39
- “Richards Curves” on page 4-39
- “Weibul Growth Curve” on page 4-40
- “Exponential Growth Curve” on page 4-41
- “Gompertz Growth Model” on page 4-41

Three Parameter Logistic Model

The three parameter logistic curve is defined by the equation

$$y_j = \frac{\alpha}{1 + \exp(-\kappa(x_j - \gamma))}$$

where α is the final size achieved, K is a scale parameter, and γ is the x-ordinate of the point of inflection of the curve.

The curve has asymptotes $y_j = 0$ as $x_j \rightarrow -\infty$ and $y_j = \alpha$ as $x_j \rightarrow \infty$. Growth rate is at a maximum when $y_j = \alpha/2$, which occurs when $x_j = \gamma$. Maximum growth rate corresponds to

$$\frac{\kappa\alpha}{4}$$

The following constraints apply to the fit coefficients:

$$\alpha > 0, K > 0, \gamma > 0.$$

The response feature vector g for the 3 parameter logistic function is defined as

$$g = \left[\alpha \quad \gamma \quad \kappa \quad \frac{\kappa\alpha}{4} \right]^T$$

Morgan-Mercer-Flodin Model

The Morgan-Mercer-Flodin (MMF) growth model is defined by

$$y_j = \alpha - \frac{\alpha - \beta}{\left(1 + (\kappa x_j)^\delta\right)}$$

where α is the value of the upper asymptote, β is the value of the lower asymptote, κ is a scaling parameter, and δ is a parameter that controls the location of the point of inflection for the curve. The point of inflection is located at

$$x = \left[\frac{\delta - 1}{\delta + 1} \right]^{1/\delta}$$

$$y = \frac{\delta - 2}{2\delta}$$

for $\delta \geq 1$

There is no point of inflection for $\delta < 1$. All the MMF curves are sublogistic, in the sense that the point of inflection is always located below 50% growth (0.5α). The following constraints apply to the fit coefficient values:

$$\alpha > 0, \beta > 0, \kappa > 0, \delta > 0$$

$$\alpha > \beta$$

The response feature vector g is given by

$$g = \left[\alpha \quad \beta \quad \kappa \quad \delta \quad \frac{\delta - 1}{2\delta} \right]^T$$

Four-Parameter Logistic Curve

The four-parameter logistic model is defined by

$$y_j = \beta + \frac{\alpha - \beta}{1 + \exp\left(-\kappa(\log(x_j) - \gamma)\right)}$$

with constraints $\alpha, \beta, \kappa, \gamma > 0$, $\beta < \alpha$ and $\beta < \gamma < \alpha$. Again, α is the value of the upper asymptote, κ is a scaling factor, and γ is a factor that locates the x-ordinate of the point of inflection at

$$\exp\left(\frac{\kappa\gamma - \log\left(\frac{1+\kappa}{\kappa-1}\right)}{\kappa}\right)$$

The following constraints apply to the fit coefficient values:

- All parameters > 0
- $\alpha > \beta$

This is the available response feature vector:

$$g = \left[\alpha \quad \beta \quad \kappa \quad \gamma \quad \frac{\left(\kappa\gamma - \log\left(\frac{1+\kappa}{\kappa-1}\right)\right)}{\kappa} \right]^T$$

Richards Curves

The Richards curves family of growth models is defined by the equation

$$y_j = \alpha \left[1 + (\delta - 1) \exp\left(-\kappa(x_j - \gamma)\right) \right]^{\frac{1}{1-\delta}} \quad \delta \neq 1$$

where α is the upper asymptote, γ is the location of the point of inflection on the x axis, K is a scaling factor, and δ is a parameter that indirectly locates the point of inflection.

The y-ordinate of the point of inflection is determined from

$$\frac{\alpha}{\delta^{1/(\delta-1)}} \quad \delta > 0$$

Richards also derived the average normalized growth rate for the curve as

$$\frac{\kappa}{2(\delta + 1)}$$

The following constraints apply to the fit coefficient values:

- $\alpha > 0$, $\gamma > 0$, $K > 0$, $\delta > 0$
- $\alpha > \gamma$
- $\delta \neq 1$

Finally, the response feature vector g for Richards family of growth curves is defined as

$$g = \left[\alpha \quad \gamma \quad \kappa \quad \delta \quad \frac{\delta \kappa}{2\delta + 1} \right]^T$$

Weibul Growth Curve

The Weibul growth curve is defined by the equation

$$y_j = \alpha - (\alpha - \beta) \exp\left(-(\kappa x_j)^\delta\right)$$

where α is the value of the upper curve asymptote, β is the value of the lower curve asymptote, κ is a scaling parameter, and δ is a parameter that controls the x-ordinate for the point of inflection for the curve at

$$\left(\frac{1}{\kappa}\right)\left(\frac{\delta-1}{\delta}\right)^{1/\delta}$$

The following constraints apply to the curve fit parameters:

- $\alpha > 0, \beta > 0, \kappa > 0, \delta > 0$
- $\alpha > \beta$

The associated response feature vector is

$$\mathbf{g} = \left[\alpha \quad \beta \quad \kappa \quad \delta \left(\frac{1}{\kappa}\right)\left(\frac{\delta-1}{\delta}\right)^{1/\delta} \right]^T$$

Exponential Growth Curve

The exponential growth model is defined by

$$y_j = \alpha - (\alpha - \beta)\exp(-\kappa x_j)$$

where α is the value of the upper asymptote, β is the initial size, and κ is a scale parameter (time constant controlling the growth rate). The following constraints apply to the fit coefficients:

- $\alpha > 0, \beta > 0, \kappa > 0$
- $\alpha > \beta$

The response feature vector \mathbf{g} for the exponential growth model is defined as

$$\mathbf{g} = \left[\alpha \quad \beta \quad \kappa \right]^T$$

Gompertz Growth Model

Another useful formulation that does not exhibit a symmetric point of inflection is the Gompertz growth model. The defining equation is

$$y_j = \alpha \exp\left(-e^{-\kappa(x_j - \gamma)}\right)$$

where α is the final size achieved, K is a scaling factor, and γ is the x-ordinate of the point of inflection. The corresponding y-ordinate of the point of inflection occurs at

$$\frac{\alpha}{e}$$

With maximum growth rate

$$\frac{\kappa \alpha}{e}$$

The following constraints apply to the selection of parameter values for the Gompertz model:

$$\alpha > 0, K > 0, \gamma > 0$$

The response feature vector g for the Gompertz growth model is defined as

$$g = \left[\alpha \gamma \kappa \frac{\kappa \alpha}{e} \right]^T$$

Local Model Class: Linear Models

Select **Linear Models** and then click **Setup**.

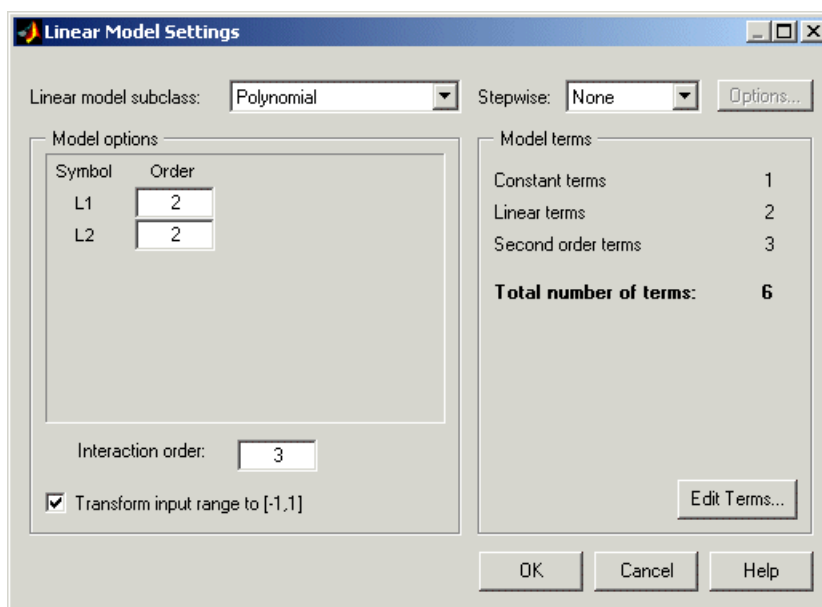
You can now set up polynomial or hybrid spline models. The settings are exactly the same as the global linear models.

These models are for multiple input factors - for single input factors you can use a different polynomial model from the Local Model Class list, where you can only change the polynomial order. See “Local Model Class: Polynomials and Polynomial Splines” on page 4-28.

If there is more than one input factor, you can only choose Linear Models, Average Fit, Multiple Models or Transient Models from the Local Model Class list. Under Linear Models you can choose Polynomial or Hybrid Spline. This

polynomial is a different model where you can change more settings such as Stepwise, the Term Editor (where you can remove any model terms) and you can choose different orders for different factors (as with the global level polynomial models).

See “Global Linear Models: Polynomials and Hybrid Splines” on page 4-8 for details.



Note that by default inputs are transformed to $[-1, 1]$ before fitting and evaluating polynomials. This is important as differences in scales between inputs can cause numerical problems. Generally it is a good idea to transform inputs as this alleviates problems with variables of different scales. In some circumstances you may be concerned with the values of polynomials (for example if your strategy requires raw polynomial coefficients). In this case you can clear the **Transform input range** check box to use untransformed units to calculate natural polynomials.

Different response features are available for this Linear Models: Polynomial model and the other Polynomial choice (for single input factors). You can view these by clicking the **Response Features** tab on the Local Model Setup

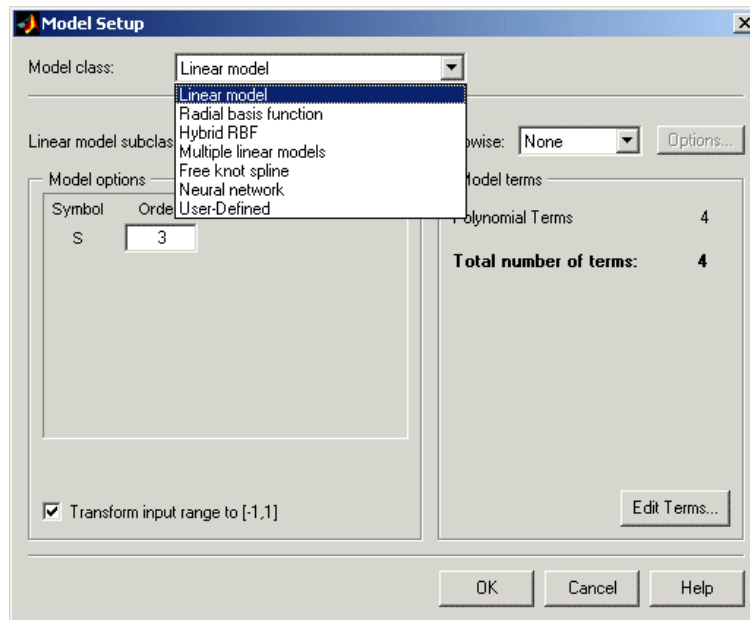
dialog. Single input polynomials can have a datum model, and you can define response features relative to the datum. See “Datum Models” on page 4-72.

These linear models are labeled Quadratic, Cubic, and so on, on the test plan block diagram, while the single input type of polynomials is labeled Poly2, Poly3, and so on. For higher orders, both types are labeled Poly n .

Local Model Class: Average Fit

You can use this local model class to fit the same model to all tests. Sometimes it is desirable to try to fit a single one-stage model to two-stage data (for example, fitting an RBF over the whole operating region of spark, speed, load, air/fuel ratio and exhaust gas recirculation). However, it can still be useful to be able to examine the model fit on a test-by-test basis. You can use Average Fit for this purpose.

Select Average Fit and click **Setup**. The Model Setup dialog appears.



In the **Model class** drop-down menu is a list of available models. This list contains the same models that you would find in the global model setup of a one-stage model. Note the number of inputs changes which models are available. A local model with only one input can access all the models seen in the example above. See “What Models Are Available?” on page 4-4.

The Average Fit local model class allows you to use any of these global model options to fit to all tests. In the same way that global models are fitted to all the data simultaneously, using average fit allows you to fit the same model to every test in your data, instead of fitting a separate local model for each test.

The advantage of this is that you can use these one stage models to fit your data while also being able to view the fit to each test individually. You should set up your global model with an input such as record number or a dummy variable. Make all the variables you want to model local inputs. It does not matter what dummy variable or model type you use for the global input - it is only there to distinguish the Local Average Fit model from a one-stage model. The dummy global variable has no influence on the model fit. The Average Fit model is fitted in the same way as a one-stage model (across all tests simultaneously) but the main difference is that you can analyze the fit to each test individually. You cannot do this when fitting one-stage models. No two-stage model is available with local Average Fit models.

Local Model Class: Multiple Models

You can use multiple models to try a variety of models for your data. These can be useful if you want to try several kinds of models, especially if you think there is a lot of variation between tests in your data. All the models you choose will be fitted and the best one chosen for each test in your data. In this way you can have a variety of local models at once, for example for some tests a spline or radial basis function may fit best, while for others a quadratic would be fine. You can choose any of the summary statistics as the selection criteria for deciding which model fits best to each test. Which statistics appear in the options is determined by your selections for the Summary Statistics.

The main purpose of this model is to provide a convenient mechanism to model a number of different tests using the same set of models. The data can be divided into tests and modeled within a single test plan rather than having a separate one-stage test plan for each test. You can use any model available as one-stage models (including radial basis functions (RBF) and hybrid RBF).

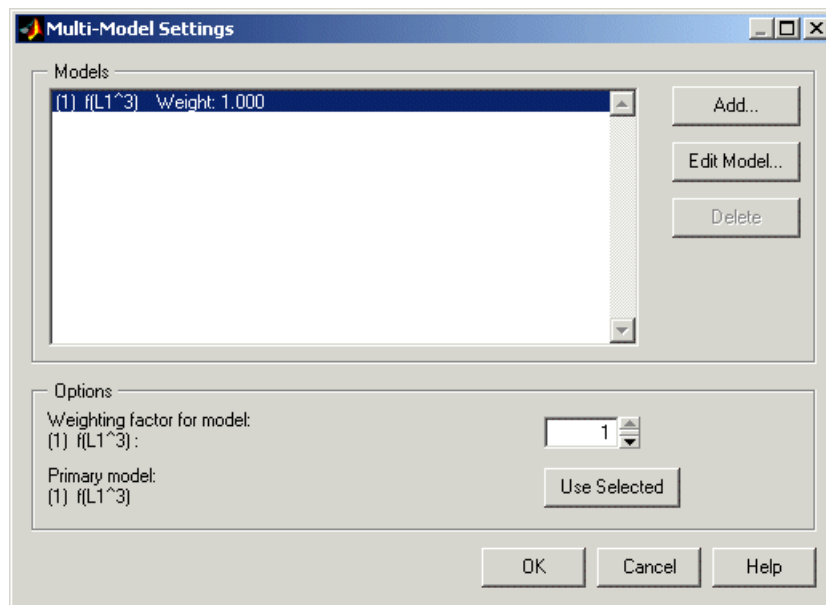
Two-stage models or response feature models are not constructed (with different model types for different tests it is impossible to choose response features that apply to all tests). You must have at least one global variable (e.g. speed, injection timing, load) and you cannot use covariance modeling.

This provides a smooth interface with the Multimodel Tradeoff in the CAGE browser part of the Model-Based Calibration toolbox. In this case two global inputs are required (often speed and load) which form the axes of the tradeoff tables. This useful application for multiple models allows you to calibrate from local maps.

You can export multimodels and import them directly into CAGE. From the Test Plan node in the Model Browser you can select **TestPlan > Export Multiple** . There will be a single CAGE model for each response. The models will only be defined at operating points corresponding to the global inputs — for example at the values of speed and load where you performed tests.

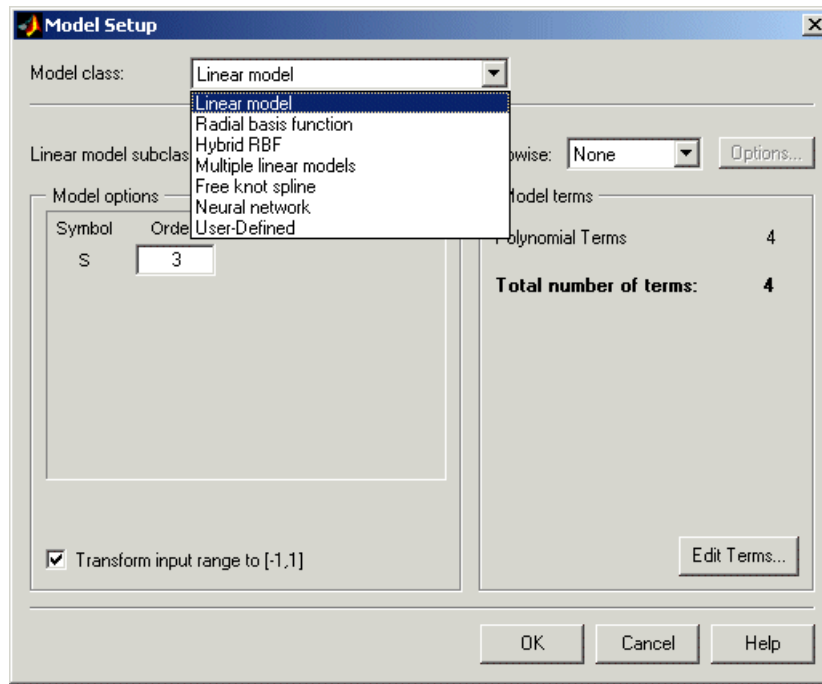
Select **Multiple Models** from the model class menu.

- Click **Edit**. The Multi-Model Settings dialog appears. This is similar to the method for using multiple models at the global level, see “Global Model Class: Multiple Linear Models” on page 4-21. Controls are described below.



- Alternatively you can click **Template**. The “Build Models Dialog” on page 4-106 opens. Here you can choose a template to build a selection of models. There are predefined templates for polynomials, radial basis functions, hybrid radial basis functions, or free knot splines, and you can also save your own templates of any models you choose.

Click **Add** to reach the Model Setup dialog. Here you can choose from all the global models that would be available for a one-stage model with the same number of inputs as your current local model. The example below shows the options for a single input local model.



You can choose and set up these models using the same global model options. Click **OK** to add the model to your Multi-Model and return to the Multi-Model Settings dialog. You can always click **Edit Model** to change the settings for any of the models in the list. Model weights are not used, the selection criteria used are the ones you choose from the drop-down menu (such as RMSE or PRESS RMSE) in the **Model options** pane of the Local Model Setup dialog. You return to this dialog when you click **OK** in the Multi-Model Settings dialog. The selection criteria that appear in this menu are determined by your selected options for the Summary Statistics displayed.

All the models you have chosen are fitted to each test individually, and the best fit to each test is chosen by the selection criteria you picked.

Local Model Class: User-Defined Models

You must check user-defined models into the Model-Based Calibration Toolbox before you can use them here. The only model checked in by default is Exponential, which requires a single input factor.

Throughout this section we refer to the template file

```
<MATLAB root>\toolbox\mbc\mbcmodel\@xregusermod\weibul.m
```

This file defines the local model Weibul in the Model-Based Calibration Toolbox and hence should not be changed. Make sure you do not overwrite the file by saving under another name if you alter it.

The `xregusermod` class allows you to define new local models. When the data being modeled at the local level matches the profile of (checked-in) user-defined local models, these models appear as options on the Local Model Setup dialog. At the end of the following section MBC offers the user-defined model `weibul` under the category User-defined in the Local Model Setup dialog.

- 1 To begin, decide on the function that you want to use. For this example, use the Weibul function:

$$y = \alpha - (\alpha - \beta) \cdot \exp(-(\kappa \cdot x)^\delta)$$

- 2 Open the template file

```
<MATLAB root>\toolbox\mbc\mbcmodel\@xregusermod\weibul.m
```

Observe that the M-file is called using

```
varargout= weibul(U,X,varargin)
```

Where the variables are given by

```
U = the xregusermod object
X = input data as a column vector for fast eval -OR-
X = specifies what sub-fcn to evaluate (not usually called
directly)
```

The first function in the template file is a vectorized evaluation of the function. First the model parameters are extracted:

```
b= double(U);
```

Then the evaluation occurs:

```
y = b(1) - (b(1)-b(2)).*exp(-(b(3).*x).^b(4));
```

Note that the parameters are always referred to and declared in the same order.

You must edit these subfunctions as follows:

```
function n= i_nfactors(U,b,varargin);  
n= 1;
```

This is the number of input factors. For functions $y = f(x)$ this is 1.

```
function n= i_numparams(U,b,varargin);  
n= 4;
```

This is the number of fitted parameters. In this example there are four parameters in the Weibul model.

```
function [param,OK]= i_initial(U,b,X,Y)  
param= [2 1 1 1]';  
OK=1;
```

This subfunction returns a column vector of initial values for the parameters to be fitted. The initial values can be defined to be data dependent; hence there is a flag to signal if the data is not worth fitting. In the template file `weibul.m` there is a routine for calculating a data-dependent initial parameter estimate. If no data is supplied, the default parameter values are used.

Optional Subfunctions

```
function [LB,UB,A,c,nlcon,optparams]=i_constraints(U,b,varargin)  
LB=[0 0 0 0]';  
UB=[1e10 1e10 1e10 1e10]';  
A= [-1 1 0 0];
```

```

c= [0];
nlcon= 0;
optparams= [];

```

Lower and upper bounds are stated for the model parameters. These are in the same order that the parameters are declared and used throughout the template file.

A linear constraint is defined:

```

b = the vector of model parameters, then the constraint is A*b < c
We define A and c in the subfunction above

```

The number of nonlinear constraints is declared to be zero. If the number of nonlinear constraints is not zero, the nonlinear constraints are calculated in `i_nlconstraints`.

No optional parameters are declared for the cost function.

```

function fopts= i_foptions(U,b,fopts)

fopts= optimset(fopts,'Display','iter');

```

The fit options are always based on the input `fopts`. See MATLAB help on the function `optimset` for more information on fit options. When there are no constraints the fitting is done using the MATLAB function `lsqnonlin`, otherwise `fmincon` is used.

```

function J= i_jacobian(U,b,x)

x = x(:);
J= zeros(length(x),4);

a=b(1); beta=b(2); k=b(3); d=b(4);

ekd= exp(-(k.*x).^d);
j2= (a-beta).*(k.*x).^d.*ekd;

J(:,1)= 1-ekd;
J(:,2)= ekd;

```

```
J(:,3)= j2.*d./k;
J(:,4)= j2.*log(k.*x);
```

To speed up the fitting algorithm an analytic Jacobian can be supplied, as it is here.

```
function c= i_labels(U,b)
c={'\alpha', '\beta', '\kappa', '\delta'};
```

These labels are used on plots and so on. Latex notation can be used and is formatted.

```
function str= i_char(U,b,fopts)

s= get(U,'symbol');
str=sprintf('%.3g - (%.3g-%.3g)*exp(- (%.3g*x)^{%.3g})',...
    b([1 1 2 3 4]));
```

This is the display equation string and can contain Latex expressions. The current values of model parameters appear.

```
function str= i_str_func(U,b)

s= get(U,'symbol');
lab= labels(U);
str= sprintf('%s - (%s - %s)*exp(- (%s*x)^{%s})',...
    lab{1},lab{1},lab{2},lab{3},lab{4});
```

This displays the function definition with labels appearing in place of the parameters (not numerical values).

```
function rname= i_rfnames(U,b)
rname= {'INFLEX'};
```

This does not need to be defined (can return an empty array). Here we define a response feature that is not one of the parameters (here it is also nonlinear).

```
function [rf,dG]= i_rfvals(U,b)

% response feature definition
rf= (1/b(3))*((b(4)-1)/b(4))^(1/b(4))
```

```

if nargout>1
% delrf/delbi
dG= [0, 0, -((b(4)-1)/b(4))^(1/b(4))/b(3)^2,...
      1/b(3)*((b(4)-1)/b(4))^(1/b(4))*...
      (-1/b(4)^2*log((b(4)-1)/b(4))+(1/b(4)-...
      (b(4)-1)/b(4)^2)/(b(4)-1))];
end

```

The response feature (labeled as INFLEX above) is defined. The jacobian is also defined here as dG.

```

function p= i_reconstruct(U,b,Yrf,dG,rfuser)
p= Yrf/dG';

f= find(rfuser>size(p,2));
if any(rfuser==4)
% need to use delta
p(:,3)= ((p(:,4)-1)./p(:,4)).^(1./p(:,4))./Yrf(:,f));
end

```

If all response features are linear in the parameters this function does not need to be defined. Here we must first find out which response features (if any) are user-defined. This subfunction allows the model parameters to be reconstructed from the response features we have been given.

Checking In To MBC

Having created a model template file, save it somewhere on the path.

To ensure that the model you have defined provides an interface that allows MBC to evaluate and fit it, we *check in* the model. If this procedure succeeds, the model is registered by MBC and is thereafter available for fitting at the local level whenever appropriate input data is being used.

At the command line, with the template file on the path, create a model and some input data, then call `checkin`. For the user-defined Weibul function, the procedure is as follows:

- 1 Create an `xregusermod` using the template model file called `weibul`.

```
m = xregusermod('name', 'weibul');
```


- 2 Call `checkin` with the transient model, its name, and some appropriate data.

```
checkin(m, 'weibul', [0.1:0.01:0.2]');
```

This creates some command line output, and a figure appears with the model name and variable names displayed over graphs of input and model evaluation output. The final command line output (if `checkin` is called with a semicolon as above) is

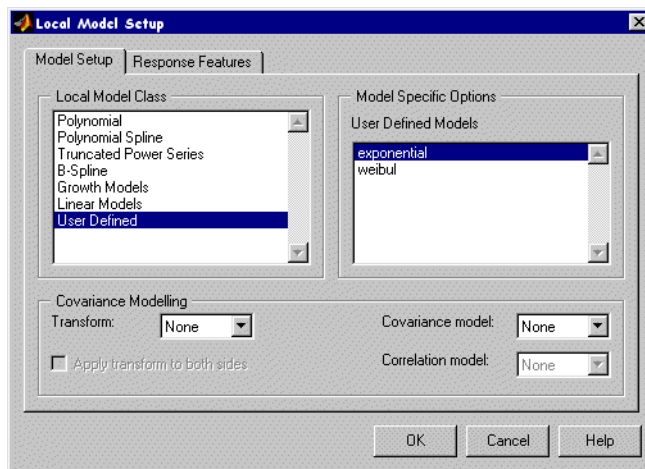
```
Model successfully registered for MBC
```

The figure displayed for `weibul` is



- 1 Start the Model Browser and load data that has the necessary format to evaluate your user-defined model.

- 2 Set up a Test Plan with one Stage1 input factor. The Local Model Setup dialog now offers the following options:



Here two user-defined models exist: `exponential` and `weibul`, as checked in.

Local Model Class: Transient Models

These are supported for multiple input factors, where time is one of the factors. You can define a dynamic model using Simulink and a template file that describes parameters to be fitted in this model. We provide an example already checked in to the Model-Based Calibration Toolbox that you can use them for modeling. The example requires two inputs.

Using the Example Transient Model

Note that the example provided requires two input factors. You can define a dynamic model using a Simulink model (.mdl file) and a template M-file that describes parameters to be fitted in this model. You must check these into the Model-Based Calibration Toolbox before you can use them for modeling.

We provide an example called `fuelpuddle` which is already checked in to the toolbox ready for modeling. The process of creating and checking in your own transient models is described in the following sections using this example.

Throughout this section you use this Simulink model:

```
<MATLAB root>\toolbox\mbc\mbcsimulink\fuelPuddle.mdl
```

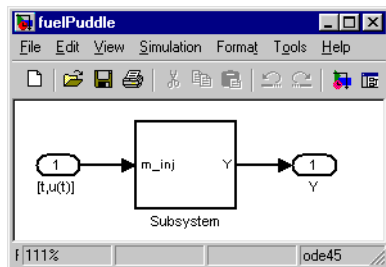
and the template file

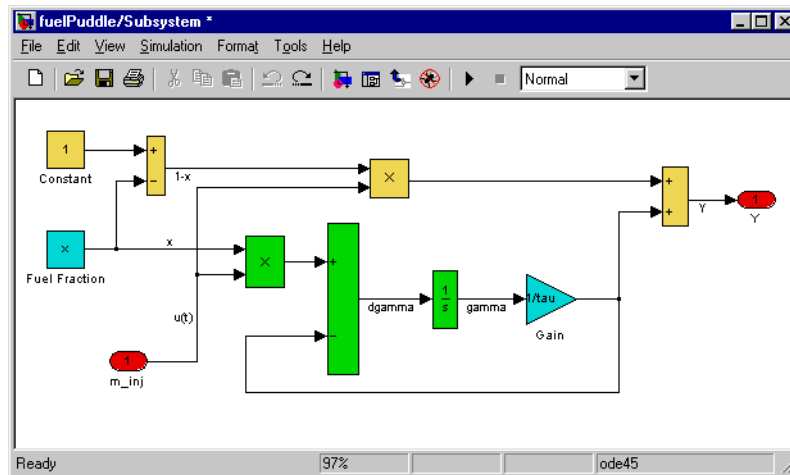
```
<MATLAB root>\toolbox\mbc\mbcmodel\@xregtransient\fuelPuddle.m
```

The Simulink Model

Create a Simulink model to represent your transient system. This has some free parameters that MBC fits by regression. There are three places to declare these variables and they all need to match (see Parameters section below).

Any subsystems can be masked so that, at the highest level, the Simulink model looks like this:





The block labels are not important. The model inputs are represented by Simulink inports at the top level. The output is represented by a Simulink output. Only one output is supported in MBC.

The model returns a single vector output (here labeled “Y”).

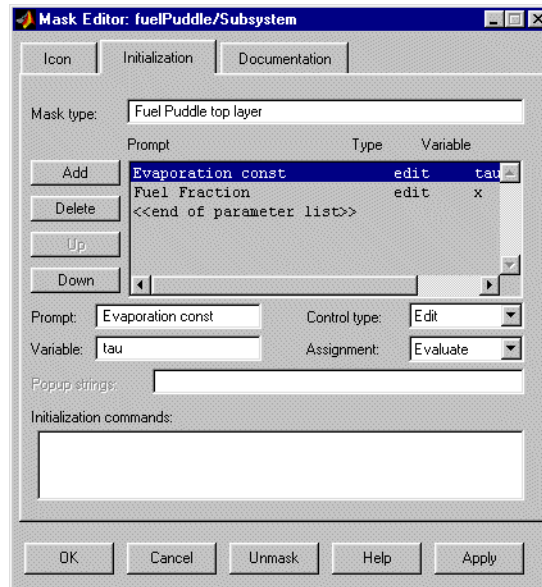
Parameters

Inside the subsystem, some blocks depend on parameters that have not been quantified. In `fuelPuddle.mdl` there is a constant block whose value is ' x ' and a gain block with the value ' $1/\tau$ '.

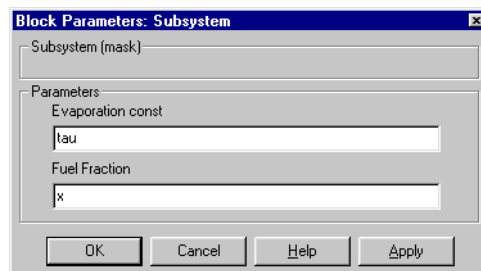
These are the parameters we will be fitting by regression using MBC. The following setup needs to be followed:

- 1 Having masked the subsystem, you need to pass the unknown parameters in to the subsystem. We need to pass them in from the workspace. This is done via a Block Parameters dialog. To create such a dialog, right-click the masked subsystem and select **Edit mask** to see the Mask Editor.

Use the **Add** button to add fields. Here two parameters are required. The **Variable** field must match the name of the corresponding parameter in the subsystem. Having declared the necessary parameters (all the unknowns of the subsystem), click **OK** to close and save.



- 2 On double-clicking the subsystem mask, the Block Parameters dialog appears with the Prompt and variable names declared in the Mask Editor. By default the variable names appear in the edit boxes, and these should not be changed. It is these names that the Simulink model tries to find in the workspace.



- 3 The same variable names must also be used in the template M-file that we shall discuss later. For this example the template M-file is `fuelPuddle.m`.

Testing the Simulink Model

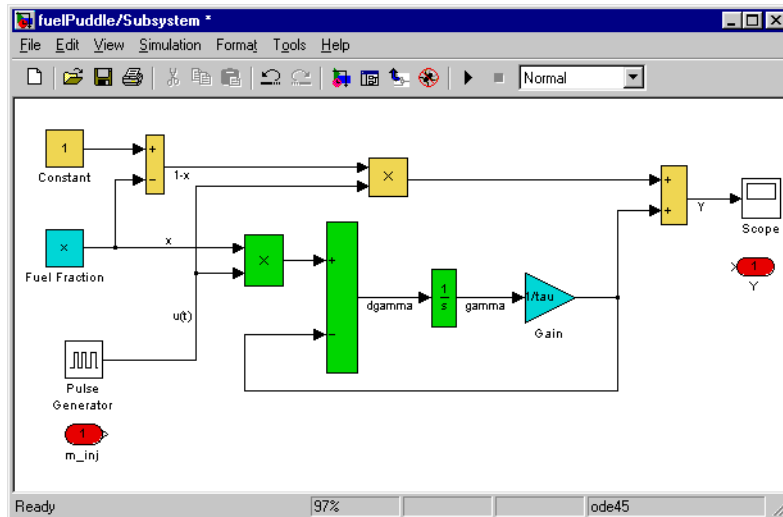
Open your Simulink model and clear the MATLAB workspace.

Define the masked block parameters in the workspace. Here, for example, you could enter at the command line

```
tau = 0.5;
x = 0.3;
```

In the subsystem, connect sources to the input factors (for example, use a pulse generator at the input) and run the model (click **Play**). If the model is taking parameters from the workspace it should run okay.

Clear the workspace again and, without defining `tau` and `x`, try to run the model; it should fail to run.



Creating Transient Models

You can create a transient model by typing the following at the command line:

```
m = xregtransient
```

You will see the following output at the command line:

```
m = functemplate([X1,X2],b=[0.5,0.5])
```

The default transient model is created. This model can be evaluated and fitted, but remember it uses the default Simulink model. To create a transient model that uses your own Simulink model, you must first define a template file for your Simulink model and this template file must have the same name as your Simulink model. The template file for `fuelPuddle.mdl` is therefore called `fuelPuddle.m`.

Template File

Note that the template file has a common layout for transient models and for user-defined models. It is not compulsory to modify every section. It is not a problem if you don't want to use several sections of the template file. The sections not actively read must have sensible defaults, so if you do not wish to use a section, leave it unmodified and it will be fine.

Open the file

```
<MATLAB root>\toolbox\mbc\mbcmodel\@xregtransient\fuelPuddle.m
```

The fast `eval` at the top of the file is not used for transient models and can be left as is.

The next section of commented code gives a summary of the functions in the template file. The function definitions follow this.

Subfunctions

Those subfunctions that must be edited are as follows:

```
function vars= i_simvars(U,b,varargin);  
vars = {'tau','x'};
```

These are the parameters of the Simulink model that will be fitted.

This subfunction must return a cell array of strings. These are the parameter names that the Simulink model requires from the workspace. These strings *must* match the parameter names declared in the Simulink model (see “Parameters” on page 4-57).

```
function [vars,vals]= i_simconstants(U,b,varargin);
vars = {};
vals = [];
```

If your Simulink model requires constant parameters to be defined you do so here. These are constant parameters required by the Simulink model from the workspace. They are not fitted. These parameters must be the same as those in the Simulink model and all names must match. Here `fuelPuddle` requires no such parameters, and hence we return an empty cell array and empty matrix.

```
function [ic]= i_initcond(U,b,X);
ic=[];
```

Initial conditions (for the integrators) are based on the current parameters, and inputs could be calculated here.

```
function n= i_numparams(U,b,varargin);
n= 2;
```

This is the number of fitted parameters. For `fuelPuddle` we have two parameters, `x` and `tau`.

```
function n= i_nfactors(U,b,t,u);
n= 2;
```

This is the number of input factors, including time. For `fuelPuddle` we input $X = [t, u(t)]$ and hence the number of input factors is 2.

```
function [param,OK]= i_initial(U,b,X,Y)
param= [.5 .5]';
OK=1;
```

This subfunction returns a column vector of initial values for the parameters that are to be fitted. You could do some calculations here to estimate good initial values from the X, Y data. The initial values can be defined to be data dependent, hence there is a flag to signal if the data is not worth fitting. Here we simply define some default initial values for x and τ .

Optional Subfunctions

The remaining subfunctions need not be edited unless they are required. The comments in the code describe the role of each function. Mostly these functions are used when you are creating a template for user-defined models. There is only one of these subfunctions relevant here.

```
function c= i_labels(U,b)
b= {'\tau', 'x'};
```

These labels are used on plots and so on. You can use LaTeX notation, and it is correctly formatted.

Checking In To MBC

Having created a Simulink model and the corresponding template file, you must save each somewhere on the path. One place to put these files is

```
<MATLAB root>\work
```

Note that you must create the folder @xregtransient somewhere on the path (such as <MATLAB root>\work) and save your `template.m` file inside this new folder.

To ensure that the transient model you have defined provides an interface that allows MBC to evaluate and fit it, we *check in* the model. If this procedure succeeds, the model is registered by MBC and is thereafter available for fitting at the local level whenever appropriate input data is being used.

Check In. At the command line, with both template file and Simulink models on the path, create a model and some input data, then call `checkin`. The procedure follows.

- 1 Create some appropriate input data. The data is not important; it is to check if the model can be evaluated.


```
t = [0:0.1:4*pi]'; u = sin(t); %% creates a sine wave input factor
X = [t,u];
```

- 2 Create a transient model using the template/Simulink model by typing the following at the command line.

```
m = xregtransient('name','MyTransient');
```

where MyTransient is the name you choose for your new transient model. For fuelPuddle you would type

```
m = xregtransient('name','fuelPuddle');
```

- 3 Call checkin with the transient model, its name, and the data.

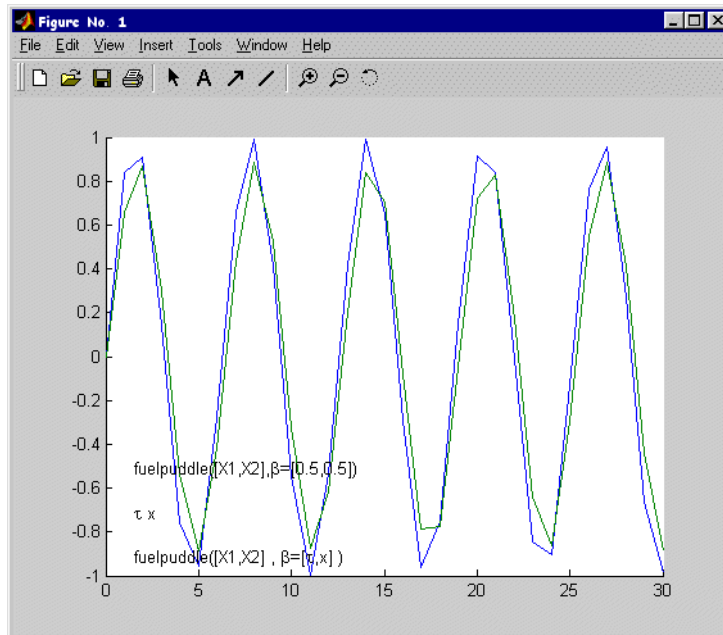
```
checkin(m, 'MyTransient', X);
```

Note that the string you use for your new transient model name is the string that will appear in the Model Setup dialog when you set up transient models using the appropriate number of inputs.

When the check in is successful, this creates some command-line output, and a figure appears with the model name and variable names displayed over graphs of input and model evaluation output. The final command line output (if checkin is called with a semicolon as above) is

```
Model successfully registered for MBC
```

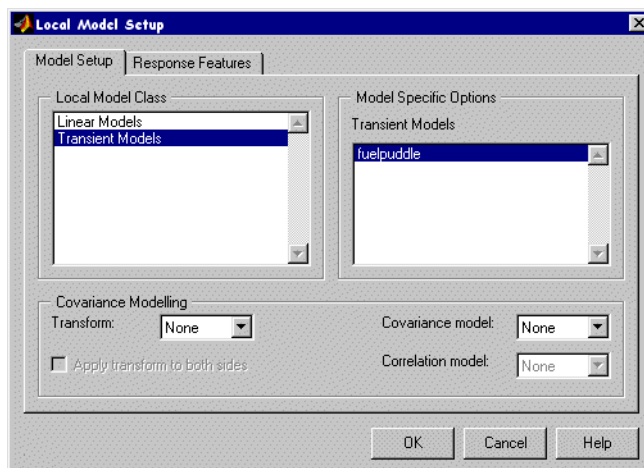
A figure is displayed for the model. Shown is the fuelPuddle example.



Using and Removing Transient Models

We provide the fuelPuddle model already checked in for you to use. Start up MBC modeling and load data that has the necessary format to evaluate this transient model — two local inputs.

Set up a Test Plan with two local input factors (the same number of input factors required by the fuelPuddle model). The Local Model Setup dialog now offers the following options:



Select the fuelpuddle model and click **OK**.

On building a response model with fuelPuddle as the local model, MBC fits the two parameters tau and x across all the tests.

Removing Checked In Models

If you wish to remove checked in models for some reason, close the toolbox and type the following at the command line:

```
mbcdeleteprefs
```

When you restart the toolbox you will see that this removes all checked in transient models. Note that this will also delete all your saved preferences for file locations.


Covariance Modeling

This frame is visible no matter what form of local model is selected in the list.

Covariance modelling

Transform: None	Covariance model: None
<input type="checkbox"/> Apply transform to both sides	Correlation model: None

Covariance modeling is used when there is heteroscedasticity. This means that the variance around the regression line is not the same for all values of the predictor variable, for example, where lower values of engine speed have a smaller error, and higher values have larger errors, as shown in the following example. If this is the case, data points at low speed are statistically more trustworthy, and should be given greater weight when modeling. Covariance models are used to capture this structure in the errors.



You can fit a model by finding the smallest least squares error statistic when there is homoscedasticity (the variance has no relationship to the variables). Least squares fitting tries to minimize the least squares error statistic

$$\sum \varepsilon_i^2, \text{ where } \varepsilon_i^2$$

is the error squared at point i .

When there is heteroscedasticity, covariance modeling weights the errors in favor of the more statistically useful points (in this example, at low engine speed N). The weights are determined using a function of the form

$$\sum \frac{\varepsilon_i^2}{W_i}$$

where W_i is a function of the predictive variable (in this example, engine speed N).

There are three covariance model types.

Power

These determine the weights using a function of the form $W_i = \hat{y}^\alpha$. Fitting the covariance model involves estimating the parameter α .

Exponential

These determine the weights using $W_i = e^{\alpha \hat{y}}$.

Mixed

These determine the weights using $W_i = \alpha + \hat{y}^\beta$. Note that in this case there are two parameters to estimate, therefore using up another degree of freedom. This might be influential when you choose a covariance model if you are using a small data set.

Correlation Models

These are only supported for equally spaced data in the Model-Based Calibration Toolbox. When data is correlated with previous data points, the error is also correlated.


There are three methods available.

- MA(1) - The Moving Average method has the form $\varepsilon_n = \alpha_1 \zeta_{n-1} + \zeta_n$.

- AR(1) - The Auto Regressive method has the form $\varepsilon_n = \alpha_1 \varepsilon_{n-1} + \zeta_n$.
- AR(2) - The Auto Regressive method of the form $\varepsilon_n = \alpha_1 \varepsilon_{n-1} + \alpha_2 \varepsilon_{n-2} + \zeta_n$
 ζ is a stochastic input, $\zeta_n \sim N(0, \sigma_\zeta^2)$.

Transforms

The following example shows the transforms available.



Input transformation can be useful for modeling. For example, if you are trying to fit

$$y = e^{a+bx+cx^2}$$

using the log transform turns this into a linear regression:

$$\log(y) = a + bx + cx^2$$

Transforms available are logarithmic, exponential, square root, y^2 , $\frac{1}{y}$, and Other. If you choose Other, an edit box appears and you can enter a function.

Apply transform to both sides is only available for nonlinear models, and transforms both the input and the output. This is good for changing the error structure without changing the model function. For instance, a log transform might make the errors relatively smaller for higher values of x . Where there is heteroscedasticity, as in the Covariance Modeling example, this transform can sometimes remove the problem.

Designing Experiments

You can design experiments after setting up models. You can design experiments for both stages, local and global. You invoke the Design Editor in several ways from the test plan level:

- Right-click a Model block in the test plan diagram and select Design Experiment.

You must select (by clicking) a stage to design for (first or second stage) or the following two options are not possible.

- Click the **Design Experiment** toolbar button .
- Select **TestPlan > Design Experiment**.

For an existing design, **View > Design Data** also launches the Design Editor (also in the right-click menu on each Model block). In this case you can only view the current data being used as a design at this stage. If you enter the Design Editor by the other routes, you can view all alternative designs for that stage.

See Chapter 2, “Designs”.

Viewing Designs

The view design facility enables the user to investigate the statistical properties of the current data.

From the test plan node, select the model stage you are interested in by clicking, then choose **View > Design Data**. Alternatively, use the right-click menu on a Model block.

This provides access to all the Design Editor and design evaluation utility functions with the current data rather than the prespecified design. If you have done some data-matching to a design, each data point is used as a design point. You can now investigate the statistical properties of this design.

For two-stage models, viewing stage one (local model) designs creates a separate design for each test.


See “The Design Editor” on page 2-3 or the step-by-step guide in “Tutorial: Design of Experiment” in the Getting Started documentation.

New Response Models

When you first set up a test plan, the Data Wizard automatically contains the response model setup after the data matching functions.

The following applies when you return to a previously setup test plan to add a new response model, or when you click the **New** button at test plan level to add new response models to an existing test plan.

Double-click the Responses output of the block diagram or use the **New Response Model** item in the **File** menu, or use the toolbar icon. (None of these is available unless you are in the Test Plan view, that is, have the test plan node selected in the model tree. This should be obvious; you can only see the test plan block diagram with the test plan node selected.)



The Response Model Setup dialog has a list box containing all the variables in the selected data set *except* the inputs to the local and global models; you cannot use an input also as a response.

You can reach the “Local Model Setup” on page 4-26 and “Global Model Setup” on page 4-6 dialogs using the **Set Up** buttons to change the local and global

models also, and you can add datum models (maximum or minimum) if the local model supports this.

You can return to the local or global setup options individually at any time by double-clicking the block in the test plan diagram.

Datum Models

Under **Datum** you can choose a datum model, but only for some local models — polysplines and polynomials (but see [Linked Datum Models](#) following). Other local models cannot have a datum model, as they do not necessarily have a unique turning point.

The datum model tracks the maximum or minimum of the local models. This is equivalent to adding the maximum or minimum as a response feature, which can be useful for analysis if those points are interesting from an engineering point of view.

The Datum options are


- None
- Maximum — This can be useful in cases using polyspline modeling of torque against spark. The maximum is often a point of engineering interest.
- Minimum — This can be useful for cases where the object is to minimize factors such as fuel consumption or emissions.
- Linked datum model — This is only available to subsequent two-stage models within a test plan in which the *first* two-stage model has a datum model defined. In this case you can make use of that datum model. The linked datum option carries the name of the response of the first two-stage model, where it originated.


If the maximum and minimum are at points of engineering interest, like MBT or minimum fuel consumption, you can add other response features later using the datum model (for example, MBT plus or minus 10 degrees of spark angle) and track these across local models too. It can be useful to know the value of MBT when modeling exhaust temperature, so you can use a linked datum model from a previous torque/spark model. Having responses relative

to datum can also be a good thing as it means the response features are more likely to relate to a feature within the range of the data points.

You can also export the datum model along with local, global, and response models if required. See “Exporting Models” on page 4-195.

Local Level

When you select a local node (with the  icon) in the model tree, this view appears.

Note that after the two-stage model is calculated the local node icon changes to a two-stage icon () to reflect this. See the model tree for clarification. The response node also has a two-stage icon, but produces the response level view instead.

See “Local Level: Toolbar” on page 4-82 and “Local Level: Menus” on page 4-83 for details on these controls.

The display features are described in the following sections:

- “Local Special Plots” on page 4-76
- “Local Scatter Plots” on page 4-77
- “Diagnostic Statistics Pane” on page 4-79 and “Pooled Statistics” on page 4-79
- “Response Features List” on page 4-78
- “Test Notes Pane” on page 4-81
- “Data Tab” on page 4-81

The following example shows a local model of torque/spark curves. In this example the two-stage model has been calculated so you can compare the local fit and the two-stage fit on the plots.

Click here to change test

Model Browser - D:\MATLAB\SandboxMBC\matlab\toolbox\mbc\mbctraining\Gasoline_project.mat*

File Model View Outliers Window Help

All Models

- Gasoline_project
 - WVT
 - BTQ
 - PS22
 - knc
 - ma
 - Bhi
 - Blo
 - EXTEMP
 - POLY2
 - Local Model: PS22
 - Test: 1 Select Test... **Model type: $f(S^2, S^2)$**
 - Model Data
 - Residuals vs SPK (S) [-]
 - Local response vs SPK (S) [-]
 - Response Features Table

Diagnostic statistics

Local parameters

| | Value | Stc |
|---------|-----------|-----|
| knot | 27.969 | |
| max | 75.688 | |
| Bhigh_2 | -5.178e-3 | 5 |
| Blow_2 | -0.029 | 7 |

Pooled statistics

| | |
|----------------|-----------|
| Local RMSE | 0.268 |
| Two-Stage RMSE | 1.38 |
| PRESS RMSE | 1.669 |
| Two-Stage T^2 | 3.514 |
| -log L | -3331.135 |

Test notes

Test number color: Set Color...

| Response Features | Observations | Parameters | Box-Cox | PRESS RMSE | RMSE |
|-------------------|--------------|------------|---------|------------|-----------|
| knot | 188 | 39 | 0 | 2.5569 | 2.2742 |
| max | 188 | 18 | 1 | 1.2407 | 1.15 |
| Bhigh_2 | 188 | 45 | 1 | 0.0037966 | 0.0030599 |
| Blow_2 | 188 | 39 | 1 | 0.005443 | 0.0048493 |

Univariate Model with knot, max, Bhigh_2, Blow_2 is the best two-stage model

Ready

Click-and-drag anywhere along this line to change the size of the Response Features pane

Click-and-drag or double-click here to expand plots/hide statistics pane

The default view is the **Model** tab, with plots described below. You can click the **Data** tab to view plots of other variables. See “Data Tab” on page 4-81.

Local Special Plots

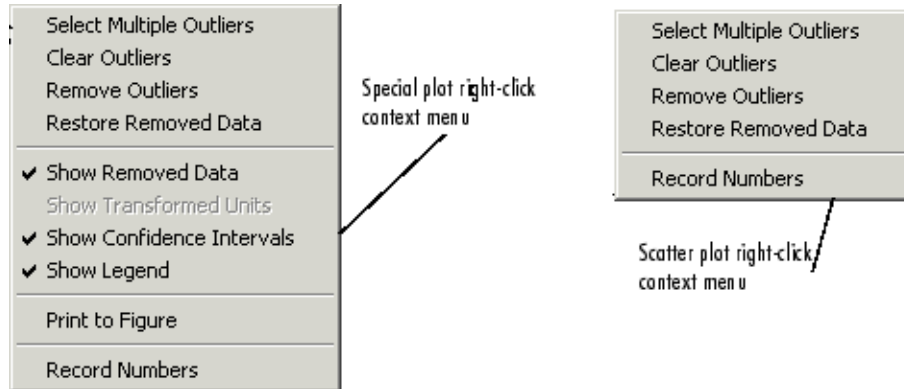
The lower plots are referred to as *special plots* as they can be different for different models.

The lower plot at the local level shows the local model fit to the data *for the current test only*, with the datum point if there is a datum model. If there are multiple inputs to the local model, a predicted/observed plot is displayed. In this case to examine the model surface in more detail you can use **Model > Evaluate**. See “Model Evaluation Window” on page 4-148.

You can scroll through all the local models by using the up/down test buttons, type directly in the edit box, or go directly to test numbers by clicking **Select Test**.

To examine the local fit in more detail, double-click the arrows (indicated in the preceding figure) to hide the scatter plot and expand the lower plot. You can zoom in on parts of the plot by **Shift**-click-dragging or middle-click-dragging on the place of interest on the plot. Return to full size by double-clicking.

You can change the lower plot from the Local Response to a Normal Plot by using the drop-down menu at the top of the plot. Normal plots are a useful graph for assessing whether data comes from a normal distribution. For more information, see “Normal Probability Plots” in the Statistics Toolbox documentation.




Above are the right-click context menus for both plots. On both plots you can manipulate outliers with all the same commands available in the **Outliers** menu. See “Outliers Menu (Local Level)” on page 4-88 for details.

The **Print to Figure** command opens a MATLAB figure plot showing the current plot. On the special plots you can switch the confidence intervals and legend on and off, and hide or show removed data. For both plots you can switch the display of **Record Numbers** on and off. This is similar to test number for global models but relates to individual records within tests.

Local Scatter Plots

The upper plots are referred to as *scatter plots*. They can show various scatter plots of statistics for assessing goodness-of-fit for the current local model shown. Upper scatter plots are replaced by an icon if you resize the Browser too small.

The statistics available for plotting are model dependent.



The preceding is an example drop-down menu on the scatter plot for changing x and y factors. In this case spark is the local input factor and torque is the response. The local inputs, the response, and the predicted response are always available in these menus. The observation number is also always available.

The other options are statistics that are model dependent, and can include residuals, weighted residuals, studentized residuals, and leverage. At local level these are internally studentized residuals.

Response Features List


Here is a list of all the response features calculated for the local model. A two-stage model using the local and global models is formed by using **Select**. Click the **Select** button here to enter the Model Selection window. This step is required before two-stage models can be calculated.

The list view displays the number of parameters and observations, the value of any Box-Cox transformation (1 indicates no transform), and the values of RMSE and PRESS RMSE (linear models only) for each response feature model. For definitions of RMSE and PRESS RMSE, see Summary Table on page 4-97. For information on Box-Cox transforms, see “Box-Cox Transformation” on page 4-176.

Click **New** to add a new response feature model (or **Delete** to remove one). For more information see the “Test Plans List Pane” on page 1-13. The contents of this pane change in different views; it always contains the child nodes of the node selected in the model tree (and the **New**, **Delete**, and **Select** buttons). At the local level it contains a list of response features.

Diagnostic Statistics Pane

The **Diagnostic Statistics** pane drop-down menu is shown, where you can select the information to be displayed in the pane.



If there is not enough room there are scroll bars.

- **Local Parameters** — Shows the values and standard errors of the parameters in the local model for the current test selected.
- **Local Correlations** — Shows a table of the correlations between the parameters.
- **Response Features** — Shows the values and standard errors of the response features defined for this local model, from the current test (often some or all of them are the same as the parameters; others are derived from the parameters).
- **Global Variables** — Shows the values and standard errors of the global variables at the position of the current test.
- **Local Diagnostics** — s_i (the standard error for the current (i^{th}) test), number of observations, degrees of freedom on the error, R squared, $\text{Cond}(J)$ and $\text{Cond}(\text{Sigma})$: the condition indices for the Jacobian matrix and the covariance matrix.
- **Global Covariance** — For MLE models, shows a covariance matrix for the response features at the global level.

Pooled Statistics

These are seen at the local node (when two-stage modeling) in the Pooled Statistics table, and at the response node in the list of local models. If you have a selection of local models, use these statistics to help you choose which model is best.

| | |
|--------------------------|---|
| Local RMSE | Root mean squared error, using the local model fit to the data for the displayed test. The divisor used for RMSE is the number of observations minus the number of parameters. |
| Two-Stage RMSE | Root mean squared error, using the two-stage model fit to the data for the displayed test. You want this error to be small for a good model fit. |
| PRESS RMSE | Root mean squared error of predicted errors, useful for indicating overfitting; see “Two-Stage Models for Engines” on page 4-181. The divisor used for PRESS RMSE is the number of observations. Not displayed for MLE models because the simple univariate formula cannot be used. |
| Two-Stage T ² | <p>T² is a normalized sum of squared errors for all the response features models. You can see the basic formula on the Likelihood view of the Model Selection window.</p> $T^2 = (y_{rf} - \hat{y})^T \Sigma^{-1} (y_{rf} - \hat{y})$ <p>Where $\Sigma = \text{blockdiag}(C_i + D)$, where C_i is the local covariance for test i. See <i>blockdiag</i> diagram following.</p> <p>A large T² value indicates that there is a problem with the response feature models.</p> |
| -log L | <p>Log-likelihood function: the probability of a set of observations given the value of some parameters. You want the likelihood to be large, tending towards -infinity, so large negative is good.</p> <p>For n observations x_1, x_2, \dots, x_n, with probability distribution $f(x, \theta)$, the likelihood is:</p> $L = \prod_{i=1}^n f(x_i, \theta)$ <p>This is the basis of “MLE” on page 4-141.</p> |

| | |
|--|---|
| | $\log L = \log(\det(\Sigma)) + (y_{rf} - \hat{y})^T \Sigma^{-1} (y_{rf} - \hat{y})$ <p>which is the same as:</p> $\log L = \log(\det(\Sigma)) + T^2$ <p>This assumes a normal distribution.</p> <p>You can view plots of $-\log L$ in the Model Selection window, see “Likelihood View” on page 4-131.</p> |
|--|---|

To explain `blockdiag` as it appears under T^2 in the Pooled statistics table: $\Sigma = \text{blockdiag}(C_i + D)$, where C_i is the local covariance for test i , is calculated as shown below.

$$\text{blockdiag}(C_i + D) = \begin{bmatrix} C_1 + D & & & \\ & C_2 + D & & \\ & & C_3 + D & \\ & & & C_i + D \end{bmatrix}$$

Test Notes Pane

You can use the **Test Notes** pane to record information on particular tests. Each test has its own notes pane. Data points with notes recorded against them are colored in the global model plots. You choose the color using the **Set Color** button in the **Test Notes** pane.

Data Tab

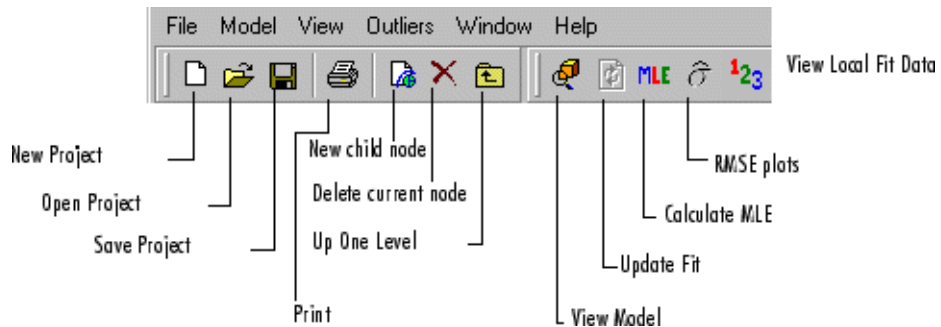
- When you click the **Data** tab at the local level or select **View > Data Plots**, you can view plots of the data for the current test.
- Use the right-click menu item **Set Up Plot Variables** to open the Plot Variables Setup dialog. The dialog appears automatically if you open the tab using the **View** menu.

Here you can choose to view any of the data signals in the data set for the current test (including signals not being used in modeling). Choose variables from the list on the left and use the buttons to move them into the

Y Variable(s) list or the **X variable** edit box. You can use the **No X Data** button to plot a variable against record number only. Remember you can also display values of global variables in the Diagnostic Statistics pane if you want to see these values at the same time as the **Model** tab.

Local Level: Toolbar

This toolbar appears when a local node is selected in the model tree.



The eight left icons remain constant throughout the levels. They are for project and node management, and the help button, and here the print icon is enabled as there are plots in this view. See “Project Level: Toolbar” on page 1-15 for details on these buttons. In the example shown the slider bar has been dragged to hide the Help button.

- View Model — Opens a dialog displaying the terms in the current model.
- Update Fit — This button is only enabled when data has been excluded from the plot using the **Remove Outliers** command (in the right-click context menu or the **Outliers** menu). At this point the local fit in the view is updated to fit only the remaining data, but this change also affects a point in all the global models. You can make this update to all the global models by using this toolbar button, or it happens automatically when another node is selected.

Note Update fit can affect several models. Removing an outlier from a best local model changes all the response features for that two-stage model. The global models all change; therefore the two-stage model must be recalculated. For this reason the local model node returns to the local (house) icon and the response node becomes blank again. If the two-stage model has a datum model defined, and other models within the test plan are using a datum link model, they are similarly affected.

- **Calculate MLE** — Calculates the two-stage model using maximum likelihood estimation. This takes correlations between response features into account. See “MLE” on page 4-141 for details.
- **RMSE Plots** — Opens the RMSE Explorer dialog, where you can view plots of the standard errors of all the response features. There is one value of standard error per test for each response feature. You can also plot these standard errors against the global variables to examine the global distribution of error.
- **View Local Fit Data** — Opens the Data Editor, where you can view a read-only version of the inputs, predicted and actual responses. This allows you to examine and export your modeling data, using all the powerful display features of the Data Editor. Also in the **View** menu.

Local Level: Menus

File Menu

Only the **New** (child node) and **Delete** (current node) functions change according to the node level currently selected. Otherwise the **File** menu remains constant.

See “File Menu” on page 1-16.


Window and Help Menus

The **Window** and **Help** menus remain throughout the Model Browser, offering access to different windows, general help and context help.

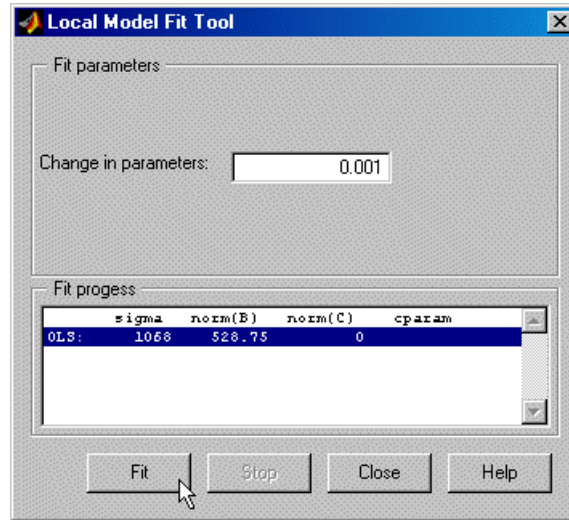
See “Window Menu” on page 1-18 and “Help Menu” on page 1-18. See also

- “Model Menu (Local Level)” on page 4-84
- “View Menu (Local Level)” on page 4-87
- “Outliers Menu (Local Level)” on page 4-88
- “Outlier Selection Criteria” on page 4-89

Model Menu (Local Level)

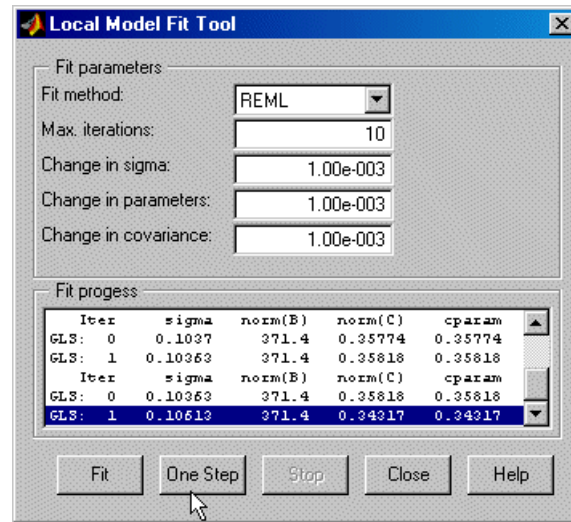


- **Set Up** — Opens the Local Model Setup dialog where you can change the model type. See “Local Model Setup” on page 4-26.
- **Fit Local** — Opens the Local Model Fit Tool dialog. Without covariance modeling, you see the following controls. This example shows the results after clicking **Fit** once. The optimization process can be stopped early by clicking **Stop** or you can wait until it finishes. The Ordinary Least Squares (OLS) parameters are displayed. You can click **Fit** to run the process again as many times as required, or **Close** to exit the dialog. You can enter a different change in parameters in the edit box.



- For covariance models this offers three different algorithms: REML (Restricted Maximum Likelihood - the default), Pseudo-likelihood, and Absolute residuals. The following example shows that there is also an additional button, **One Step**. Using the **Fit** button might take several steps to converge, but if you use the **One Step** button only one step in the optimization process is taken.


Every time you run a process, the initial and final Generalized Least Squares parameter values are displayed for each iteration.



- **Update Fit** — Only enabled when outliers have been removed. The **Update Fit** option updates the global models to take these changes into account. This happens automatically when you select a different node. Duplicated in the toolbar. This can affect many models; see “Local Level: Toolbar” on page 4-82.
- **Calculate MLE** — Calculates the two-stage model using maximum likelihood estimation. This takes interactions between response features into account. Duplicated in the toolbar. See “MLE” on page 4-141 for details.
- **Evaluate** — Opens the Model Evaluation window.
- **Select** — Available whenever the **Select** button is also enabled in the lower right pane (when it is titled **Local Models**, **Response Features**, or **Models**). This item opens the Model Selection window to allow you to choose the best model. See “Select Button” on page 4-111.
- **Assign Best** — Selects the current model as best. If it is one of several child node models of a response model, selecting it as best means that this local model (and associated response features) is used for the two-stage model. Note that this is only enabled if the local node selected has a

two-stage model calculated; that is, if the local node still has a local icon (a house) you cannot use **Assign Best**. See “Model Tree” on page 1-19.

View Menu (Local Level).




- **Model Definition** — Opens the Model Viewer dialog, showing the terms in the model.
- **Data Plots** — Opens the **Data** tab on the display and the Plot Variables Setup dialog, where you can choose to view any of the data signals in the data set for the current test (including signals not being used in modeling). Choose variables from the list on the left and use the buttons to move them into the **Y Variable(s)** list or the **X variable** edit box. You can use the **No X Data** button to plot a variable against record number only.
- **RMSE Plots** — Opens the RMSE Explorer where you can view plots of the standard errors of all the tests, both overall and by response feature. There is one value of standard error per test, overall and for each response feature. It is most useful to first plot plain s.e. against test number to get an idea of how the error is distributed and locate any tests with much higher errors. Right-click to select display of test numbers. This can be useful for quickly identifying problem tests. Ideally, all the standard errors should be roughly the same value otherwise some of the statistics theory breaks down. You can also plot these standard errors against the global variables to examine the global distribution of error.
- **Local Fit Data** — Opens the Data Editor, where you can view a read-only version of the inputs and predicted values. This allows you to examine your

modeling data and inputs using all the powerful display features of the Data Editor. You can also export your modeling data to the workspace or Excel. Also in the toolbar.

- **Response Data** — Opens the Data Editor, where you can view a read-only version of the values of all the response features for the local models. You can also use the Data Editor export facilities.
- **Record numbers** — Toggles the display of record numbers on plots.
- **Next Test, Previous Test, and Select Test** — Duplicate the buttons for changing tests above the plots in the top left of the **Local Model** display pane.

Outliers Menu (Local Level).



All the commands except **Remove All Data** and **Copy Outliers** are also available in the right-click context menus on all plots.

- **Select Multiple Outliers** — Use this item to draw a selection box around as many data points as required to select them all as outliers. This is useful for removing many data points at once.
- **Clear Outliers** — Returns all data points to the unselected state (that is, no points outlined in red) as possible outliers.
- **Remove Outliers** — Removes red-outlined data points from the fit and refits the current local fit only. Use the Update Fit toolbar button or **Model > Update Fit** to refit all the global models also. This also happens automatically when another node is selected.
- **Restore Removed Data** — Opens the Restore Removed Data dialog, where you can choose the points to restore from the list by record number,

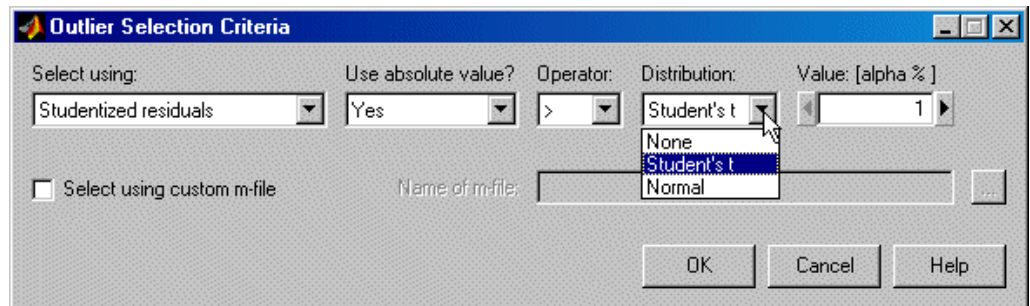
or restore all available points. You can also press **Ctrl+Z**. Select points in the left list and use the buttons to move points between the lists. When you click **OK** this refits the local model, including all data points previously removed as outliers. Use **Update Fit** once more to refit the global models, or it happens automatically when a new node is selected.

- **Copy Outliers** — Opens the Copy Outliers dialog, where you can choose which model's outliers to copy. Select a model (of the same type — local or global) in the tree and click **OK**, and the current model (and other models affected) are refitted using the outlier selections for that model.
- **Selection Criteria** — Opens the Outlier Selection Criteria dialog where you can set the criteria for the automatic selection of outliers. This is disabled for MLE models.
- **Remove All Data** — Leaves the current local model with no data, so entirely removes the current test. This test is removed from all the global models.

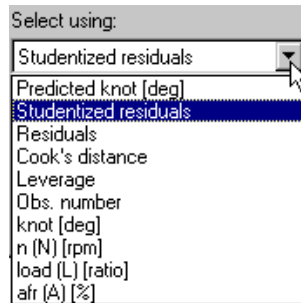
Outlier Selection Criteria

You can select outliers as those satisfying a condition on the value of some statistic (for example, $\text{residual} > 3$), or by selecting those points that fall in a region of the distribution of values of that statistic.

For example, assume that residuals are normally distributed and select those with $p\text{-value} > 0.9$. You can also select outliers using the values of model input factors.



The drop-down menu labeled **Select using** contains all the available criteria, shown in the following example.



The options available in this menu change depending on the type of model currently selected. The options are exactly the same as those found in the drop-down menus for the x - and y -axis factors of the scatter plots in the Model Browser (local level and global level views).

In the preceding example, the model selected is the knot response feature, so knot and Predicted knot appear in the criteria list, plus the global input factors; and it is a linear non-MLE model, so Cook's Distance and Leverage are also available.

The range of the selected criteria (for the current data) is indicated above the **Value** edit box, to give an indication of suitable values. You can type directly in the edit box. You can also use the up/down buttons on this box to change the value (incrementing by about 10% of the range).


Distribution

You can use the **Distribution** drop-down menu to remove a proportion of the tail ends of the normal or t distribution. For example, to select residuals found in the tails of the distribution making up 10% of the total area:

- Select Normal in the **Distribution** drop-down menu.
- Select the operator >.
- Enter 10 as the $\alpha\%$ value in the edit box.

Residuals found in the tails of the distribution that make up 10% of the total area are selected. If you had a vast data set, approximately 10% of the residuals would be selected as outliers.

As shown, residuals found beyond the value of Z_{α} in the distribution are selected as outliers. α is a measure of significance; that is, the probability of finding residuals beyond Z_{α} is less than 10%. Absolute value is used (the modulus) so outliers are selected in both tails of the distribution.



The t distribution is used for limited degrees of freedom.

If you select None in the **Distribution** drop-down menu, you can choose whether or not to use the absolute value. That is, you are selecting outliers using the actual values rather than a distribution. Using **absolute value** allows you to select using magnitude only without taking sign into account (for example, both plus and minus ranges). You can select No here if you are only interested in one direction: positive *or* negative values, above *or* below the value entered. For example, selecting only values of speed below 2000 rpm.

The **Select using custom M-file** check box enables the adjacent edit box. Here you can choose an M-file that selects outliers. Type the name of the file and path into the edit box, or use the browse button.

In this M-file you define a MATLAB function of the form:

```
function outIndices = funcname (Model, Data, Names)
```


Model is the current MBC model.

Data is the data used in the scatter plots. For example, if there are currently 10 items in the drop-down menus on the scatter plot and 70 data points, the data make up a 70 x 10 array.

Names is a cell array containing the strings from the drop-down menus on the scatter plot. These label the columns in the data (for example, spark, residuals, leverage, and so on).

The output, outIndices, must be an array of logical indices, the same size as one column in the input Data, so that it contains one index for each data point. Those points where index = 1 in outIndices are highlighted as outliers; the remainder are not highlighted.

Global Level

When you select a response feature node or one-stage model node in the model tree, this view appears. Both kinds of models have a global icon (), to reflect the fact that a global model is fitted, so this is referred to as *global level*. Plots shown here are referred to as global plots. Child nodes of these models also have global icons. Any node with a global icon produces this view.

For one-stage models, this view shows the functionality available at all model nodes. For two-stage models there are other levels with different functionality for local and response models.

See “Global Level: Toolbar” on page 4-99 and “Global Level: Menus” on page 4-103 for details on these controls.

This view is similar in format to the local level view, which also contains scatter plots above special plots. The statistical information panes on the right side are different, there is a **Removed Data** pane, and there are no test number controls. There is an edit box for model comments.

The display features are described in the following sections:

- “Global Special Plots” on page 4-95
- “Global Scatter Plots” on page 4-96
- Summary Table on page 4-97
- ANOVA Table on page 4-98
- “Removed Data Pane” on page 4-98
- “Models Comments” on page 4-98
- “Models List” on page 4-98

Model Browser - D:\MATLAB\Sandbox\MBC\matlab\toolbox\mbc\mbctraining\Gasoline_project.mat*

File Model View Outliers Window Help

All Models

- Gasoline_project
 - WT
 - BTQ
 - PS22
 - knot
 - ma
 - Bhi
 - Blo
 - EXTEMP
 - POLY2
 - Model type: Radial Basis Function Box Cox = 0

Response Feature: knot

Model type: Radial Basis Function Box Cox = 0

Studentized residuals

Predicted knot (transformed)

X-axis factor: Predicted knot (t... Y-axis factor: Studentized resi...

Predicted/Observed

Predicted knot

knot

Legend:

- Data
- Predicted = Observed
- * RBF center

Summary table

| | |
|--------------|-------|
| Observations | 188 |
| Parameters | 39 |
| Box-Cox | 0 |
| PRESS RMSE | 2.557 |
| RMSE | 2.274 |

Anova table

| Source | SS | df |
|------------|---------|-----|
| Regression | 1.376e6 | 38 |
| Error | 393.514 | 149 |
| Total | 1.377e6 | |

Removed data
* Data not restorable

Model comments

| Models | Observations | Parameters | Box-Cox | PRESS RMSE | RMSE |
|------------------|--------------|------------|---------|------------|--------|
| Quadratic | 188 | 10 | 1 | 4.2655 | 4.0951 |
| Cubic | 188 | 24 | 1 | 3.3424 | 3.0617 |
| RBF-multiquadric | 188 | 39 | 0 | 2.5569 | 2.2742 |

New Delete Select... 'Gasoline_project\WT\BTQ\PS22\knot\RBF-multiquadric' is the best model for knot

Ready

Click-and-drag anywhere along this line to change the size of the Models pane

Click-and-drag or double-click here to expand plots/hide statistics panes

Global Special Plots

The lower plots in the global level view are referred to as *special plots*, as they can be different for different models (for example, none at all for neural net models).

The special plot at the global level shows a Predicted/Observed plot. Where there is only one input factor, the plot shows the model fit and the data against the input factor (as in most local model special plots, which often have only one input factor).

For response feature models, each data point is the value taken by this response feature for some local model fit (of this two-stage model). Note that response features are not necessarily coefficients of the local curves, but are always derived from them in some way.

When there is more than one input factor it becomes impossible to display the fit in the same way, so the data for the response feature is plotted against the values predicted by the global model. The line of *predicted=observed* is shown. With a perfect fit, each point would be exactly on this line. The distances of the points from the line (the residuals) show how well the model fits the data.

To examine the fit in more detail, double-click the arrows (indicated in the figure in “Global Level” on page 4-93) to hide the scatter plot and expand the lower plot. You can also zoom in on parts of the plot by **Shift**-click-dragging or middle-click-dragging on the place of interest on the plot. Return to full size by double-clicking.


Note Right-click a point in either the special or scatter plot to open a figure plot of that particular test (for example, torque against spark).

You can change the lower plot from the Predicted/Observed to a Normal Plot by using the drop-down menu at the top of the plot. Normal plots are a useful graph for assessing whether data comes from a normal distribution. For more information, see “Normal Probability Plots” in the Statistics Toolbox documentation.

Global Scatter Plots

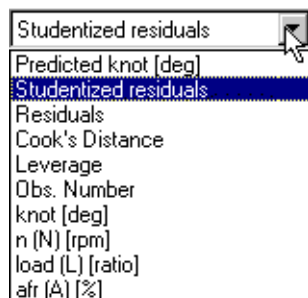
The upper plots in the global level view are referred to as *scatter plots*. They can show various scatter plots of statistics for assessing goodness-of-fit for the current model shown. The upper scatter plots are replaced by an icon if you resize the Browser too small.

The statistics available for plotting are model dependent.



The preceding are the context menus for both plots. On both plots you can manipulate outliers with all the same commands available in the **Outliers** menu except **Outlier Selection Criteria** and **Copy Outliers**. See “Outliers Menu (Local Level)” on page 4-88 for details. For one-stage models you can select **Plot Variables** on the scatter plot.

You can choose the x - and y -axis factors using the drop-down menus. The available statistics and factors are model dependent. Following is an example.



Shown is an example drop-down menu on the scatter plot for changing x and y factors. In this example knot is the response feature node selected. Therefore the model output is knot, so knot and Predicted knot are available in the menu. (For child nodes of knot, the model output is still knot.) The global inputs, the model output, and the predicted model output are always available in these menus. The observation number is also always available.

The other options are statistics that are model dependent, and can include: Residuals, Weighted Residuals, Studentized Residuals, Leverage, and Cook's Distance. These statistics and the other factors are also used as the available criteria for selection of outliers, so the options in the Outlier Selection Criteria dialog are similarly model dependent. At global (or one-stage) level these are externally studentized residuals.

Summary Table

| | |
|--------------|---|
| Observations | Number of observations used to estimate model |
| Parameters | Number of parameters in model |
| Box-Cox | Power transform used for box-cox transformation. A value of zero means a log transform is used. A value of 1 means there is no transformation. |
| PRESS RMSE | Root mean squared error of predicted errors. The divisor used for PRESS RMSE is the number of observations. The residuals are in untransformed values to enable comparison between alternative models with different Box-Cox transformations. |
| RMSE | Root mean squared error. The divisor used for RMSE is the number of observations minus the number of parameters. The residuals are in untransformed values, to enable comparison between alternative models with different Box-Cox transformations. |

ANOVA Table

| | SS | df | MS |
|------------|-----------|-----------|-----------|
| Regression | SSR | p-1 | SSR/(p-1) |
| Error | SSE | n-p | SSE/(N-p) |
| Total | SST | n-1 | |

This table is not present after you have calculated MLE. See also “Toolbox Terms and Statistics Definitions” on page 4-192.

Removed Data Pane

Tests you have removed here (using the **Remove Outliers** menu item) are shown in a list. Select **Outliers > Restore Removed Data** to select some or all of these to restore. Tests marked with an asterisk (*) are not restorable here. This means the entire test has been removed at the local level (using the **Remove All Data** menu item), or the local model could not be fitted. Removed tests can only be restored at the local level.

Double click on any removed test number to display a plot of the test in a figure window.

Models Comments

You can enter comments in the edit box.

Models List

Here is a list of any child models of the currently selected global model. This list is empty if there are no child nodes to compare. If there are child nodes you can click the **Select** button here to enter the Model Selection window to compare the child models and choose the best.

Click **New** to add a new child global model (or **Delete** to remove one). For more information see the “Test Plans List Pane” on page 1-13. The contents of this pane change in different views; it always contains the child nodes of the node selected in the model tree (and the **New**, **Delete**, and **Select** buttons). At the global level it contains a list of any child models.

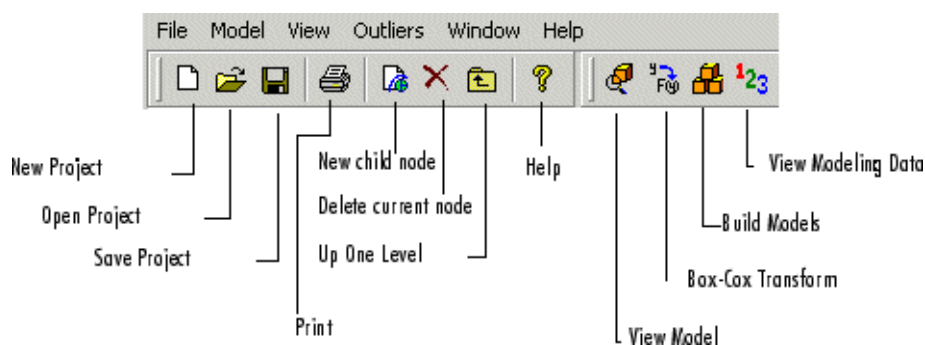
The list view displays the number of parameters and observations, the value of any Box-Cox transformation (1 indicates no transform), and the values of PRESS RMSE (for linear models only) and RMSE for each child model. Note that you can specify additional statistics to display here and in the Summary table by selecting **Model > Summary Statistics**. Use these statistics to compare the fit of different child models and help you choose the best. For definitions of RMSE and PRESS RMSE, see Summary Table on page 4-97 and “Model Selection Guide” on page 4-113. For other statistics see “Summary Statistics” on page 1-32 and “Model Menu (Global Level)” on page 4-104. For information on Box-Cox transforms, see “Box-Cox Transformation” on page 4-176.

Note that the Summary Statistics settings are inherited from parent global models and from the test plan level *only if* set before the current node was created, when they also appear in the bottom list pane if there are child nodes to compare.

Global Level: Toolbar

This toolbar appears when a response feature node or one-stage model node (both have a global icon) is selected in the model tree. Note that for one-stage models all model child nodes of the one-stage test plan are of this type.

Further buttons appear on the right depending on the type of model at the node selected. See “Global Level: Model-Specific Tools” on page 4-100.



The eight left icons remain constant throughout the levels. They are for project and node management, Help, and here the print icon is enabled, as

there are plots in this view. See “Project Level: Toolbar” on page 1-15 for details on these buttons.

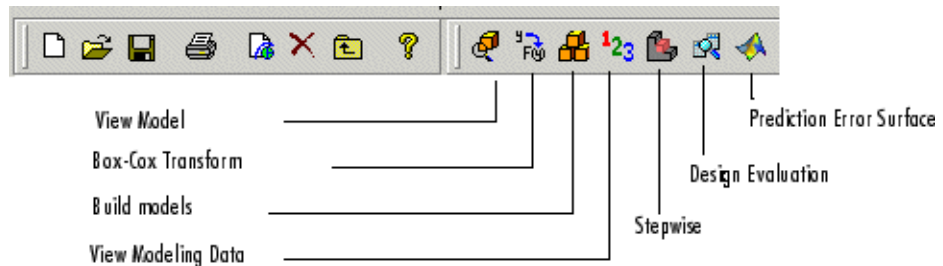
- **View Model** — Opens a dialog displaying the terms in the current model.
- **Box-Cox Transform** — Opens the Box-Cox Transformation plots, where you can minimize SSE to try to improve the fit. See “Box-Cox Transformation” on page 4-176 for statistical details.
- **Build Models** — Opens the “Build Models Dialog” on page 4-106, where you can choose a template for the type of models you want to build. There are predefined templates for polynomials, RBFs and hybrid RBFs, and free knot splines, and you can create new ones. You can also save templates of whatever models you choose using the **Model > Make Template** menu item. 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.
- **View Modeling Data** — Opens the Data Editor, where you can view a read-only version of the inputs, predicted and actual responses. This allows you to examine and export your modeling data, using all the powerful display features of the Data Editor. Also in the **View** menu.

These four toolbar icons appear for every global model node (although Box-Cox is not enabled for neural net models). The icons that appear to the right are model specific.

Global Level: Model-Specific Tools

All twelve left buttons (up to View Modeling Data) appear for all response feature models and one-stage models. See “Global Level: Toolbar” on page 4-99 for details on these buttons. The right buttons change according to model type.

Linear Model and Multiple Linear Models

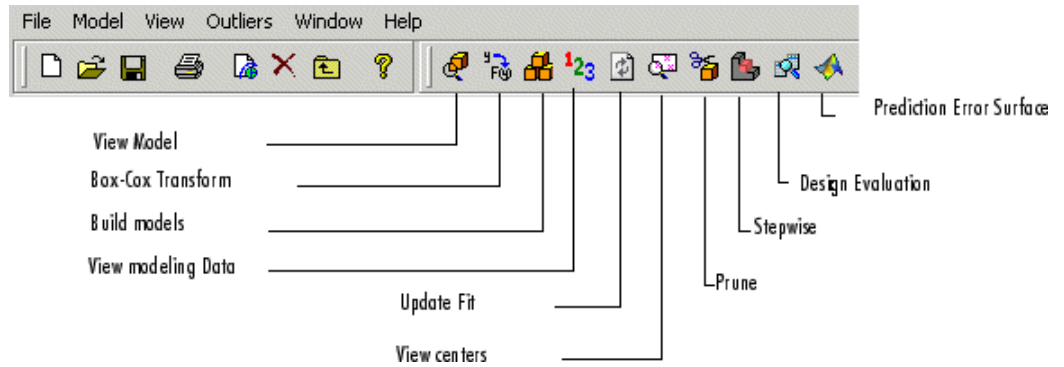


- Stepwise — This opens the Stepwise Regression window, where you can view the effects of removing and restoring model terms on the PRESS statistic (Predicted Error Sum of Squares), which is a measure of the predictive quality of a model. You can also use Min PRESS to remove all at once model terms that do not improve the predictive qualities of the model. See “Stepwise” on page 4-166 for further discussion of the statistical effects of the Stepwise feature.
- Design Evaluation — Opens the Design Evaluation tool, where you can view properties of the design. See “Design Evaluation Tool” on page 2-62.
- Prediction Error Surface - Opens the Prediction Error Variance Viewer. See “Prediction Error Variance Viewer” on page 2-54.

Free-Knot Spline Models

Free knot spline models do not have any model-specific tools, just the standard View Model, Box-Cox Transform, Build Models and View Modeling Data.

Radial Basis Function Models




- Update Fit refits the RBF widths and centers. See “Tips for Modeling with Radial Basis Functions” on page 5-34 and “Fitting Routines” on page 5-11 in the Radial Basis Functions chapter.
- View Centers opens a dialog where you can view the position of the radial basis function’s centers graphically and in table form.
- Prune opens the Number of Centers Selector where you can minimize various error statistics by decreasing the number of centers. See “Prune Functionality” on page 5-25.
- Stepwise opens the Stepwise Regression window.
- Design Evaluation - Opens the Design Evaluation tool, where you can view properties of the design. See “Design Evaluation Tool” on page 2-62.
- Prediction Error Surface - Opens the “Prediction Error Variance Viewer” on page 2-54.

Hybrid RBFs have the same toolbar buttons as linear models.

MLE Models

This toolbar appears when you select any response feature that is an MLE model (purple icon). See “Global Level” on page 4-93 for other functionality in this view.



At this point the New child node, Box-Cox and Build Models buttons are disabled.

- Recalculate MLE returns to the MLE dialog, where you can perform more iterations to try to refine the MLE model fit. See “MLE” on page 4-141 for more details.

Neural Networks

Neural net models have the View Model, Build Models, and Update Fit tools.

Global Level: Menus

File

Only the **New** (child node) and **Delete** (current node) functions change according to the node level currently selected. Otherwise the **File** menu remains constant.

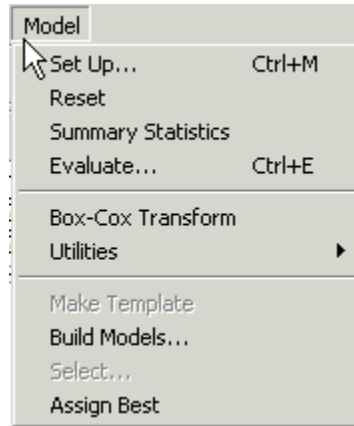
See “File Menu” on page 1-16.

Window and Help Menus

The **Window** and **Help** menus have the same form throughout the Model Browser.

See “Window Menu” on page 1-18 and “Help Menu” on page 1-18.

Model Menu (Global Level)



- **Set Up** opens the Global Model Setup dialog, where you can change the model type. See “Global Model Setup” on page 4-6.
- **Reset** returns to the global model defaults, that is, the global model specified at the test plan stage, restoring any removed outliers and removing any transforms.
- **Summary Statistics** opens the Summary Statistics dialog where you can select which statistics to display to help you evaluate models. In the global level view these appear in the **Summary Table** to the right of the plots, and in the **Models** list pane if there are child node models to compare. They also appear in the Model Selection window, and can be used to automatically select the best child node when using the Build Models dialog and Local Multiple Models. The standard summary statistics are PRESS RMSE (for linear models only) and RMSE, and these are always displayed. You can choose additional statistics for display in the Summary Statistics dialog by selecting the check boxes. You can also reach the dialog from the test plan. When you create child nodes the selected summary statistics are inherited. See “Summary Statistics” on page 1-32 for more information.
- **Evaluate** opens the Model Evaluation window.
- **Box-Cox Transform** opens the Box-Cox Transformation plots, where you can minimize SSE to try to improve the fit. See “Box-Cox Transformation” on page 4-176 for statistical details.

- **Utilities** opens a submenu showing the model-specific options available, duplicating the model-specific toolbar buttons (for example, Stepwise, Design Evaluation, View Centers, Prediction Error Surface, and so on).
- **Make Template** is available when child nodes exist. This opens a file browser where you can choose to save all the current child node models as a template, which you can then access using the Build Models menu item or toolbar button.
- **Build Models** opens the Build Models dialog. Here you can create a template or use a predefined template to build a selection of models as child nodes of the current node. The best model of this selection of child nodes will be automatically selected (it will have a blue icon), based on the selection criteria you choose in the following dialog (such as PRESS RMSE, RMSE, Box-Cox, Observations or Parameters).

See “Build Models Dialog” on page 4-106 for details.

- **Select** is available whenever the **Select** button is also enabled in the lower right pane (when it is titled **Local Models, Response Features, or Models**). This item opens the Model Selection window to allow you to choose the best model. See “Select Button” on page 4-111.
- **Assign Best** selects the current model as best. If it is one of several child node models of a global model, selecting it as best means that it is duplicated at the parent global model. See “Model Tree” on page 1-19.

View Menu (Global Level)

- **Model Definition** opens the Model Viewer dialog displaying the model terms.
- **Modeling Data** opens the Data Editor to display a read only version of the inputs and predicted and actual responses. Here you can view and export your modeling data. Also in the toolbar.
- **Test Numbers** turns test numbers on and off for both the special and scatter plots. Also available in the right-click plot menus.


Outliers Menu (Global Level)

This is the same as the local level **Outliers** menu, except that there is no **Remove All Data** command. All items are duplicated in the right-click

context menu on the plots, except **Selection Criteria** and **Copy Outliers**. See “Outliers Menu (Local Level)” on page 4-88.

At global level, as at local level, the **Restore Removed Data** item opens the Restore Removed Data dialog, where you can choose the points to restore from the list, or restore all available points. Select points in the left list and use the buttons to move points between the lists. Note that entire tests removed at the local level (using the **Remove All Data** item) cannot be restored at global level.


Build Models Dialog

The option **Build Models** () is in the toolbar and the **Model** menu for global models. You can use this to quickly build a selection of models as child nodes of the current node. You can use predefined templates, create templates within the Build Models dialog, or create templates from a global model with existing child nodes by selecting **Model > Make Template**.

Once you have used the Build Models dialog to create a number of child nodes, the best model of this selection will be automatically selected (it will have a blue icon), based on the selection criteria you choose (such as PRESS RMSE, RMSE, Box-Cox, or Parameters).

You can use the following templates:

- “New Template” on page 4-107
- “User-Defined Template” on page 4-108
- “Polynomial Template” on page 4-108
- “RBF Template” on page 4-108
- “Hybrid RBF Template” on page 4-109
- “Free Knot Spline Template” on page 4-110



See the Quick Start tutorial section “Creating Multiple Models to Compare” in the Getting Started documentation for a step by step example of using the Build Models dialog to quickly create many models to search for the best fit.

New Template

To create a template containing the models you want to build,

- 1 Click **New** then click **OK**.
- 2 Click **Add** in the following Multi-Model Settings dialog to reach the Model Setup dialog, where you can select any model type available for the number of inputs. Click **OK** to add the model and return to the Multi-Model Settings dialog.
- 3 Click **Add** again to repeat the process to add as many different models as you like. You can click **Edit Model** to change the settings for any models in the list.
- 4 When you are satisfied, click **OK** in the Multi-Model Settings dialog, and your chosen selection of model types are built as a selection of child nodes of the currently selected model node.
- 5 The Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node. Use the drop-down menu

to select from the available criteria (such as from PRESS RMSE, RMSE, Box-Cox, Observations or Parameters). You can select additional criteria to appear here using the Summary Statistics options, from the **Model** menu. See “Model Menu (Global Level)” on page 4-104 and Summary Statistics. Click **OK** to accept the chosen criterion.

User-Defined Template

You can save templates of models you have already built by using the **Model > Make Template** menu item from any global modeling node with child nodes. 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. Select your template and click **OK**. The models are built and the Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node.

Polynomial Template

You can use the polynomial template to build a variety of polynomials of different orders.

- 1** Select **Polynomials** and click **OK**. The Model Building Options dialog opens where you can specify the model settings.
- 2** Choose the minimum and maximum order of the polynomials you want to build, and whether to use Stepwise settings. For example, if you choose 1 and 5 as the minimum and maximum polynomial order, 5 child node models are built (linear, quadratic, cubic, and so on). If you choose a Stepwise setting (e.g. Minimize PRESS) it is applied to all child models.
- 3** Click **Build** and the models are built. The Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node.

RBF Template

You can use the RBF template to build a variety of radial basis function models with varying maximum numbers of centers and/or different kernels.

- 1** Select **RBF** and click **OK**. The Model Building Options dialog opens where you can specify the model settings.

- 2** Enter a vector in the edit box to specify the maximum numbers of centers for each child model. This can be a MATLAB expression including the number of observations, e.g. 10:10:nObs/2.
- 3** Click **Model Settings** to open the Radial Basis Function Model Settings dialog, where you can change all the model parameters such as kernel and widths. See “Types of Radial Basis Functions” on page 5-5.
- 4** Select the check box **Build all kernels** to create child models with the specified range of centers for each kernel type.
- 5** Click **Build** and the models are built. The Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node.

Hybrid RBF Template

You can use the Hybrid RBF template to build a variety of hybrid radial basis function models with varying maximum numbers of centers and/or different kernels.

- 1** Select **Hybrid RBF** and click **OK**. The Model Building Options dialog opens where you can specify the model settings.
- 2** Enter a vector in the edit box to specify the maximum numbers of centers for each child model. This can be a MATLAB expression including the number of observations, e.g. 10:10:nObs/2.
- 3** Click **Model Settings** to open the Hybrid RBF Model Settings dialog, where you can change all the model parameters such as kernel and widths (and the order of the polynomial part of the model on the **Linear Part** tab). See “Hybrid Radial Basis Functions” on page 5-32.
- 4** Select the check box **Build all kernels** to create child models with the specified range of centers for each kernel type.
- 5** Click **Build** and the models are built. The Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node.

Free Knot Spline Template

You can use the Free Knot Spline template to build a variety of free knot spline models of different numbers of knots. Only available for models with a single input factor.

- 1** Select **Free Knot Spline** and click **OK**. The Model Building Options dialog opens where you can specify the model settings.
- 2** Choose the initial and final number of knots. For example if you specify the initial and final numbers of knots as 1 and 5, five child nodes are built, one with one knot, one with two and so on.
- 3** Click **Build** and the models are built. The Model Selection dialog appears, where you can select the criterion for automatically choosing the best child node.

Selecting Models

The Model Selection window appears when you click the **Select...** button. This window is intended to help you select a best model by comparing several candidate models. For information on how to select best models, see “Model Selection Guide” on page 4-113.

The Model Selection window allows visual comparison of several models in many ways, depending on the type of model:

- “Tests View” on page 4-122
- “Predicted/Observed View” on page 4-125
- “Response Surface View” on page 4-127
- “Likelihood View” on page 4-131
- “RMSE View” on page 4-133
- “Residuals View” on page 4-135
- “Cross Section View” on page 4-136

Select Button

The **Select** button is under the list view in the pane at the bottom of the Model Browser display. This pane is the Test Plans list pane at startup and changes title depending on the level in the model tree that is selected. The list box in this pane always contains the child nodes of whichever node in the tree is selected.

The pane also always contains three buttons: **New**, **Delete**, and **Select**.

Select is only available when the lower pane lists local models, response models, or models.

You can select among the following:

- Local models
- Response features
- Submodels of response features

However, you cannot select between response models or test plans.

Note To get a two-stage model at the response node, you must use the **Select** button at the local node (when the lower pane is the **Response Features** pane) to assign a model (even if it is the only one) as best. This step then calculates the two-stage model.

Select might not be available if you are not ready to choose among the child nodes. For example, at the response node, the child nodes must have models assigned as best (using the **Select** feature at those levels) before you can select among them. Also, if a response feature has child nodes of alternate models, you must select the best, or the Browser cannot tell which to use to calculate that response feature. After calculating MLE, **Select** compares the MLE model with the previous univariate model, and you can choose the best.

The Model Selection window allows visual comparison of several models. From the response level you can compare several two-stage models. From the local level, if you have added new response features you can compare the different two-stage models (constructed using different combinations of response feature models). If you have added child nodes to response feature models, you can compare them all using the Model Selection window.

When a model is selected as best it is copied up a level in the tree together with the outliers for that model fit.

A tree node is automatically selected as best if it is the only child, except two-stage models which are never automatically selected - you must use the Model Selection window.

If a best model node is changed the parent node loses best model status (but the automatic selection process will reselect that best model if it is the only child node).

Note You can bypass the Model Selection window if you want to quickly assign a model as best (unless you need to calculate a two-stage model). Select the model node in the tree, then select **Model > Assign Best**.

Model Selection Guide

First it is important to point out that there is no recipe for model selection. It is not possible to cover the entire topic of Model Selection in a few paragraphs. Instead, we outline some general guidelines which should be helpful in using the Model-Based Calibration Toolbox to choose the best model for a given data set. There are many books you can go to for a fuller account of statistical modeling; see “References” on page 4-117.

Overfitting and Underfitting

When fitting a model to noisy data, we effectively make the fundamental assumption that the data have been generated from some model (the “truth”) by making predictions at given values of the inputs, then adding some amount of noise to each point, where the noise is drawn from a normal distribution with an unknown variance.

Our task is to discover both this model and the width of the noise distribution. In doing so, we aim for a compromise between *bias*, where our model does not follow the right trend in the data (and so does not match well with the underlying truth), and *variance*, where our model fits the data points too closely, and so “chases” the noise rather than trying to capture the true trend. These two extremes are known as underfitting and overfitting.

An important concept in this context is the *number of parameters* in a model. As this number increases, the model can bend in more complicated ways. If the number of parameters in our model is larger than that in the truth, then we risk overfitting, and if our model contains fewer parameters than the truth, we could underfit.

RMSE


Our basic measure of how closely a model fits some data is the Root Mean Squared Error (RMSE), 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. When two-stage modeling, use the RMSE Explorer to quickly investigate the local models with highest RMSE.

The smaller the RMSE, the closer our model follows the data; if a model goes through each data point exactly, then the RMSE is zero. The illustration

shows how increasing the number of parameters in the model can result in overfitting. The 9 data points (shown as black circles) are generated from a cubic polynomial which contains 4 parameters (the “truth”, shown as the black curve) by adding a known amount of noise. We can see that by selecting candidate models containing more parameters than the truth, we can reduce, and even eliminate, any mismatch between the data points and our model, causing the RMSE to vanish. This latter case occurs when the number of parameters in the model is the same as the number of data points (an 8th order polynomial has 9 parameters).

This does not mean that we have obtained a good fit - the model is overfitting, as we can see from the large difference between the model and the truth in the regions between the data points. By forcing our model to go through all the data points, we have included too much structure in the curve, which reduces the quality of the fit away from the data points.

Similarly, if we use a model with fewer parameters than in the truth, we risk underfitting; our model is not flexible enough to match the truth well. This is shown in the following illustration.



Other Statistics

As illustrated above, relying solely on RMSE can result in overfitting, which leads to poor model performance away from regions containing data points. In general, this problem is tackled by replacing RMSE with some other statistic, which also must be reduced to improve the fit, but which is designed to rise when we start overfitting. This is why you should consider RMSE *and* another tool such as the PRESS statistic to help you decide on the best model.

PRESS RMSE (Predicted sum of squares) is calculated in a similar way to RMSE, except we remove a data point from the fit, and ask the model to predict where that point lies with no knowledge of the data in that area. To

calculate PRESS RMSE, this process is repeated for each point in the data set and the results are averaged. If the value of PRESS RMSE is much bigger than the RMSE then we are overfitting. Weighted PRESS and GCV are also derived from this idea.

A different approach to solving the problem of overfitting results in statistics known as “Information Criteria”, such as AIC and BIC. Here, we combine a term involving RMSE with a term that rises with the number of parameters in the model. This explicitly penalizes a model for an increase in its level of structure. Try to minimize the information criteria statistics. Both AIC and BIC are approximations, which get more accurate as the number of observations increases. In general, we do not recommend using them unless the ratio of the numbers of observations to parameters is greater than 40:1 (see Section 2.4 of Burnham and Anderson, “References” on page 4-117). AICc, however, can be used with smaller samples and is the most appropriate information criterion for most problems in engine calibration.

The absolute value of AICc for a given model includes an arbitrary constant, and so is of no direct use. However, the difference between the AICc value for two models is meaningful: one rule of thumb says that if this difference is greater than about 10, then the worse model can be neglected in the selection process (see Section 2.6 of Burnham and Anderson, “References” on page 4-117).

Validation

All of the statistics mentioned above attempt to yield a model which makes good predictions both at the data points, and in the regions in between the data points. The simplest way to confirm that this is the case, is to collect additional data and test (or “validate”) the model against this new data, by evaluating the new RMSE based on these data. Comparing a validation RMSE with the RMSE based on the modeling data is a good model selection statistic. You can use the Model Evaluation window to validate models against other data, see “Model Evaluation Window” on page 4-148.

How much validation data to collect (or whether it is feasible to collect any at all) are matters governed primarily by practical considerations.

Trends

Throughout this discussion, we have focussed on using statistics for model selection. It is advisable, however, to combine a study of the model statistics with a careful examination of the trends present in the models. It would be a mistake to underestimate the importance of using engineering knowledge as a tool for comparing models. In addition, if two or more models of a different type (e.g. two different RBF kernel functions) follow the same trend, then that lends confidence to those models, because they are likely to be picking up real structure in the data. You can use the cross-section view in the Model Selection window to plot multiple models on the same axes to aid this process.

Where to Find Statistics For Comparing Models

- Look for RMSE and PRESS RMSE values in the model lists at the bottom of the Model Selection window and the Model Browser local, global and response views. Use these to compare models.
- To use other statistics such as information criteria (e.g. AICc) to compare models, you can specify extra statistics at the test plan and global levels. From any global model node select **Model > Summary Statistics**, and use the check boxes to select additional statistics to display. The choices you make in the Summary Statistics dialog determine what summary statistics are displayed in the global model views and the Model Selection window. These choices also determine what criteria you can use for automatically selecting best models in the Local Model Setup dialog for Multiple Models, and the Build Models dialog. Note, you must select the summary statistics before creating the child models to compare. See “Summary Statistics” on page 1-32.
- Use the RMSE Explorer for local models; look for tests with high RMSE and investigate. Select **View > RMSE Plots**, or use the toolbar button.
- Use the Model Evaluation window for validating against other data and you will see a Fit RMSE reported alongside the validation RMSE for comparison.

References

- Classical Regression: Draper and Smith, Applied Regression Analysis (3rd edition), John Wiley & Sons 1998

- BIC Methods: Kass and Raftery, Bayes factors. *Journal of the American Statistical Association*, 90(1995):773-795.
- AIC Methods: Burnham and Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (2nd edition), Springer-Verlag 2002.

Using Information Criteria to Compare Models

There are information criteria available as additional summary statistics for comparison of models. See “Summary Statistics” on page 1-32 for information on how to display these. This section provides some statistical background to help you compare the Akaike Information Criteria (AIC and AICc) and the Bayes Information Criterion (BIC). See also “Model Selection Guide” on page 4-113 for practical guidelines on using AIC and BIC.

AIC-type criteria are based on the difference in Kullback-Leibler information between two models, or their K-L distance. K-L distance is an appealing measure because it essentially compares the information content of two curves, by calculating the entropy in each. Akaike and others found ways to estimate K-L distance based on the results of a maximum likelihood estimate of the parameters of a model, given some data. These estimates are the information criteria, and become more accurate as the sample size increases.

BIC is derived from Bayes’ theorem, and essentially just applies the Occam effect to select a preferred model; the idea that if two models provide an equally good fit with some data then the simpler model is the likelier. This can be understood in the following sense: for models with greater complexity (both in terms of the number of parameters and the set of values those parameters can take) it is less remarkable that they are able to fit a given data set well. Conversely, for a simple model, if you happen to encounter a data set for which the model provides an acceptable fit, it would seem a lucky coincidence. Therefore, for data matching both models well, the odds are that the simpler one is closer to the truth.

Quantifying these ideas leads to Bayes factors (evidence ratios) which measure the relative probabilities of two models. In the context of MBC, BIC is an estimate of Bayes factors based on the results of a maximum likelihood estimate, and, like AIC, increases in accuracy in the limit of large sample size. Although priors often spring to mind in the context of Bayes theorem, all of

the above can proceed with uniform priors on everything, and the Occam effect still applies.

There is a degree of controversy over which approach gives the best results. Copious literature exists on the subject of Bayesian model selection, a smaller amount on K-L distance based techniques and a still smaller amount comparing the two approaches. Bayesian authors consistently find that BIC performs better in Monte Carlo simulations (e.g. Leonard and Hsu 1999) whereas Burnham and Anderson 2002 (the main proponents of K-L distance techniques) reach conclusions which favour AIC.

Such tests can be set up to favour either criterion, and there are two main effects relevant to understanding this. Differences arise due to the assumptions made about the truth in each case (relevant to choosing Bayes factors or K-L distance), and due to the number of samples relative to the number of parameters in the candidate models.

Regarding the former effect: Bayes factors always seek the simplest model consistent with the data. K-L distance also has this tendency, although not as strongly as the Bayesian approach. As a result, the simulations in which BIC does well tend to be based on simple models with few parameters (Leonard and Hsu choose a simple quadratic, then consider polynomials of order 1 to 7 in their candidate set). Although both approaches choose the correct model more often than any other, AIC gives slightly more weight to the higher order models than does BIC. For this reason, Bayesians often accuse AIC of overfitting.

Burnham and Anderson, however, are biologists and as such they abandon all hope of actually finding the true model in their candidate set - they simply attempt to find the best approximation to the truth. A typical simulation of theirs considers linear models with up to 13 possible variables in the context of predicting body fat. They are not concerned with the subtle shape of curves, only with which variables they can safely throw away. In this scenario, they find that BIC favours too simple a model and hence underfits the data.

Although for BIC most authors assume that the true model is contained within the candidate set, this is not necessary for model comparison — it just concerns the normalisation of the probabilities, and hence not the ratios that form the Bayes factors.

AIC and BIC both improve as estimators of their respective statistical measures as the sample size increases, with relative errors of $O(n^{-1})$, where n is the sample size. AIC is obtained from a first order Taylor expansion, and AICc is a second order correction to that for the special case of Gaussian Likelihood (there is no general second order correction) and should be used when the ratio of data samples to model parameters (in the largest model for nested sets) is less than about 40:1. For very small sample sizes, even Bayesian authors do not seem to trust BIC, but do consider AICc.

In terms of the complexity of the truth, most problems in MBC probably lie in between the two extremes described above: internal combustion engines are not so simple that we assume that our model set really contains the precise, closed-form solutions to the relevant dynamical equations, but we are dealing with a mechanical system, not trying to predict, for example, characteristics of the human body. In terms of the number of samples per model parameter, AIC is seldom likely to be a reliable statistic; AICc should be used instead. But if you have reason to prefer a more conservative estimate of the complexity of the model, BIC should be considered.

For a discussion of Bayes factors, see:

Kass and Raftery (1995). Bayes factors. *Journal of the American Statistical Association* 90, 773-795

See also Chapter 28 from the following book: *Information Theory, Inference, and Learning Algorithms*, available from

- <http://www.inference.phy.cam.ac.uk/mackay/itila/book.html>
- *Bayesian Methods*, Leonard and Hsu, Cambridge University Press 1999
- *Model Selection and Multimodel Inference*, Burnham and Anderson, Second Edition, Springer-Verlag 2002.

Model Selection Window

The Model Selection window comprises several different views depending on the type of models being compared:

- “Tests View” on page 4-122
- “Predicted/Observed View” on page 4-125

- “Response Surface View” on page 4-127
- “Likelihood View” on page 4-131
- “RMSE View” on page 4-133
- “Residuals View” on page 4-135
- “Cross Section View” on page 4-136

You can change to any available view in the Model Selection window using the **View** menu or by clicking the buttons of the toolbar.

The **Assign Best** button at the bottom of the window marks the currently selected model as best or you can double-click a model in the list.

Information about each candidate model is displayed in the list at the bottom. The information includes categories such as the number of observations and parameters, and various diagnostic statistics such as RMSE and PRESS RMSE. You can click on column headers in this list to sort models by that category — for example, clicking on the column header for PRESS RMSE sorts the models in order of increasing PRESS RMSE. As this statistic is an indication of the predictive power of the model, it is a useful diagnostic statistic to look at (the lower the better), but remember to also look at other factors.

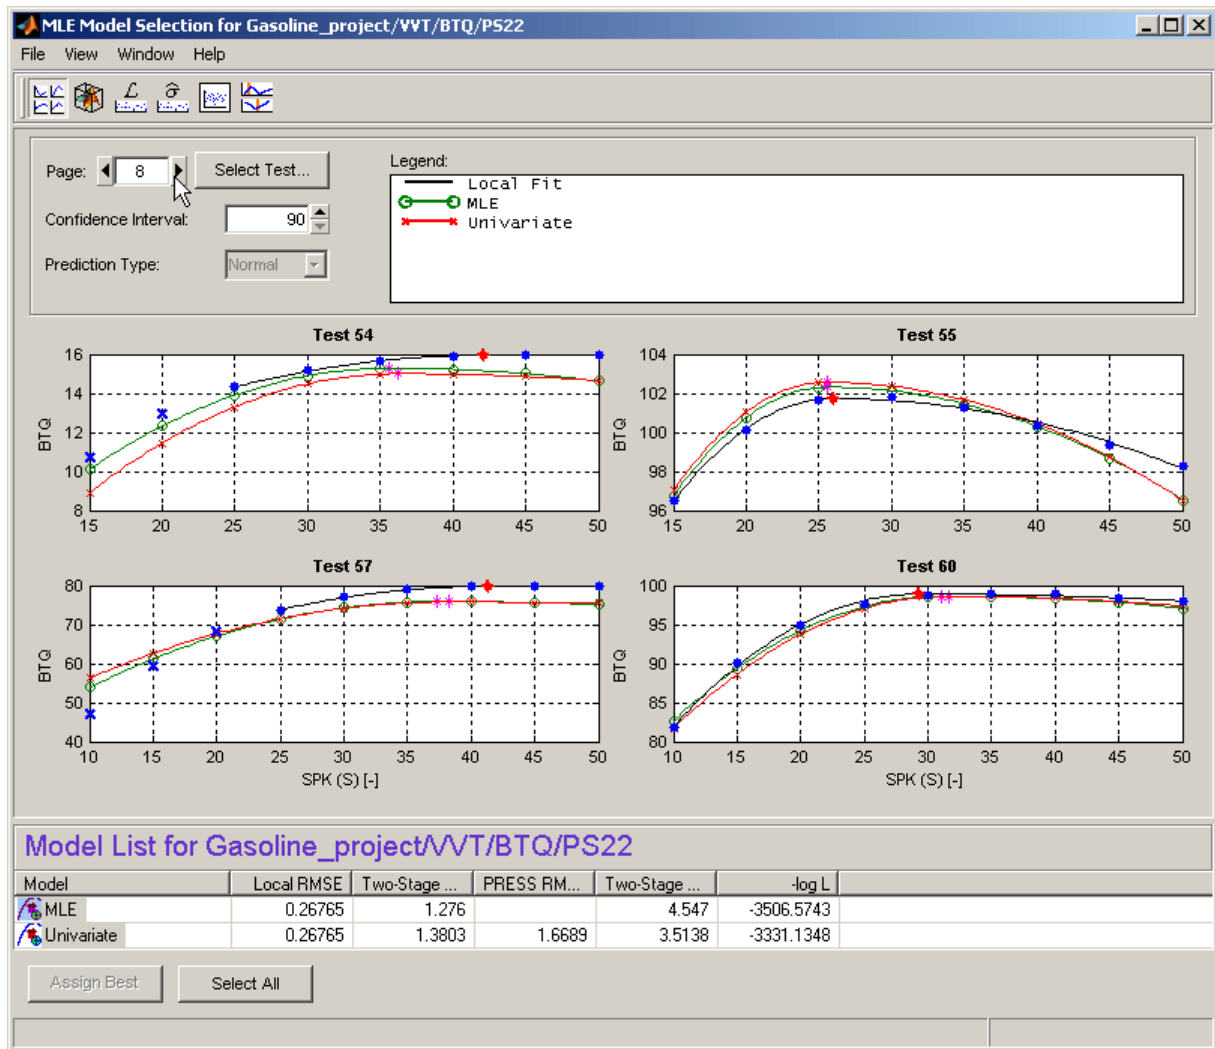
Note The diagnostic statistics displayed depend on your choices in the Summary Statistics dialog for the parent model node. See “Summary Statistics” on page 1-32 for details.

To print the current view, use the **File > Print** menu item or its hot key equivalent **Ctrl+P**. In the Response Surface view you can also use the right-click context menu.

To close the Model Selection window, use the **File > Close** menu item or its hot key equivalent **Ctrl+W**. This window is intended to help you select a best model by comparing several candidate models. On closing the figure, you are asked to confirm the model you chose as best.

Tests View

For a two-stage model the initial view is as follows:




The tests view shows the data being modeled (blue dots) and models that have been fitted to this data. The black line shows the local model that has been fitted to each test separately. The green line and red lines in this case show an

MLE two-stage model and the Univariate two-stage model: you can see the local model curves reconstructed using response feature values taken from the global models, and compare the fits.

This view allows you to compare several models simultaneously. Using standard Windows multiselect behavior (**Shift**+click and **Ctrl**+click) in the list view, or by clicking the **Select All** button, you can view several two-stage models together. A maximum of five models can be selected at once. The legend allows you to identify the different plot lines.

If the local input has more than one factor, a “Predicted/Observed View” on page 4-125 appears instead.

Clicking one of the plots (and holding the mouse button down) displays information about the data for that test. For example:

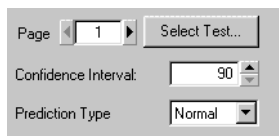


Here you see the values of the global variables for this test and some diagnostic statistics describing the model fit. Also displayed are the values (for this test) of the response features used to build this two-stage model and the two-stage model’s estimation of these response features.

The controls allow navigation between tests.

You can change the size of the confidence intervals; these are displayed using a right-click menu on the plots themselves.

The prediction type allows a choice of Normal or PRESS (Predicted Error Sum of Squares) — although not if you entered this view through model evaluation (rather than model selection). PRESS predictions give an indication of the model fit if that test was not used in fitting the model. For more on PRESS see “PRESS statistic” on page 4-174, Summary Table on page 4-97, and “Stepwise” on page 4-166.

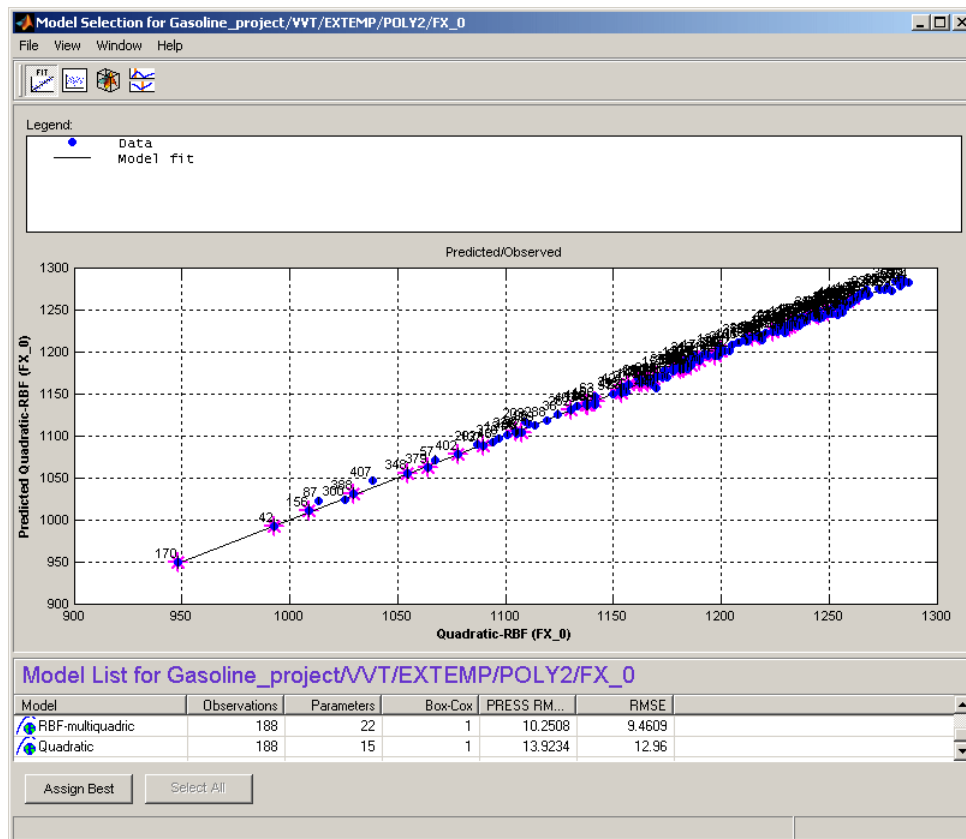


The image shows a software interface with three rows of controls:

- Row 1: A "Page" field with a left arrow, the number "1", and a right arrow, followed by a "Select Test..." button.
- Row 2: A "Confidence Interval:" label followed by a text box containing "90" and a vertical spinner control.
- Row 3: A "Prediction Type" label followed by a dropdown menu showing "Normal" and a downward arrow.

Predicted/Observed View

For a one-stage model, or when you are comparing different models for one Response Feature, the initial view is as follows:



The plot shows the data used to fit this model, against the predicted values found by evaluating the model at these data points. The straight black line is the plot of $y=x$. If the model fitted the data exactly, all the blue points would lie on this line. The error bars show the 95% confidence interval of the model fit.

For single inputs, the response is plotted directly against the input.

The Predicted/Observed view only allows single selection of models for display. Right-click to toggle test number display, as you can on most plots.

Response Surface View

This view shows the model surface in a variety of ways.

Model Selection for May2002/dualInd/tq/POLY2

File View Window Help

Input factors:

N: 2895.79
 L: 0.548762
 EXC Min: -5 Max: 50 Pts: 70
 ICA Min: -5 Max: 50 Pts: 70

Display using (S - datum)
 Display boundary constraint

Display type:
 Table
 2-D Plot

Prediction Error shading
 Prediction Error threshold: _____

Export model values:
 Export to workspace
 Export to mat file...

Save input as: X
 Save result as: Y

Export

X-axis factor: ICAMPH
 Y-axis factor: EXCAMPH

Model List for May2002/dualInd/tq/POLY2

| Model | Local RMSE | Two-Stage ... | PRESS RM... | Two-Stage ... | -log L |
|----------------------|------------|---------------|-------------|---------------|---------|
| FX_0 FX_plus10 FX... | 5.0277 | 30.1265 | 21.267 | 4.2861 | 756.965 |

Assign Best Select All

Ready

The default view is a 3-D plot of the model surface, as in the example. This model has five dependent factors; you can see these in the controls at the top left (there is a scroll bar as only four can be seen at once at this size of window).

You can choose which input factors to display by using the drop-down menus below the plot. The unselected input factors are held constant and you can change their values using the controls at the top left of the view (either by clicking the arrow buttons or by typing directly in the edit box).

Display using (S - datum) — If a datum model is being displayed, this check box appears. The datum variable here is spark angle, S. When you select this box, the model is displayed in terms of spark angle relative to the datum. The appropriate local variable name appears here. See “Datum Models” on page 4-72.

Display boundary constraint — If you have boundary models you can display them by selecting the check box. Areas outside the boundary are yellow, as shown in the example. Areas outside the boundary are yellow (or gray in table view). They are shown on all display types (contour, 2-D, surface, movie and table).

Display Type— Changes the model plot. Display options are available for some of these views and are described under the relevant view. The choices are as follows:

- A table showing the model evaluated at a series of input factor values.
- A 2-D plot against one input factor.
- A 2-D plot with several lines on it (called a multiline plot); this shows variation against two input factors.
- A contour plot.

The **Contours.** button opens the Contour Values dialog. Here you can set the number, position, and coloring of contour lines.

Fill Contour colors each space between contours a different color.

Contour Labels toggles the contour value numbers on and off. Without labels a color bar is shown to give you a scale.

Auto (the default) automatically generates contours across the model range.

N Contour Lines opens an edit box where you can enter any number of contour lines you want.

Specify values opens an edit box where you can enter the start and end values where you want contour lines to appear, separated by a colon. For example, entering 5:15 gives you 10 contour lines from 5 to 15. You can also enter the interval between the start and end values; for example 1:100:600 gives you contour lines between 1 and 600 at intervals of 100.

- A surface (shown in the example).

Prediction Error shading — Colors the surface in terms of the prediction error ($\sqrt{\text{Prediction Error Variance}}$) of the model at each point. A color bar appears, to show the value associated with each color.

Note For datum models, Prediction Error shading is only available when the **Display using** (*local variable - datum*) check box is not selected.

Prediction Error threshold — To see good color contrast in the range of PE of interest, you can set the upper limit of the coloring range. All values above this threshold are colored as maximum P.E.

- A movie: this is a sequence of surfaces as a third input factor's value changes.
 - **Replay** replays the movie.
 - **Frame/second** selects the speed of movie replay.
 - The number of frames in the movie is defined by the number of points in the input factor control (in the array at the top left) that corresponds to the **Time factor** below the plot.

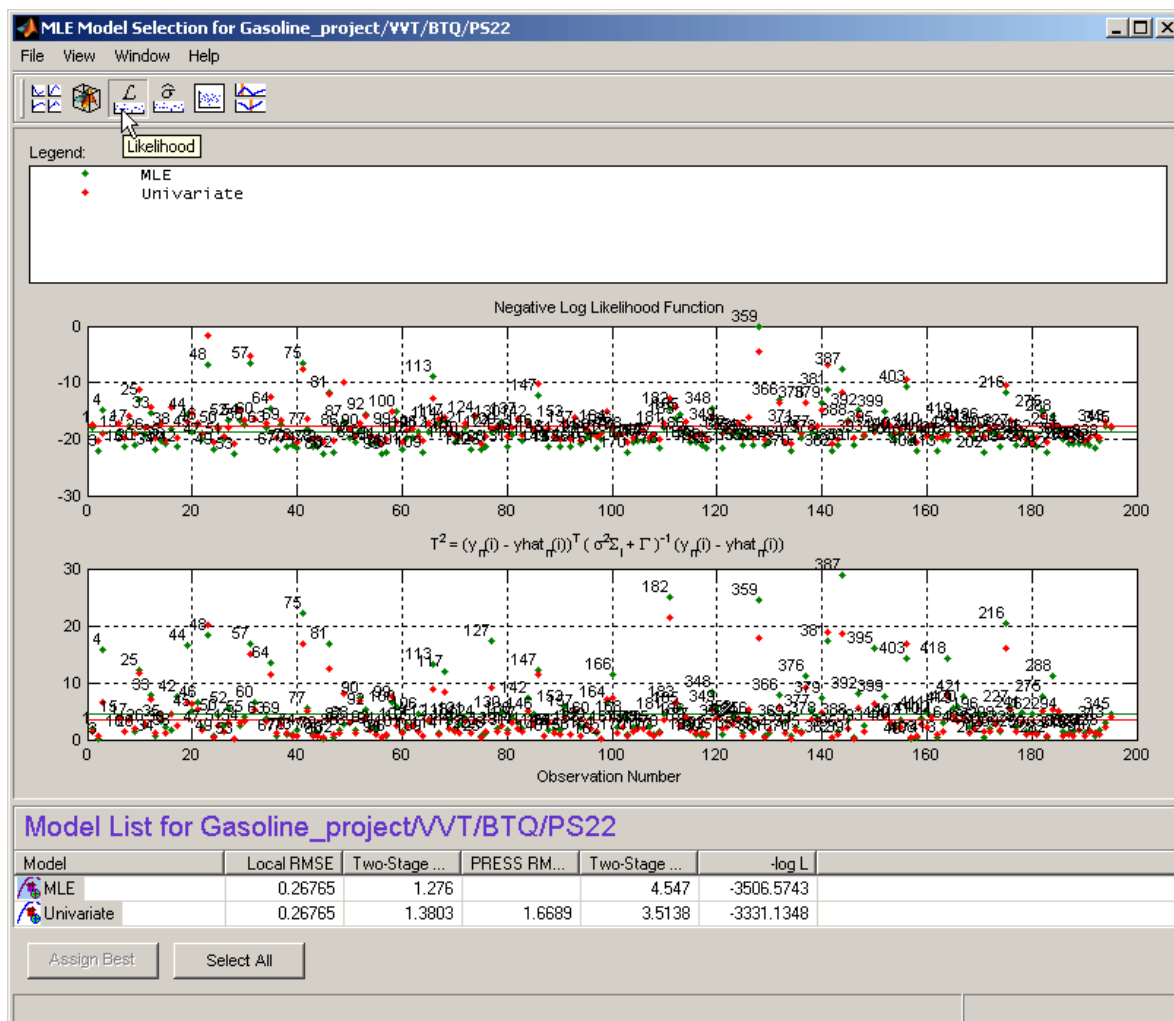
Export model values allows the currently displayed model surface to be saved to a MAT file or to the MATLAB workspace.

Right-click on the plot to reach the context menu and change many display properties (lighting, colormap etc.) and print to figure.

Within a test plan the memory is retained of the evaluation region, plot type and the number of points resolution last displayed in the Response Surface view.

Likelihood View

The likelihood view shows two plots relating to the log likelihood function evaluated at each test. It is useful for identifying problem tests for maximum likelihood estimation (MLE).



Each plot has a right-click menu that allows test numbers to be displayed on the plots and also offers autoscaling of the plots. You can also **Print to Figure**.

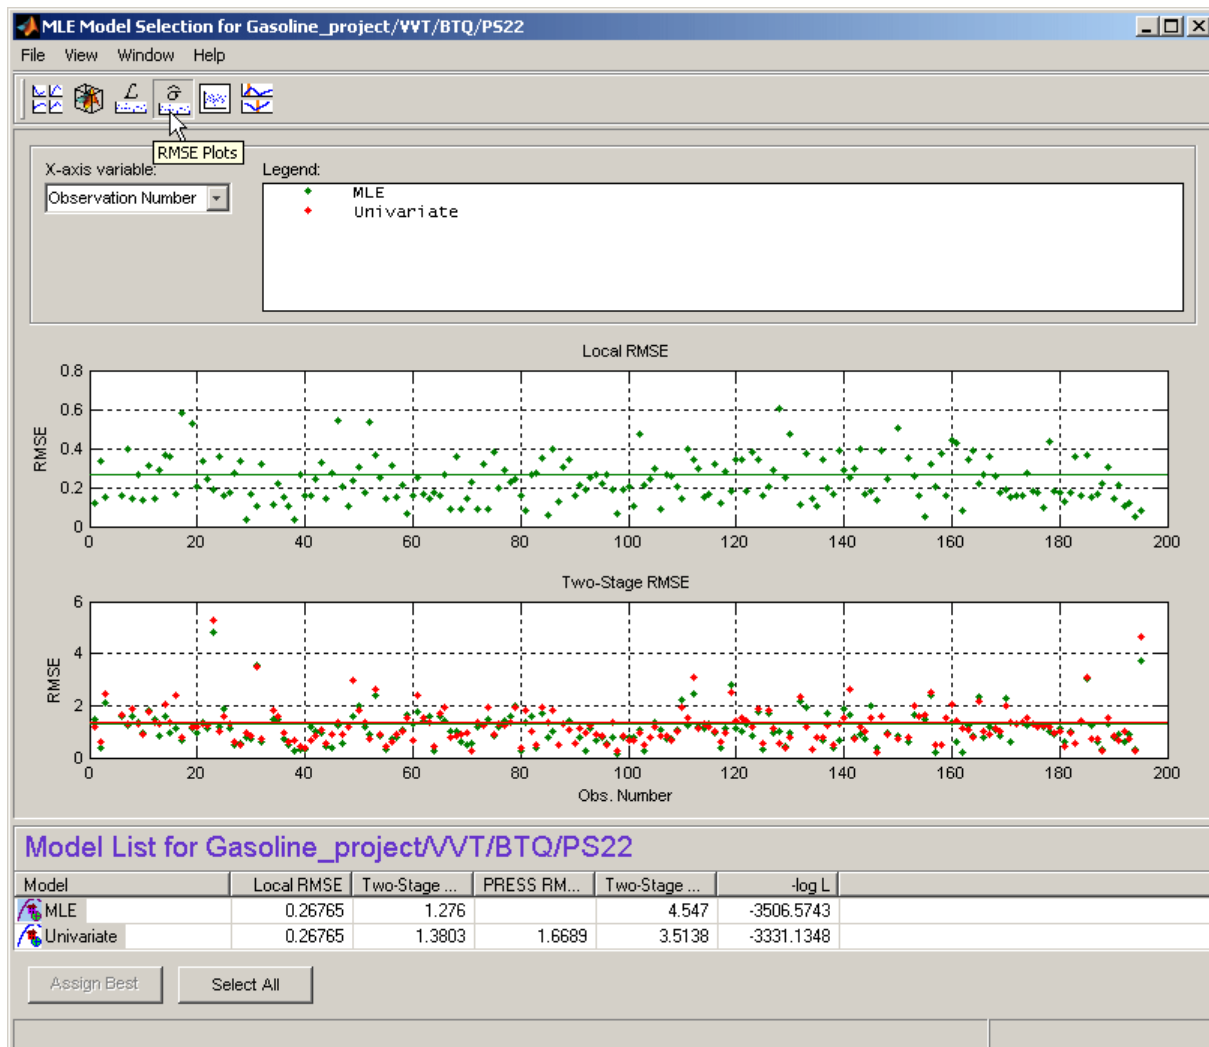
The likelihood view allows several models to be displayed simultaneously; click the **Select All** button at the bottom of the window or, in the model list view, **Shift**+click or **Ctrl**+click to select the models for display.

The upper plot shows values of the negative log likelihood function for each test. This shows the contribution of each test to the overall negative log likelihood function for the model, as compared with the average, as indicated by the horizontal green line.

The lower plot shows values of the T-squared statistic for each test. This is a weighted sum squared error of the response feature models for each test. As above, the purpose of this plot is to show how each test contributes to the overall T-squared statistic for this model. The horizontal line indicates the average.

RMSE View

The Root Mean Square Errors view has three different plots, each showing standard errors in the model fit for each test.



Each plot has a right-click menu that allows test numbers to be displayed on the plots, and you can **Print to Figure**.

The **X variable** menu allows you to use different variables as the *x*-axis of these plots.

The RMSE view allows several models to be displayed simultaneously; click the **Select All** button at the bottom of the window or, in the model list view, **Shift**+click or **Ctrl**+click to select the models for display.

Local RMSE shows the root mean squared error in the local model fit for each test.

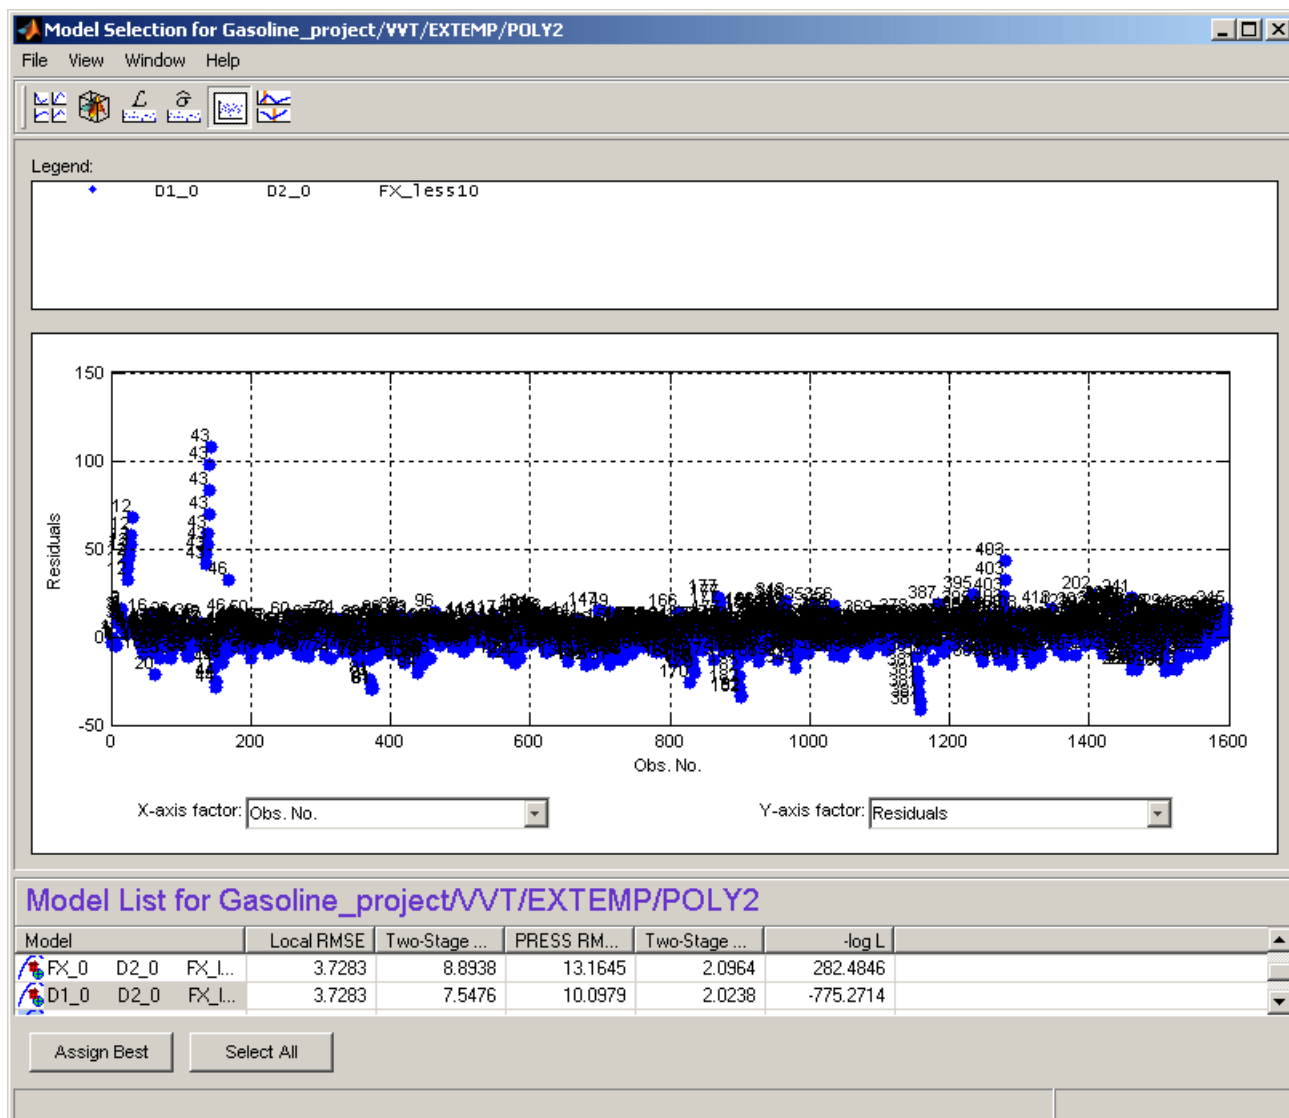
Two-Stage RMSE shows the root mean squared error in the two-stage model fit to the data for each test. You should expect this to be higher than the local RMSE.

PRESS RMSE is available when all response feature models are linear. This plot shows the root mean squared error in the PRESS two-stage model fit at each test.

For information on PRESS RMSE see Summary Table on page 4-97 and “Model Selection Guide” on page 4-113.

Residuals View

The residuals view shows the scatter plots of observation number, predicted and observed response, input factors, and residuals.



This view allows several models to be displayed simultaneously; click the **Select All** button at the bottom of the window or, in the model list view, **Shift**+click or **Ctrl**+click to select the models for display.

A right-click menu allows the test number of each point to be displayed when only one model is being displayed, as shown.

The **X-axis factor** and **Y-axis factor** menus allow you to display various statistics.

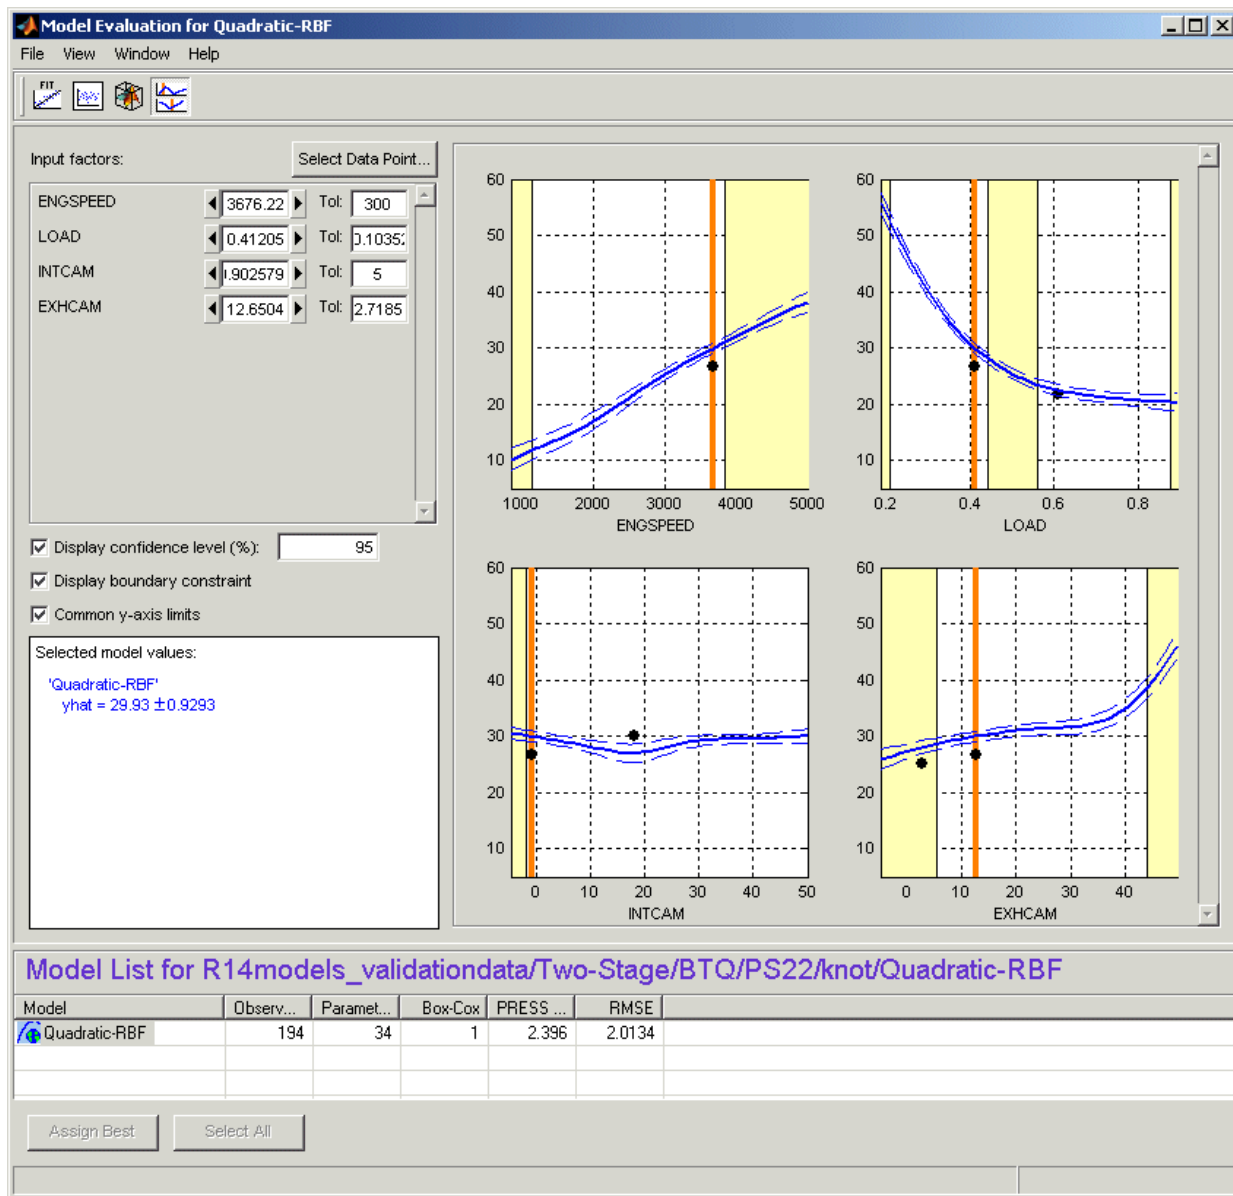
Cross Section View

The cross-section view shows an array of cross sections through the model surface. You can choose the point of cross section in each factor. Data points near cross sections are displayed, and you can alter the tolerances to determine how much data is shown. The only exception is when you evaluate a model without data; in this case no data points are displayed.

You can select individual data points by test number (using the **Select Data Point** button). You can double-click a data point in a graph to take the display directly to that point. You can choose to use a common Y-axis limit for all graphs using the check box.

If you have boundary models you can choose to display them here using the check box; regions outside the boundary are yellow, as shown in the example.

Within a test plan the memory is retained of the point last displayed in the Cross Section view; when you reopen the view you return to the same point.



The number of plots is the same as the number of input factors to the model. The plot in S shows the value of the model for a range of values of S while the other input factors are held constant. Their values are displayed in the controls at the top left, and are indicated on the plots by the vertical orange bars.

- You can change the values of the input factors by dragging the orange bars on the plots, using the buttons on the controls, or by typing directly into the edit boxes.
- For example, changing the value of N to 1000 (in any of these ways) does nothing to the graph of N, but all the other factor plots now show cross sections through the model surface at $N = 1000$ (and the values of the other variables shown in the controls).

On the plots, the dotted lines indicate a confidence interval around the model. You define the confidence associated with these bounding lines using the **Display confidence level (%)** edit box. You can toggle confidence intervals on and off using the check box on this control.

For each model displayed, the value of the model and the confidence interval around this are recorded in the legend at the lower left. The text colors match the plot colors. In the example shown, two models are selected for display, resulting in blue (PS22 model) and green (POLY2 model) legends on the left to correspond with the blue and green plots. You can select multiple models to display in the list at the bottom using **Ctrl**+click, or click **Select All**. The values of the input factors (for which the model is evaluated) can be found in the controls (in the **Input factors** pane) and seen as the orange lines on the plots.

Data points are displayed when they fall within the tolerance limit near each cross section. You can set the tolerance in the **Tol** edit boxes.

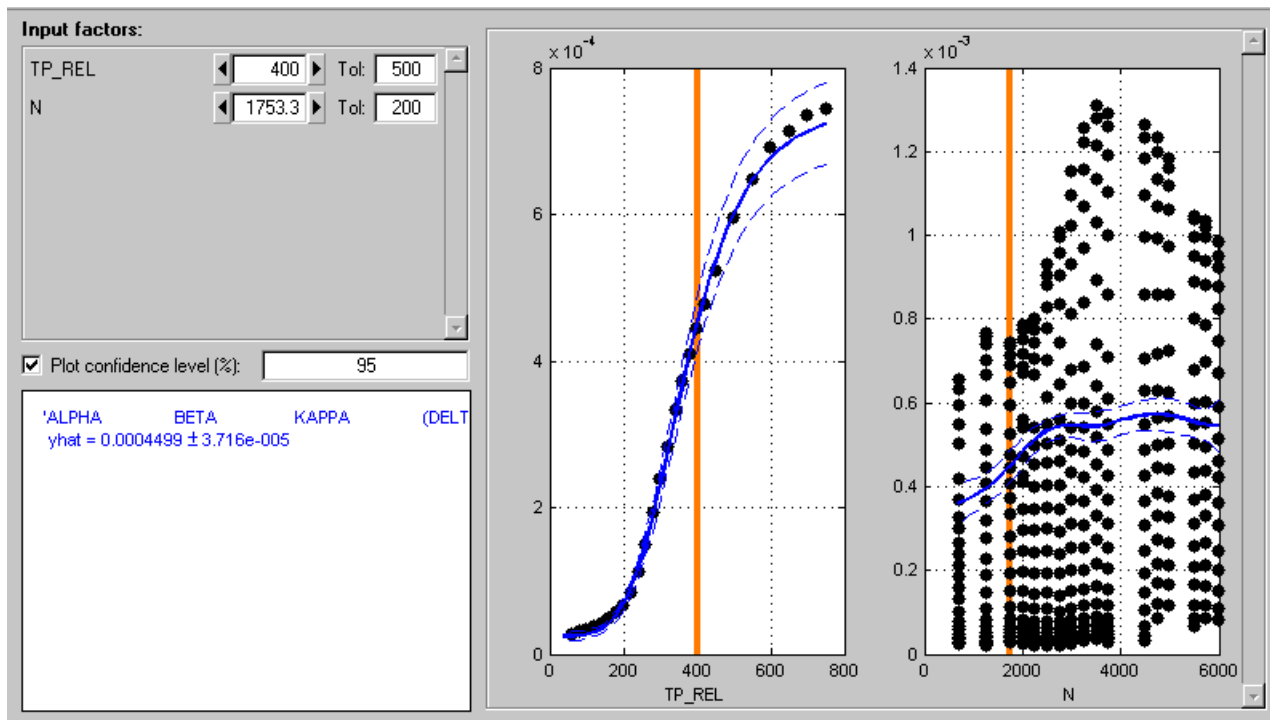
- For example, if N is set to 1000, and the tolerance for N is set to 500, all data points with values between $N = 500$ and $N = 1500$ appear on the plots of the other factors.
- This means that changing the tolerance in one factor affects the data points that appear on the plots of all the other factors. It does not affect the plots of that factor.

- You can click data points in the plots to see their values. Several points can mask each other; in this case the values of all coincident data points are displayed. Double-click to move the display directly to a data point.

The following example illustrates how the tolerance level determines which data points are displayed. The tolerance for TP_REL (500) includes all points in the data set (this is an extreme example). The plot for N therefore shows the data points for all the tests. Note that you can see the structure of the data as each test shows as a vertical line of points.

You can see that the orange line on the N plot passes through a test. This orange line shows the value of N for the cross-section plot of TP_REL. You can also read the value in the edit box (N=1753.3). The tolerance for N (200) only includes data points of this test. Data in adjacent tests fall outside this tolerance. Therefore the TP_REL plot shows the data points from one test only.

Increasing the tolerance on N will mean that more data points fall within the tolerance and so would appear on the TP_REL plot.




MLE

For an ordinary (univariate) two-stage model, the global models are 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.

Calculating MLE

When you close the Model Selection window, a dialog asks if you want to calculate MLE. If you click **Cancel** at this point, you can calculate MLE later as follows:

- 1 From the local node, click the MLE icon in the toolbar .

Alternatively, choose **Model > Calculate MLE**.

- 2 The MLE dialog appears. Click **Start**.

You can alter various MLE settings on this dialog.

- 3 After you click **Start** a series of progress messages appears, and when finished a new Two-Stage RMSE (root mean square error) value is reported.
- 4 You can perform more iterations by clicking **Start** again to see how the RMSE value changes, or you can click **Stop** at any time.
- 5 Clicking **OK** returns you to the Model Browser, where you can view the new MLE model fit.

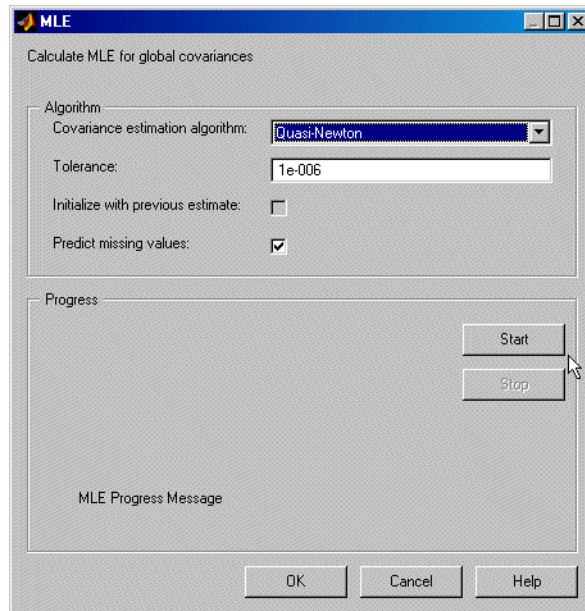
Note After calculating MLE, you will notice that the plots and the icons in the model tree for the whole two-stage model (response node, local node, and all response feature nodes) have turned purple. See “Icons: Blue Backgrounds and Purple Worlds” on page 1-22.

You can select all response features in turn to inspect their properties graphically; the plots are all purple to symbolize MLE. At the local node the plots show the purple MLE curves against the black local fit and the blue data.

- From the response feature nodes, at any time, click the **MLE** toolbar icon to recalculate MLE and perform more iterations.
- From the local node, you can click **Select** to enter the Model Selection window, compare the MLE model with the previous univariate model (without correlations), and choose the best. Here you can select the univariate model and click **Assign Best** to “undo” MLE and return to the previous model.

Note If there are exactly enough response features for the model, you can go straight to MLE calculation after model setup without going through the Select process. The MLE toolbar button and the **Model > Calculate MLE** menu item are both active in this case. If you add new response features, you cannot create MLE until you go through model selection to choose the response features to use.

MLE Settings



Algorithm

The algorithm drop-down menu offers a choice between two covariance estimation algorithms, Quasi-Newton and Expectation Maximization. These are algorithms for estimating the covariance matrix for the global models.

Quasi-Newton is recommended for smaller problems (< 5 response features and < 100 tests). Quasi-Newton usually produces better answers (smaller values of $-\log L$) and hence is the default for small problems.

Expectation Maximization is an iterative method for calculating the global covariance (as described in Davidian and Giltinan (1995); see References in “Two-Stage Models for Engines” on page 4-181). This algorithm has slow convergence, so you might want to use the **Stop** button.

Tolerance

You can edit the tolerance value. Tolerance is used to specify a stopping condition for the algorithm. The default values are usually appropriate, and if calculation is taking too long you can always click **Stop**.

Initialize with previous estimate


When you recalculate MLE (that is, perform more iterations), there is a check box you can use to initialize with the previous estimate.

Predict missing values

The other check box (selected by default) predicts missing values. When it is selected, response features that are outliers for the univariate global model are replaced by the predicted value. This allows tests to be used for MLE even if one of the response features is missing. If all the response features for a particular test are missing or the check box is unselected, the whole test is removed from MLE calculation.

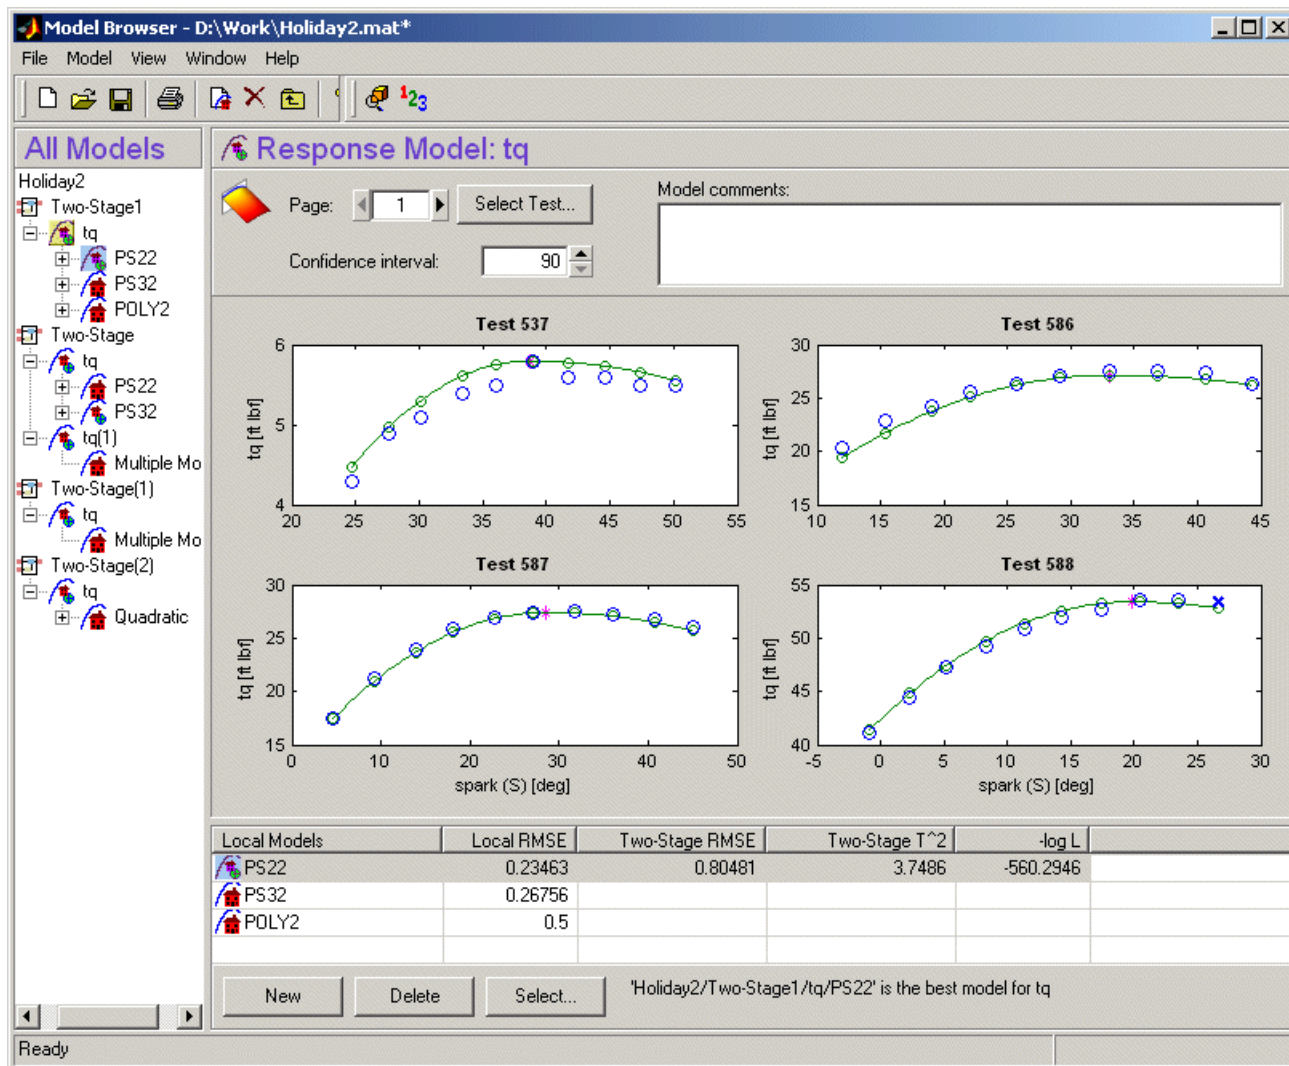
Response Level

Note The response node remains empty until you have used the Model Selection window at the local level. Exiting this window then copies the best two-stage model to the response node.

Selecting a response model node (with a two-stage icon  — a house *and* a globe) in the model tree produces this view.


Note that local model nodes also have the same icon *after* calculation of the two-stage model (see the model tree for clarification) but selecting them produces the local level view instead.

These plots show the data and the two-stage model fit. You can scroll through the tests using the test controls, as at the local level: by clicking the up/down page number buttons, typing directly in the edit box, or clicking **Select Test** to go directly to a particular test.



The Local Models list view shows all child models and summary statistics for comparison of child models. See “Pooled Statistics” on page 4-79 for information on the diagnostic statistics in the list (such as log likelihood and T^2). The statistics in this list can be seen in the Pooled Statistics table at the local model level.

Response Level: Toolbar and Menu



- View Model is the only toolbar icon after the standard project and node management and print buttons (see “Project Level: Toolbar” on page 1-15 for details on these buttons). View Model opens the Model Viewer dialog displaying the model terms.
- **File**, **Window**, and **Help** menus remain constant throughout the Model Browser. See “Project Level: Menus” on page 1-16.
- **Model > Evaluate** opens the Model Evaluation window, for examining the fit without data, against current data, or against additional data.
- **Model > Select** is the same as the **Select** button in the **Local Models** pane and opens the Model Selection window. Here you can examine the fit of the two-stage model against the local fit and the data.
- **View > View Model Definition** duplicates the toolbar button and opens the Model Viewer dialog.

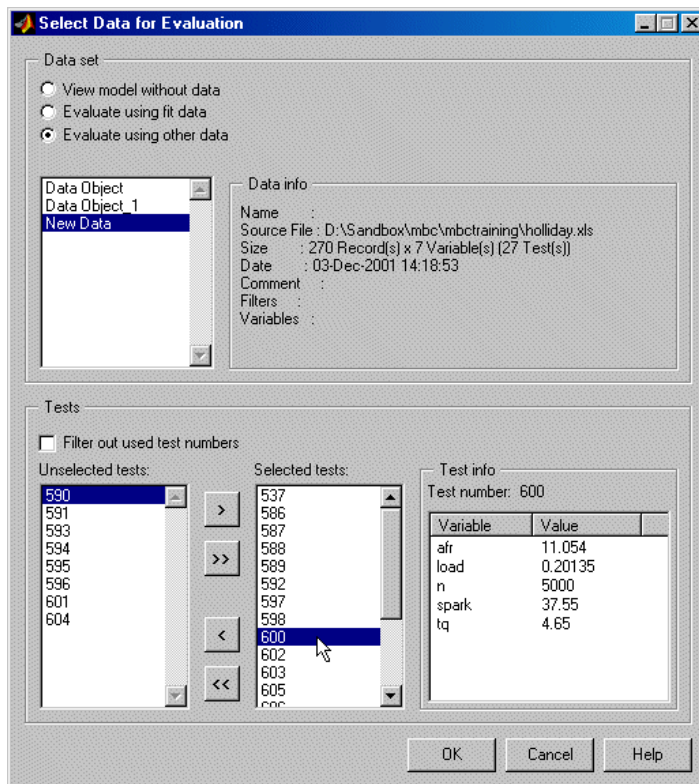
Model Evaluation Window

You can access the Model Evaluation window via the menu item **Model > Evaluate** (or the hot key **Ctrl+E**) from any of the modeling nodes: the one-stage or two-stage model node, local model node, response feature nodes, or any child model nodes.

Recall that the Model Selection window allows you to compare different models with each other and with the data used to create these models. The Model Evaluation window also allows you either to examine a model without data or to validate it against data other than that used in creating the model. For any model node, model evaluation is a quick way to examine the model in more ways than those available in the main Browser views. For example, local models with more than one input factor can only be viewed in the Predicted/Observed view in the main Browser, and the Model Selection window only shows you the two-stage model, so you go to Model Evaluation to view the local model itself in more detail. For other models such as childless response feature nodes, or their child nodes, **Select** is not available, so Model Evaluation is the only way to view these models in detail.

The Model Evaluation window comprises some of the same views you see in the Model Selection window. The views available depend on what kind of model you are evaluating (two-stage, local, global, or one-stage) and the evaluation mode you choose.

The following example illustrates the evaluation process using other data to validate a model. From a model node, you reach the Model Evaluation window via the **Model > Evaluate...** menu item. The Data Selection dialog appears.



There are three modes for evaluation, each represented by a radio button choice:

- **View model without data** — The evaluation window appears with only the model surface view and the cross-section view. These do not show the model fit against data. Here you can investigate the shape of the model surface.
- **Evaluate using fit data** — The evaluation window appears, and the data shown along with the model surface is the data that was used to create this model. The views available are residuals, response surface, and cross

section. If you are evaluating a two-stage model, you can also have the tests view, where the local model fit is shown for each test. If you are evaluating a one-stage model or a two-stage response feature model, the predicted/observed view is also available.

- **Evaluate using other data** — The data that is shown along with the model surface is chosen in this dialog. The evaluation window then appears with the residuals, response surface, and cross-section view. If you are evaluating a two-stage model, the tests view is also available, although the local fit is not shown, as the local model was fitted to different data.

In the tests view (two-stage models only) you can click and hold on a plot to see the values of the global variables for that test.

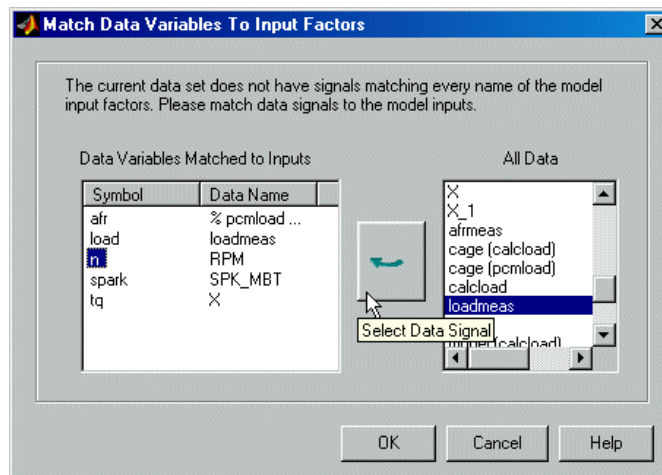
Model Evaluation Using Other Data

Prediction performance creates a resource dilemma for the experimenter. You could keep some data back to use later on to get a measure of the predictive capacity of the model, but why withhold data from the model building process if its addition is likely to improve the predictive capability of the model?

There is always a trade-off in statistics when you decide how much of your data to use for model fitting and whether to leave some data to test your model against.

The Model Evaluation window is intended to help you validate your model against other data, although you can also evaluate the fit against the original data or without any data. The rest of the Model Evaluation dialog is only enabled when you choose **Evaluate using other data**, the last of three choices in the Data Selection dialog.

When you select the **Evaluate using other data** choice, if the input factors required to evaluate the model do not appear in the selected data set, another dialog appears to match signal names in the selected data set to those required to evaluate the model.



Match signal names from the data set (on the right) to the symbols required to evaluate the model (on the left) and click **OK** to accept this assignment.

In the example shown in “Model Evaluation Window” on page 4-148, the Data Selection dialog shows three data sets that you can use to evaluate this model: Data Object, Data Object_1, and New Data. Traceability information about the data set currently selected in the list box is always displayed in the **Data info** pane on the right.

In the lower **Tests** frame you can choose the tests from this data set to use. Certain test numbers might have been used to create the model (as in this case) and you can eliminate these test numbers from the evaluation data set if you want, by selecting the **Filter out used log numbers** check box.

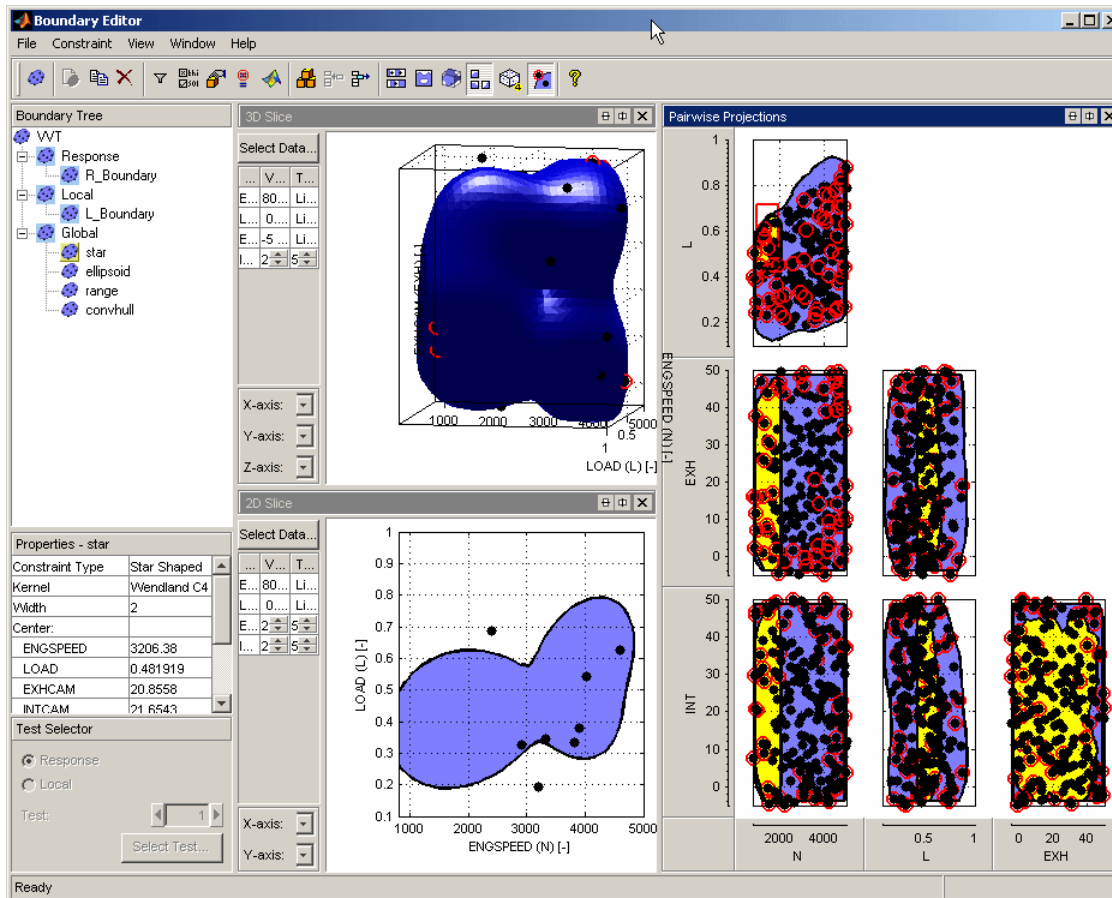
Of the remaining tests, you can select some and reject others using the **Selected Tests/Unselected Tests** list boxes. You move tests between them using the arrow buttons. For the currently selected test, the mean test values of all variables in this data set are displayed on the right.

Click **OK** to use the selected tests to evaluate this model. The Model Evaluation window appears.

For information about the available views, see “Selecting Models” on page 4-111.

Boundary Constraint Modeling

From the test plan level, if you select **TestPlan > Boundary Constraints** the Boundary Constraint Editor appears.



Here you can construct boundary constraint models from your data. You can construct nonparametric surfaces as a visual aid to understanding complex operating envelopes, and use these boundaries to guide modeling and constrain optimizations.

- “Creating A Boundary Constraint” on page 4-154
- “Boundary Constraint Toolbar” on page 4-157
- “Boundary Constraint Menus” on page 4-159
- “Complete Constraint Fit Dialog” on page 4-161

These boundary constraints are integrated with the rest of the toolbox functionality. You can view them in the Model Selection window (in the cross section view and response surface view) and in CAGE in optimizations, tradeoff, model views and the Surface Viewer . You can import boundary constraints into the Design Editor. They are useful for clipping models to view only the area of interest, constraining modeling and designs to realistic engine operating envelopes, and for designating the most valid areas for optimization, tradeoff and calibration.


- From the Model Browser you can export boundary constraints with your models, from the test plan level — use the **File** menu to export either to CAGE, to the Workspace or as an exported model file (.exm). See “Exporting Models” on page 4-195.

You can use the function `ceval` to evaluate a boundary constraint exported to the Workspace. For example, if your exported model is `M`, then `ceval(M, X)` evaluates the boundary constraint attached to `M` at the points given by the matrix `X` (values less than zero are inside the boundary). See “Evaluating Models in the Workspace” on page 4-200 for more details.


- From the Boundary Editor you can export boundary constraint models to Simulink.
- You can import boundary constraints into the Design Editor.

Creating A Boundary Constraint

To fit a boundary model,

- 1 Get started by clicking **Make Boundary Constraint**  in the toolbar, or select **File > Make Boundary**.
- 2 Select a radio button to specify whether you want to model the boundary of the response, local or global values and click **OK**.

3 For global and response boundary models:



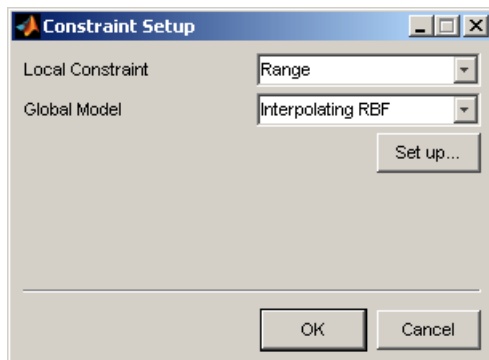
Select a type of constraint: Range, Star-shaped, Ellipsoid, or Convex Hull.

- The Range constraint finds the furthest extent of points for each variable and draws a hyper-rectangle to enclose all points.
- The Ellipsoid constraint forms an ellipse to enclose all points.
- The Convex Hull constraint forms the minimal convex set containing the data points.
- The Star Shaped constraint is a more complex constraint with various settings that determine how your boundary model is calculated. This happens in three stages: determining the center of the data; deciding which points are on the boundary, and interpolating between those points. The star-shaped constraint is the only constraint type that can fit non-convex regions.

You can select a set of input factors to model using the **Active factors** check boxes. The required and selected number of inputs is displayed on the right. It can be useful to build constraint models using two or three

input factors and you can combine them for the most accurate boundary; this can be more effective than including all factors.

For *local* boundary models you can set up a two-stage boundary:



You can select **Range** or **Ellipsoid** for the **Local Constraint** (to fit to the local factors). The global model must be an interpolating RBF that interpolates across the global factors between these local constraints. You can click **Set Up** to change the parameters for the interpolating RBF.

Click **OK**.

- 4** The options in the next dialog depends on which type of constraint you selected. **Range**, **Convex Hull** and **Two-Stage** constraints do not have any further settings you can alter. **Star-shaped** and **ellipsoid** constraints have a variety of parameters you can alter, see “**Complete Constraint Fit Dialog**” on page 4-161. Try the defaults before experimenting with these. Click **OK** and the constraint model is calculated.





You can reach the settings in Step 3 by selecting **Constraint > Set Up** or the equivalent toolbar button to set or edit the constraint type and active factors.

You can reach the settings in Step 4 by selecting **Constraint > Full Constraint Fit** or the equivalent toolbar button to alter the constraint settings where available.


You can view the type and details of the selected constraint in the Properties pane. This contains information about the constraint such as: number of data points; number of boundary, interior, and exterior points; the type of constraint

and summary of settings (e.g. centre point of star). For root & branch nodes you can see which constraint or combination of constraints is selected as best.


Boundary Constraint Toolbar

-  Make Boundary Constraint — click here to set up a new boundary constraint. See “Creating A Boundary Constraint” on page 4-154. Also in the **File** menu.
-  New Constraint— — Adds a new child node to the current node. As for the model tree in the Model Browser, these child nodes are different in different places.
 - a For one-stage test plans, new child nodes of the root (top) node are boundary constraints (leaf nodes).
 - b For two-stage test plans, new child nodes differ depending on the parent node. From the top or root node, a new child node is a response, local or global node (you can only have one of each). From local, global or response nodes, new child nodes are boundary constraints of the same type as their parent node - local, global or response. You can add as many boundary constraints of each type as you want.
-  Duplicate Boundary Constraint— Duplicates the current node. Only available for leaf nodes.
-  Delete Boundary Constraint— Deletes the current node. Not available for the root node.
- Constraint Fit Settings — the following buttons open dialogs where you can edit different settings of a constraint. Different settings are available for different constraints. Also in the **Constraint** menu.
 - Set Up — see step 3 of “Creating A Boundary Constraint” on page 4-154.
 - Full Constraint Fit — see “Complete Constraint Fit Dialog” on page 4-161
 - Find Special Points — see “Star Shaped Constraint Settings” on page 4-162


- Find Boundary Points — see “Complete Constraint Fit Dialog” on page 4-161
- Fit Constraint — see “Star Shaped Constraint Settings” on page 4-162

-  1-D Slice — shows a 1-D slice through your model and data. Also in the **View** menu. You can select values for variables in the edit boxes. These values determine the point at which the slice through the constraint is plotted. You can change the tolerance values in the **Tolerance** edit boxes next to each variable. This determines the ‘width’ of the slice. Data points within the tolerances will be displayed with the slice. This is very similar to the display in the Cross Section model selection view, where the Tolerance either side of the displayed model slice determines how near data points must be to the model slice to be displayed. See “Cross Section View” on page 4-136 for comparison.

You can click and hold on data points to view the values of the factors at that point, and the distance from the boundary. Double-click a data point to move the slice view to that point, or click **Select Data Point** to choose a particular point.




-  2-D View — shows a 2-D slice through your model and data. Also in the **View** menu. You can choose which variables to plot on the X and Y drop-down menus. Similarly to the 1-D Slice view, you can change the values and tolerances of the other variables in the edit boxes to determine where the constraint slice is plotted and how much data is also displayed.

You can click and hold on data points to view the values of the factors at that point, and the distance from the boundary. Double-click a data point to move the slice view to that point, or click **Select Data Point** to choose a particular point.

-  3-D Slice — shows a 3-D slice of your model and data. Also in the **View** menu. You can choose which variables to use for the three axes using the drop-down menus, and set the resolution of the grid (number of points) to display in the value edit boxes for each factor. You can set the value of other variables in the edit boxes as for the 2-D view.

You can click and hold on data points to view the values of the factors at that point, and the distance from the boundary. Double-click a data point to

move the slice view to that point, or click **Select Data Point** to choose a particular point.

-  Pairwise View — shows a pairwise projection of your constraint and data. Also in the **View** menu. There is a plot of the entire range for each pair of variables in turn. You can click and drag to select a region. Do not click on points, click in the blue or white regions and then drag. The same region is then highlighted in yellow in each projection, so you can see how your data is distributed in each dimension.
-  4-D Data Projection — a 3-D plot of the data points colored by a 4th factor. You can use the drop-down menus to select plot and color bar factors. Also in the **View** menu.
-  Highlight Boundary Points — this highlights in red all data points that are on the boundary surface. This applies to all views. The points highlighted red (on the boundary) should have a **Distance** value of zero (distance to the boundary) when you click and hold on these points. Due to rounding errors the value may not be exactly zero, just very small. Highlight Boundary Points is also in the **View** menu and the right-click context menu on views.

Boundary Constraint Menus

- In the **File** menu:
 - **Make Boundary** — select this to set up a new boundary constraint. See “Creating A Boundary Constraint” on page 4-154.
 - **New, Duplicate, and Delete Constraint** — add, copy or remove a node. Use **Duplicate** to make a selection of leaf nodes to make changes and compare many constraint models. You cannot duplicate or delete the root node, and you cannot duplicate the Local or Global nodes. See the toolbar.
 - **Export to Simulink** — the currently selected node is exported as a block in a Simulink model. If the currently selected node is a leaf node, that is a single boundary model. If it is a parent node (root, local, global or response) it may be one boundary or a combination of several depending on which nodes you have assigned best and added to best.

- **Close** — this closes the Boundary Constraint Editor. Your constraint models are saved with the test plan.

- In the **Constraint** menu:

The following menu choices open dialogs where you can edit different settings of a constraint. Different settings are available for different constraints. These functions are also in the toolbar.

- **Constraint Type** — see step 3 of “Creating A Boundary Constraint” on page 4-154.
- **Full Constraint Fit** — see “Complete Constraint Fit Dialog” on page 4-161
- The following three choices give you access to specific parts of the Complete Constraint Fit dialog:
 - a** Find Special Points
 - b** Find Boundary Points
 - c** Fit Constraint
- **Assign Best** — you can choose a best constraint. In the same way as in the model tree in the Model Browser, this best constraint is copied up the tree, and the icons of constraints selected as best become outlined in blue. For example, if you have a selection of constraints under the root node for a one-stage model and select one as best, the icon becomes blue and this constraint is copied to the root node above. Similarly, for two-stage models, if you assign a best constraint from a variety of leaf nodes under the Global node, that best model is copied up to the Global node. You can see which model is currently best from the parent node (e.g.l Global) by looking under **Best Constraint** in the bottom left **Properties** pane.

Note If you have built multiple constraints you *must* select one (or a combination) as best or you will not be able to export constraints with your models.

- **Add to Best** — you can also combine constraints. For example, for a one-stage model, if you pick two leaf nodes and select **Add to Best** for each, the root node will show a combination of the two constraints. For

example you may want your final constraint to combine the boundaries of a star shaped and range constraint. You can view the results: the combined model is clipped to fit within the ranges defined by both constraints. You can combine as many of the same type of leaf nodes as you like.


Similarly you can combine best models for both local and global leaf nodes. Note you can only combine local leaf nodes with other local leaf nodes, and global leaf nodes with other global leaf nodes, because you can only have one best model at the Local node, and one best model at the Global node. However you can choose to combine the Local and Global constraints for the root node. You can always see which constraints have been combined at the currently selected node in the **Properties** pane.

- **Remove From Best** — use this to reverse your selections of best and combined best models.
- In the **View** menu:
 - Under **Current View** and **Split View** there are submenus where you can select any of the views available in the toolbar — 1-D, 2-D, 3-D slices, the pairwise view, and the 4D data view.
 - You can select graph size.
 - You can toggle boundary point highlighting on and off.
 - You can split the currently selected view horizontally and vertically, and close the current view, as in the Design Editor and Data Editor.

All these options are available in the right-click context menu on the views.

Complete Constraint Fit Dialog

You can reach these settings by either

- Selecting **Constraint > Full Constraint Fit** or the equivalent toolbar button, or
- Selecting  in the toolbar, or **File > Make Boundary**. At step 4 you reach the Complete Constraint Fit dialog, see “Creating A Boundary Constraint” on page 4-154.

Range, Convex Hull and Two-Stage constraints do not have any further settings you can alter in this dialog. Star-shaped and ellipsoid constraints have a variety of parameters you can alter, detailed below.

Note that the following three toolbar and **Constraint** menu choices give you access to specific parts of the Complete Constraint Fit dialog:

- 1 Find Special Points
- 2 Find Boundary Points
- 3 Fit Constraint

Ellipsoid Constraint Settings

The ellipsoid constraint has a variety of optimization settings you can alter if you are having problems getting a good fit. Change the display setting to `iter` or `final` to see output messages at the command line during the fit, and try different tolerances or numbers of iterations. For details on these settings see `optimset` in the MATLAB Reference.

Star Shaped Constraint Settings

The Star Shaped constraint is a more complex constraint with various settings that determine how your boundary model is calculated. This is done in three stages: determining the center of the data; deciding which points are on the boundary, and interpolating between those points. These settings are detailed below.

- **Special Points > Center** — This is not the same as the Center Selection settings for RBFs, this is the method for determining the center of the boundary model sphere. The boundary model can be thought of as a deformed sphere. You can choose Mean, Median, Mid Range, MinEllipse or User Defined. If you select User Defined you can enter a value for each input.

Note You can also reach the **Special Points** settings by selecting **Constraint > Find Special Points** or the equivalent toolbar button.


- **Boundary Points** These settings determine how to decide which points are on the boundary.

Note You can also reach the **Boundary Points** settings by selecting **Constraint > Find Boundary Points** or the equivalent toolbar button.

Interior — choose this if you don't want all your points to be on the boundary.

Boundary Only — places all points on the boundary. If this is suitable for your data then this saves time on this stage.

- **Dilation Radius** — If you choose **Interior** points, then **Dilation Radius** is used to determine which points are on the boundary. Each point is expanded to a sphere until a boundary shape is formed by the intersection of all those spheres. See the illustration below.

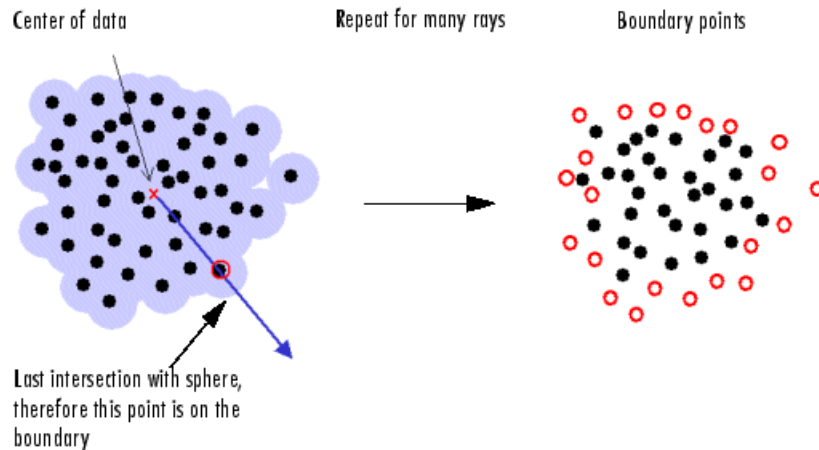


Dilation Radius settings:

- **Auto** — This setting selects the dilation radius (how much to expand each point) by checking all the minimum distances between points, then choosing the largest of those.
- **Manual** — You can manually set the dilation radius in the edit box. The default is 1. This may seem large as model range is between -1 and 1, but all points are expanded equally so you will still detect the points on the edge. However, very large spheres will intersect and obscure points that should be detected as boundary points.

- **Ray Casting** — Rays are drawn from the centre of the boundary model to determine which points are on the edge. The last point intersected on each ray is a boundary point (remember that the ray actually intersects a sphere, given that each point has been expanded by the dilation radius). See the following illustration.

RayCasting



Ray casting settings:


- **From Data** — This uses the same number of rays as there are data points and sends one ray in the direction of each point. If you have very dense data or a very large number of points it might be better to use the **Manual** setting to choose a smaller number of rays.
- **Manual** — You can set a value in the **Number of Rays** edit box. This number of rays will then be used in random directions. A good guide is about twice the number of data points, though if you have a large number (many hundreds) it will become very slow and you may run out of memory. In most situations more than 1000 is too many.
- **Constraint Fit Options.**

Note You can also reach the **Constraint Fit Options** settings by selecting **Constraint > Fit Constraint** or the equivalent toolbar button.

- **Transform** — None, Log, or McCallum. The default is None. Depending on the shape of your boundary, you might need to use a transform to prevent self intersections near the center of the model.
- **Radial Basis Function (RBF)** settings — you can choose RBF kernels, width and continuity as when setting up models. See “Global Model Class: Radial Basis Function” on page 4-14 for more information. Once the boundary points have been determined, each of those points is used as an RBF center and the boundary surface is obtained by interpolating radial basis functions between all those centers. The width and continuity settings depend on which kernel you choose.
- **RBF Algorithm**

These options control the interpolating RBF model settings. You can leave the defaults unless you have a very large data set (several thousand points). With very large data sets it can help speed and robustness of fitting to try a different Algorithm setting here (e.g. try GMRES first) and vary the tolerance and number of iterations.

Stepwise

You can open the stepwise regression window through the  toolbar icon when you are in the global level view (that is, with a response feature selected in the “Model Tree” on page 1-19). The Stepwise tool provides a number of methods of selecting the model terms that should be included. You can use the **Stepwise** menu in the Model Setup dialogs to run stepwise automatically when building linear models.

Minimizing Predicted Error Sum of Squares (PRESS) is a good method for working toward a regression model that provides good predictive capability over the experimental factor space. See “PRESS statistic” on page 4-174.


The use of PRESS is a key indicator of the predictive quality of a model. The predicted error uses predictions calculated without using the observed value for that observation. PRESS is known as Delete-1 statistics in the Statistics Toolbox. See also “Two-Stage Models for Engines” on page 4-181.

The next section describes the displays and controls in the Stepwise window.

Using the Stepwise Regression Window

- 1 Use the stepwise command buttons at the bottom of the following figure (and duplicated in the **Regression** menu) as follows:
 - Click **Min. PRESS** to automatically include or remove terms to minimize PRESS. This procedure provides a model with improved predictive capability.
 - **Include All** terms in the model (except the terms flagged with **Status** as Never). This option is useful in conjunction with **Min. PRESS** and backward selection. For example, first click **Include All**, then **Min. PRESS**. Then you can click **Include All** again, then **Backwards**, to compare which gives the best result.
 - **Remove All** terms in the model (except the terms flagged with **Status** as Always). This option is useful in conjunction with forward selection (click **Remove All**, then **Forwards**).

- **Forwards** selection adds all terms to the model that would result in statistically significant terms at the $\alpha\%$ level. The addition of terms is repeated until all the terms in the model are statistically significant.



- **Backwards** selection removes all terms from the model that are not statistically significant at the $\alpha\%$ level. The removal of terms is repeated until all the terms in the model are statistically significant.

- 2 Terms can be also be manually included or removed from the model by clicking on the Term, Next PRESS, or coefficient error bar line.

The confidence intervals for all the coefficients are shown to the right of the table. Note that the interval for the constant term is not displayed, as the value of this coefficient is often significantly larger than other coefficients.

Terms that are currently not included in the model are displayed in red.

- 3 A history of the PRESS and summary statistics is shown on the right of the stepwise figure. You can return to a previous model by selecting an item in the list box or a point on the stepwise plot.

The ANOVA Table on page 4-98 for the current model is shown on the right side of the stepwise figure.

- 4 The critical values for testing whether a coefficient is statistically different from zero at the $\alpha\%$ level are displayed at the bottom right side of the stepwise figure. The value of α can be entered in the edit box to the left of the critical values. The default is 5%.

Any changes made in the stepwise figure automatically update the diagnostic plots in the Model Browser.

You can revert to the starting model when closing the Stepwise window. When you exit the Stepwise window, the Confirm Stepwise Exit dialog asks Do you want to update regression results? You can click **Yes** (the default), **No** (to revert to the starting model), or **Cancel** (to return to the Stepwise window).

Stepwise Table

| Term | Label for Coefficient |
|------------|---|
| Status | Always. Stepwise does not remove this term.
Never. Stepwise does not add this term.
Step. Stepwise considers this term for addition or removal. |
| B | Value of coefficient. When the term is not in the model the value of the coefficient if it is added to the model is displayed in red. |
| stdB | Standard error of coefficient. |
| t | <i>t</i> value to test whether the coefficient is statistically different from zero. The <i>t</i> value is highlighted in blue if it is less than the critical value specified in the $\alpha\%$ edit box (at bottom right). |
| Next PRESS | The value of PRESS if the status of this term is changed at the next iteration. The smallest Next PRESS is highlighted with a yellow background, The column header is also yellow if there is a model change that results in a smaller PRESS value. |

The preceding table describes the meanings of the column headings in the Stepwise Regression window.

Automatic Stepwise

You can set the Minimize PRESS, Forward, and Backward selection routines to run automatically without the need to enter the stepwise figure.

You can set these options in the Model Setup dialog, when you initially set up your test plan, or from the global level as follows:

- 1 Select **Model > Set Up**.

- 2** The Global Model Setup dialog has a drop-down menu **Stepwise**, with the options None, Minimize PRESS, Forward selection, and Backward selection.

Stepwise in the Model Building Process

Once you have set up a model, you should create several alternative models, use the Stepwise functions and examine the diagnostic statistics to search for a good model fit. For each response feature,

- 1** Begin by conducting a stepwise search.

You can do this automatically or by using the Stepwise window.

The goal of the stepwise search is to minimize PRESS. Usually not one but several candidate models per response features arise, each with a very similar PRESS R^2 . The predictive capability of a model with a PRESS R^2 of 0.91 cannot be assumed superior in any meaningful engineering sense to a model with a PRESS R^2 of 0.909. Further, the nature of the model building process is that the “improvement” in PRESS R^2 offered by the last few terms is often very small. Consequently, several candidate models can arise. You can store each of the candidate models and associated diagnostic information separately for subsequent review. Do this by making a selection of child nodes for the response feature.

However, experience has shown that a model with a PRESS R^2 of less than 0.8, say, is of little use as a predictive tool for engine mapping purposes. This criterion must be viewed with caution. Low PRESS R^2 values can result from a poor choice of the original factors but also from the presence of outlying or influential points in the data set. Rather than relying on PRESS R^2 alone, a safer strategy is to study the model diagnostic information to discern the nature of any fundamental issues and then take appropriate corrective action.

- 2** Once the stepwise process is complete, the diagnostic data should be reviewed for each candidate model.

It might be that these data alone are sufficient to provide a means of selecting a single model. This would be the case if one model clearly exhibited more ideal behavior than the others. Remember that the interpretation of diagnostic plots is subjective.

- 3** You should also remove outlying data at this stage. You can set criteria for detecting outlying data. The default criterion is any case where the absolute value of the external studentized residual is greater than 3.
- 4** After removing outlying data, continue the model building process in an attempt to remove further terms.

High-order terms might have been retained in the model in an attempt to follow the outlying data. Even after removing outlying data, there is no guarantee that the diagnostic data will suggest that a suitable candidate model has been found. Under these circumstances,

- 5** A transform of the response feature might prove beneficial.

A useful set of transformations is provided by the Box and Cox family, see “Box-Cox Transformation” on page 4-176. Note that the Box-Cox algorithm is model dependent and as such is always carried out using the $(N \times q)$ regression matrix X .

- 6** After you select a transform, you should repeat the stepwise PRESS search and select a suitable subset of candidate models.
- 7** After this you should analyze the respective diagnostic data for each model.

It might not be apparent why the original stepwise search was carried out in the natural metric. Why not proceed directly to taking a transformation? This seems sensible when it is appreciated that the Box-Cox algorithm often, but not always, suggests that a contractive transform such as the square root or log be applied. There are two main reasons for this:

- The primary reason for selecting response features is that they possess a natural engineering interpretation. It is unlikely that the behavior of a transformed version of a response feature is as intuitively easy to understand.
- Outlying data can strongly influence the type of transformation selected. Applying a transformation to allow the model to fit bad data well does not seem like a prudent strategy. By “bad” data it is assumed that the data is truly abnormal and a reason has been discovered as to why the data is outlying; for example, “The emission analyzer was purging while the results were taken.”

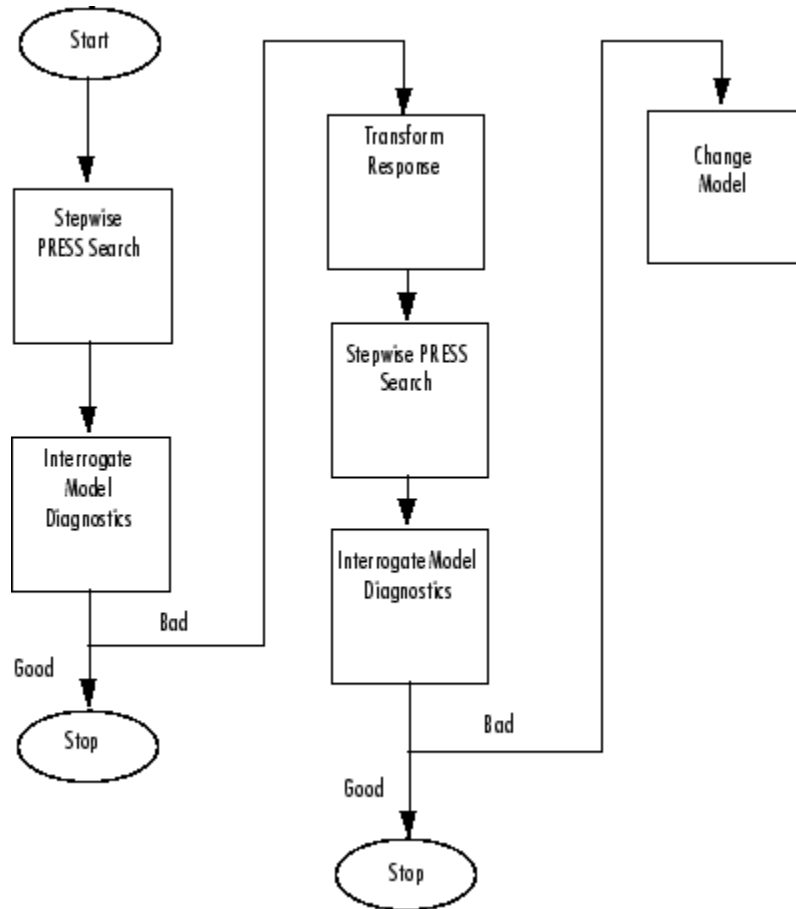
Finally, if you cannot find a suitable candidate model on completion of the stepwise search with the transformed metric, then a serious problem exists either with the data or with the current level of engineering knowledge of the system. Model augmentation or an alternative experimental or modeling strategy should be applied in these circumstances.

After these steps it is most useful to validate your model against other data (if any is available). See “Model Evaluation Window” on page 4-148.

See also these guideline pages with links to information about each of the steps involved in creating one and two-stage models and then searching for the best fit:

- “How To Set Up a One-Stage Model” on page 1-6
- “How To Set Up a Two-Stage Model” on page 1-6
- “Creating Multiple Models To Compare” on page 1-8

The recommended overall stepwise process is best viewed graphically, as shown in the following flow chart.



Note that the process depicted in the preceding diagram should be carried out for each member of the set of response features associated with a given response and then repeated for the remaining responses.

PRESS statistic

With n runs in the data set, the model equation is fitted to $n-1$ runs and a prediction taken from this model for the remaining one. The difference between the recorded data value and the value given by the model (at the value of the omitted run) is called a prediction residual. PRESS is the sum of squares of the prediction residuals. The square root of PRESS/n is PRESS RMSE (root mean square prediction error).

Note that the prediction residual is different from the ordinary residual, which is the difference between the recorded value and the value of the model when fitted to the whole data set.

The PRESS statistic gives a good indication of the predictive power of your model, which is why minimizing PRESS is desirable. It is useful to compare PRESS RMSE with RMSE as this may indicate problems with overfitting. RMSE is minimized when the model gets very 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.

For more information see "Stepwise" on page 4-166 and "Toolbox Terms and Statistics Definitions" on page 4-192.

Note that calculating PRESS for the two-stage model applies the same principle (fitting the model to $n-1$ runs and taking a prediction from this model for the remaining one) but in this case the predicted values are first found for response features instead of data points. The predicted value, omitting each test in turn, for each response feature is estimated. The predicted response features are then used to reconstruct the local curve for the test and this curve is used to obtain the two-stage predictions. This is applied as follows:

To calculate two stage PRESS:

- 1 For each test, S , do the following steps:

- For each of the response features, calculate what the response feature predictions would be for S (with the response features for S removed from the calculation).
- This gives a local prediction curve C based on all tests except S.
- For each data point in the test, calculate the difference between the observed value and the value predicted by C.

2 Repeat for all tests.

Sum the square of all of the differences found and divide by the total number of data points.

Box-Cox Transformation

You might want to transform a response feature either to correct for nonnormality and/or a heteroscedastic variance structure. A useful class of transformations for this purpose is the power transform y^λ , where λ is a parameter to be determined. Box and Cox (1964) showed how λ and the regression coefficients themselves could be estimated simultaneously using the method of maximum likelihood. The procedure consists of conducting a standard least squares fit using

$$y^{(\lambda)} = \frac{y^\lambda - 1}{\lambda y^{\lambda-1}} \text{ for } \lambda \neq 0$$

$$y^{(\lambda)} = y \ln(y) \text{ for } \lambda = 0$$

where the so called geometric mean of the observations is given by

$$y = \exp \left[\frac{\sum_{i=1}^N \ln(y_i)}{N} \right]$$

The maximum likelihood estimate of λ corresponds to the value for which the $SSE(\lambda)$ from the fitted model is a minimum. This value of λ is determined by fitting a model (assumed throughout to be defined by the regression matrix for the full model - X) for various levels of λ and choosing the value corresponding to the minimum $SSE(\lambda)$. A plot of $SSE(\lambda)$ versus λ is often used to facilitate this choice.

The parameter λ is swept between the range of -3 to 3 in increments of 0.5.

You can enter a value for lambda in the edit box that approaches the point on the plot with the smallest SSE.

Although $SSE(\lambda)$ is a continuous function of λ , simple choices for λ are recommended. This is because the practical difference between 0.5 and 0.593,

say, is likely to be very small but a simple transform like 0.5 is much easier to interpret.

You can also find an approximate $100(1-\alpha)$ confidence interval on l by computing

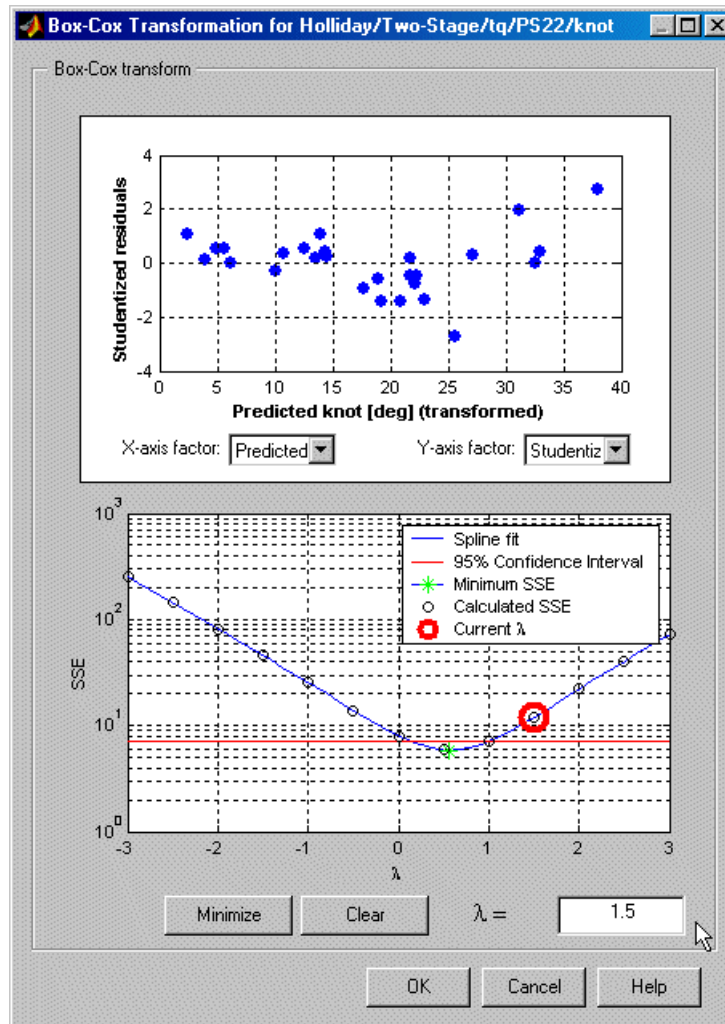
$$SS^* = SSE(\lambda) \left[1 + \frac{t_{\alpha/2, \nu}^2}{\nu} \right]$$

where ν is the number of residual degrees of freedom equal to $(N-q)$.

In this formula λ is understood to be the value that minimizes $SSE(\lambda)$. Note that this confidence interval might encompass more than one incremental value for λ . In this case, any of these values is as valid as any other and you can select any of these transformations from which to develop trial models.

- You should always look at the residuals plots at the top to see the effect of different transforms.
- You can create several child nodes of a single model and choose different transforms for each to compare them using the rest of the Model Browser tools.

For the sake of clarity, consider the following example, which illustrates the results of applying the Box-Cox algorithm to a polyspline torque model.



In this example the minimum value of $SSE(\lambda)$ occurs near to $\lambda=0$. The minimum is marked in green. The 95% confidence limit has been calculated and drawn on the figure as a red solid line. It is apparent in this example that, after rounding to the nearest incremental value contained within the

confidence interval, any λ in the range $0 \leq \lambda \leq 1$ is appropriate. Of the three possible increments, 0, 0.5, and 1, $\lambda = 0.5$ is the closest to the minimum SSE.

You can select any point on the plot by clicking. The chosen point (current lambda) is then outlined in red. You can also enter values of lambda directly in the edit box and press **Enter**.

Maximum Likelihood Estimation for Nonlinear Repeated Measures

This section contains an overview of the mathematics of two-stage models. A comprehensive reference for two-stage modeling is Davidian and Giltinan [3]. The information is divided into the following sections:

- “Two-Stage Models for Engines” on page 4-181
- “Local Models” on page 4-182
 - “Local Covariance Modeling” on page 4-183
 - “Response Features” on page 4-184
- “Global Models” on page 4-185
- “Two-Stage Models” on page 4-186
 - “Initial Values for Covariances” on page 4-188
 - “Quasi-Newton Algorithm” on page 4-188
 - “Expectation Maximization Algorithm” on page 4-189
 - “Global Model Selection” on page 4-188
- “References” on page 4-189

Two-Stage Models for Engines

Lindstrom and Bates [6] define repeated measurements as data generated by observing a number of individuals repeatedly under various experimental conditions, where the individuals are assumed to constitute a random sample from a population of interest. An important class of repeated measurements is longitudinal data where the observations are ordered by time or position in space. More generally, longitudinal data is defined as repeated measurements where the observations on a single individual are not, or cannot be, randomly assigned to the levels of a treatment of interest.

Modeling data of this kind usually involves the characterization of the relationship between the measured response, y , and the repeated measurement factor, or covariate x . Frequently, the underlying systematic relationship between y and x is nonlinear. In some cases the relevant nonlinear model can be derived on physical or mechanistic grounds. However, in other contexts a nonlinear relationship might be imposed simply to provide a convenient empirical description for the data. The presence of repeated observations on an individual requires particular care in characterizing the variation in the experimental data. In particular, it is important to represent two sources of variation explicitly: random variation among measurements within a given individual (*intraindividual*) and random variation among individuals (*interindividual*). Inferential procedures accommodate these different variance components within the framework of an appropriate hierarchical statistical model. This is the fundamental idea behind the analysis of repeated measurement data.

Holliday [1,2] was perhaps the first to apply nonlinear repeated measurements analysis procedures to spark injection engine data. The focus of Holliday's work was the modeling of data taken from engine mapping experiments. In these experiments, engine speed, load, and air/fuel ratio were held constant while spark was varied. Various engine response characteristics, for example, torque or emission quantities, were measured at each spark setting. Holliday modeled the response characteristics for each sweep as a function of spark advance. Variations in the individual sweep parameters were then modeled as a function of the global engine operating variables speed, load, and air/fuel ratio. Conceptually, variations in the measurements taken within a sweep represent the intraindividual component of variance. Similarly, variation in the sweep-specific parameters between sweeps represents the interindividual component of variance. You can generalize these principles

to other steady-state engine modeling exercises where the nature of data collection usually involves sweeping a single engine control variable while the remainder are held at fixed values. These points suggest that nonlinear repeated measurements analysis represents a general approach to the parameterization of mean value engines models for controls-oriented development.

Another application for models of this form is the flow equations for a throttle body. Assuming the flow equations are based upon the usual one-dimensional isentropic flow principle, then they must be modified by an effective area term, A_e , which accounts for the fact that the true flow is multidimensional and irreversible. You can map the throttle flow characteristics by sweeping the throttle position at fixed engine speed. This data collection methodology naturally imposes a hierarchy the analysis of which is consistent with the application of nonlinear repeated measures. Experience in modeling effective area suggests that free knot spline or biological growth models provide good local predictions. The global phase of the modeling procedure is concerned with predicting the systematic variation in the response features across engine speed. A free knot spline model has proven useful for this purpose.

Local Models

Modeling responses locally within a sweep as a function of the independent variable only. That is,

(4-1)

$$y_i^j = f_i(s_i^j, \theta_i) + \varepsilon_i^j \quad \text{for } j = 1, 2, \dots, m_i \quad (4-1)$$

where the subscript i refers to individual tests and j to data within a test, s_i^j is the j^{th} independent value, θ_i is a (rx1) parameter vector, y_i^j is the j^{th} response, and ε_i^j is a normally distributed random variable with zero mean and variance σ^2 . Note that equation (4-1) can be either a linear or a nonlinear function of the curve fit parameters. The assumption of independently normally distributed errors implies that the least squares estimates of θ are also maximum likelihood parameters.

Local Covariance Modeling

The local model describes both the systematic and random variation associated with measurements taken during the i^{th} test. Systematic variation is characterized through the function f while variation is characterized via the distributional assumptions made on the vector of random errors e_i . Hence, specification of a model for the distribution of e_i completes the description of the intratest model. The Model-Based Calibration Toolbox allows a very general specification of the local covariance:

(4-2)

$$e_i \sim N(0, \sigma^2 C_i(\beta_i, \xi_i)) \quad (4-2)$$

where C_i is an $(n_i \times n_i)$ covariance matrix, σ^2 is the coefficient of variation, and ξ_i is a $(q\text{-by-}1)$ vector of dispersion parameters that account for heterogeneity of variance and the possibility of serially correlated data. The specification is very general and affords considerable flexibility in terms of specifying a covariance model to adequately describe the random component of the intratest variation.

The Model-Based Calibration Toolbox supports the following covariance models:

- Power Variance Model:

(4-3)

$$C_i = \text{diag}\left\{f(x_i, \beta_i)^{\xi_{i1}}\right\} \quad (4-3)$$

- Exponential Variance Model:

(4-4)

$$C_i = \text{diag}\{\exp(f(x_i, \beta_i)\xi_{i1})\} \quad (4-4)$$

- Mixed Variance Model:

(4-5)

$$C_i = \text{diag}\left\{\xi_1 + f(x_i; \theta_i) \xi_2\right\} \quad (4-5)$$

where $\text{diag}\{x\}$ is a diagonal matrix.

Correlation models are only available for equispaced data in the Model-Based Calibration Toolbox. It is possible to combine correlation models with models with the variance models such as power.

One of the simplest structures that can be used to account for serially correlated errors is the AR(m) model (autoregressive model with lag m). The general form of the AR(m) model is

(4-6)

$$e_j = \phi_1 e_{j-1} + \phi_2 e_{j-2} + \dots + \phi_m e_{j-m} + v_j \quad (4-6)$$

where ϕ_k is the k^{th} lag coefficient and v_j is an exogenous stochastic input identically and independently distributed as $N(0, \sigma_v^2)$. First- and second-order autoregressive models are implemented in the Model-Based Calibration Toolbox.

Another possibility is a moving average model (MA). The general structure is

(4-7)

$$e_j = \phi_1 v_{j-1} + \phi_2 v_{j-2} + \dots + \phi_m v_{j-m} + v_j \quad (4-7)$$

where ϕ_k is the k^{th} lag coefficient and v_j is an exogenous stochastic input identically and independently distributed as $N(0, \sigma_v^2)$. Only a first-order moving average model is implemented in the Model-Based Calibration Toolbox.

Response Features

From an engineering perspective, the curve fit parameters do not usually have any intuitive interpretation. Rather characteristic geometric features of the curve are of interest. The terminology “response features” of Crowder and

Hand [7] is used to describe these geometric features of interest. In general, the response feature vector p_i for the i^{th} sweep is a nonlinear function (g) of the corresponding curve fit parameter vector θ_i , such that

(4-8)

$$P_i = g(\theta_i) \quad (4-8)$$

Global Models

Modeling the variation in the response features as a function of the global variables. The response features are carried through to the second stage of the modeling procedure rather than the curve fit parameters because they have an engineering interpretation. This ensures that the second stage of the modeling process remains relatively intuitive. It is much more likely that an engineer will have better knowledge of how a response feature such as MBT behaves throughout the engine operating range (at least on a main effects basis) as opposed to an esoteric curve fit parameter estimate.

The global relationship is represented by one of the global models available in the Model-Based Calibration Toolbox. In this section we only consider linear models that can be represented as

(4-9)

$$P_i = X_i\beta + \gamma_i \text{ for } i=1,2,\dots,r \quad (4-9)$$

where the X_i contains the information about the engine operating conditions at the i^{th} spark sweep, β is the vector of global parameter estimates that must be estimated by the fitting procedure, and γ_i is a vector of normally distributed random errors. It is necessary to make some assumption about the error distribution for γ , and this is typically a normal distribution with

(4-10)

$$\gamma_i \sim N_r(0,D) \quad (4-10)$$

where r is the number of response features. The dimensions of D are $(r \times r)$ and, being a variance-covariance matrix, D is both symmetric and positive

definite. Terms on the leading diagonal of D represent the test-to-test variance associated with the estimate of the individual response features. Off-diagonal terms represent the covariance between pairs of response features. The estimation of these additional covariance terms in a multivariate analysis improves the precision of the parameter estimates.

Two-Stage Models

To unite the two models, it is first necessary to review the distributional assumptions pertaining to the response feature vector \mathbf{p}_i . The variance of \mathbf{p}_i ($\text{Var}(\mathbf{p}_i)$) is given by

$$(4-11) \quad \text{Var}(\mathbf{p}_i) = \left[\frac{\partial \mathbf{g}(\theta_i)}{\partial \theta} \right] \sigma^2 C_i \left[\frac{\partial \mathbf{g}(\theta_i)}{\partial \theta} \right]^T \quad (4-11)$$

For the sake of simplicity, the notation $\sigma^2 C_i$ is to denote $\text{Var}(\mathbf{p}_i)$. Thus, \mathbf{p}_i is distributed as

$$(4-12) \quad \mathbf{p}_i \sim N_r(\mathbf{p}_i, \sigma^2 C_i) \quad (4-12)$$

where C_i depends on f_i through the variance of θ_i and also on g_i through the conversion of θ_i to the response features \mathbf{p}_i . Two standard assumptions are used in determining C_i : the asymptotic approximation for the variance of maximum likelihood estimates and the approximation for the variance of functions of maximum likelihood estimates, which is based on a Taylor series expansion of g_i . In addition, for nonlinear f_i or g_i , C_i depends on the unknown θ_i ; therefore, we will use the estimate $\hat{\theta}_i$ in its place. These approximations are likely to be good in the case where σ^2 is small or the number of points per sweep (m_i) is large. In either case we assume that these approximations are valid throughout.

We now return to the issue of parameter estimation. Assume that the γ_i are independent of the ϵ_i^j . Then, allowing for the additive replication error in response features, the response features are distributed as

(4-13)

$$p_i \sim N(X_i\beta, \sigma^2 C_i + D) \quad (4-13)$$

When all the tests are considered simultaneously, equation (6-13) can be written in the compact form

(4-14)

$$P \sim N(Z\beta, W(\omega)) \quad (4-14)$$

where P is the vector formed by stacking the n vectors p_i on top of each other, Z is the matrix formed by stacking the n X_i matrices, W is the block diagonal weighting matrix with the matrices on the diagonal being $\sigma^2 C_i + D$, and ω is a vector of dispersion parameters. For the multivariate normal distribution (6-14) the negative log likelihood function can be written:

(4-15)

$$\log L(\beta, \omega) = \log |W| + (P - Z\beta)' W^{-1} (P - Z\beta) \quad (4-15)$$

Thus, the maximum likelihood estimates are the vectors β_{ML} and ω_{ML} that minimize $\log L(\beta, \omega)$. Usually there are many more fit parameters than dispersion parameters; that is, the dimension of β is much larger than ω . As such, it is advantageous to reduce the number of parameters involved in the minimization of $\log L(\beta, \omega)$. The key is to realize that equation (6-15) is conditionally linear with respect to β . Hence, given estimates of ω , equation (6-15) can be differentiated directly with respect to β and the resulting expression set to zero. This equation can be solved directly for β as follows:

(4-16)

$$\beta = (Z'W^{-1}Z)^{-1} (Z'W^{-1}P) \quad (4-16)$$

The key point is that now the likelihood depends only upon the dispersion parameter vector ω , which as already discussed has only modest dimensions. Once the likelihood is minimized to yield ω_{ML} , then, since $W(\omega_{ML})$ is then known, equation (6-16) can subsequently be used to determine β_{ML} .

Global Model Selection

Before undertaking the minimization of Equation 4-15 (see “Two-Stage Models” on page 4-186) it is first necessary to establish the form of the X_i matrix. This is equivalent to establishing a global expression for each of the response features a priori. Univariate stepwise regression is used to select the form of the global model for each response feature. Minimization of the appropriate PRESS statistic is used as a model building principle, as specified in “Stepwise in the Model Building Process” on page 4-170. The underlying principle is that having used univariate methods to establish possible models, maximum likelihood methods are subsequently used to estimate their parameters.

Initial Values for Covariances

An initial estimate of the global covariance is obtained using the standard two-stage estimate of Steimer *et al.* [10],

(4-17)

$$D_{\text{STS}} = \frac{1}{r-1} \sum_{i=1}^r (p_i - X_i\beta)(p_i - X_i\beta)^T \quad (4-17)$$

where β are the estimates from all the univariate global models. This estimate is biased.

Quasi-Newton Algorithm

Implicit to the minimization of equation (6-17) is that D is positive definite. It is a simple matter to ensure this by noting that D is positive definite if and only if there is an upper triangular matrix, G , say, such that

(4-18)

$$D = G'G \quad (4-18)$$

This factorization is used in the Quasi-Newton algorithm. Primarily, the advantage of this approach is that the resulting search in G , as opposed to D , is unconstrained.

Expectation Maximization Algorithm

The expectation maximization algorithm is an iterative method that converges toward the maximal solution of the likelihood function. Each iteration has two steps:

- 1 Expectation Step — Produce refined estimates of the response features given the current parameter estimates.
- 2 Maximization Step — Obtain new estimates of the parameters (global model parameters and covariance matrix) for the new response features.

These steps are repeated until the improvement in value of the log likelihood function is less than the tolerance. Details of the algorithm can be found in [3, Chapter 5].

References

- 1 Holliday, T., *The Design and Analysis of Engine Mapping Experiments: A Two-Stage Approach*, Ph.D. thesis, University of Birmingham, 1995.
- 2 Holliday, T., Lawrance, A. J., Davis, T. P., Engine-Mapping Experiments: A Two-Stage Regression Approach, *Technometrics*, 1998, Vol. 40, pp 120-126.
- 3 Davidian, M., Giltinan, D. M., *Nonlinear Models for Repeated Measurement Data*, Chapman & Hall, First Edition, 1995.
- 4 Davidian, M., Giltinan, D. M., Analysis of repeated measurement data using the nonlinear mixed effects model, *Chemometrics and Intelligent Laboratory Systems*, 1993, Vol. 20, pp 1-24.
- 5 Davidian, M., Giltinan, D. M., Analysis of repeated measurement data using the nonlinear mixed effects model, *Journal of Biopharmaceutical Statistics*, 1993, Vol. 3, part 1, pp 23-55.
- 6 Lindstrom, M. J., Bates, D. M., Nonlinear Mixed Effects Models for Repeated Measures Data, *Biometrics*, 1990, Vol. 46, pp 673-687.
- 7 Davidian, M., Giltinan, D. M., Some Simple Methods for Estimating Intraindividual Variability in Nonlinear Mixed Effects Models, *Biometrics*, 1993, Vol. 49, pp 59-73.

- 8** Hand, D. J., Crowder, M. J., *Practical Longitudinal Data Analysis*, Chapman and Hall, First Edition, 1996.
- 9** Franklin, G.F., Powell, J.D., Workman, M.L., *Digital Control of Dynamic Systems*, Addison-Wesley, Second Edition, 1990.
- 10** Steimer, J.-L., Mallet, A., Golmard, J.L., and Boisvieux, J.F., Alternative approaches to estimation of population pharmacokinetic parameters: Comparison with the nonlinear mixed effect model. *Drug Metabolism Reviews*, 1984, 15, 265-292.

Linear Regression

Building a regression model that includes only a subset of the total number of available terms involves a trade-off between two conflicting objectives:

- Increasing the number of model terms always reduces the Sum Squared Error.
- However, you do not want so many model terms that you overfit by chasing points and trying to fit the model to signal noise. This reduces the predictive value of your model.

The best regression equation is the one that provides a satisfactory trade-off between these conflicting goals, at least in the mind of the analyst. It is well known that there is no unique definition of *best*. Different model building criteria (for example, forward selection, backward selection, PRESS search, stepwise search, Mallows C_p Statistic...) yield different models. In addition, even if the optimal value of the model building statistic is found, there is no guarantee that the resulting model will be optimal in any other of the accepted senses.

Principally the purpose of building the regression model for calibration is for predicting future observations of the mean value of the response feature. Therefore the aim is to select the subset of regression terms such that PRESS is minimized. Minimizing PRESS is consistent with the goal of obtaining a regression model that provides good predictive capability over the experimental factor space. This approach can be applied to both polynomial and spline models. In either case the model building process is identical.

- 1 The regression matrix can be viewed in the Design Evaluation Tool. Terms in this matrix define the *full model*. In general, the stepwise model is a subset of this full term set.
- 2 All regressions are carried out with the factors represented on their coded scales (-1,1).

Toolbox Terms and Statistics Definitions

Definitions

| Symbol | Definition |
|-----------|--|
| N | Number of data points |
| p | Number of terms currently included in the model |
| q | Total number of possible model parameters ($q=p+r$) |
| r | Number of terms not currently included from the model |
| y | (N \times 1) response vector |
| X | Regression matrix. X has dimensions (N \times q) |
| X_p | (N \times p) model matrix corresponding to terms currently included in the model |
| X_r | (N \times r) matrix corresponding to terms currently excluded from the model |
| β_p | (p \times 1) vector of model coefficients
$\beta_p = \{\beta_1, \beta_2, \dots, \beta_p\}$ $\hat{\beta} = (X^T X)^{-1} X^T y$ $\text{var}\hat{\beta} = (X^T X)^{-1} \text{MSE}$ |
| PEV | Prediction Error Variance
$\text{PEV}(x) = \text{var}y(\hat{y}) = x(X^T X)^{-1} x^T \text{MSE}$ |
| α | User-defined threshold criteria for automatically rejecting terms |
| \hat{y} | (N \times 1) vector of predicted responses.
$\hat{y} = X_p \beta_p$ |

Definitions (Continued)

| Symbol | Definition |
|--------------------|---|
| e | (N×1) residual vector.
$e = (y - \hat{y})$ |
| e _(i) | (N×1) vector of PRESS residuals.
$e_{(i)} = e_i / (1 - H_{ii})$ |
| H | Hat matrix.
$X'(X'X)^{-1}X$ |
| L | (N×1) vector of leverage values.
$L = \{l_1, l_2, \dots, l_N\}' = \{H_{11}, H_{22}, \dots, H_{NN}\}$ |
| VIF | Variance Inflation Factors |
| SSE | Error Sum of Squares. SSE = e'e |
| SSR | Regression Sum of Squares. SSE =
$\sum e_i^2$ |
| SST | Total Sum of Squares. SST = y'y - N \bar{y}^2 |
| MSE | Mean Square Error. MSE = SSE/(N-p) |
| MSR | Mean Square of Regression. MSR = SSR/P |
| F | F-statistic. F = MSR/MSE |
| MSE _(i) | MSE calculated with i th point removed from the data set.
$MSE_{(i)} = \frac{(N-p)MSE - e_i^2 / (1 - H_{ii})}{N-p-1}$ |
| RMSE | Root Mean Squared Error: the standard deviation of regression. $RMSE = \sqrt{MSE}$ |

Definitions (Continued)

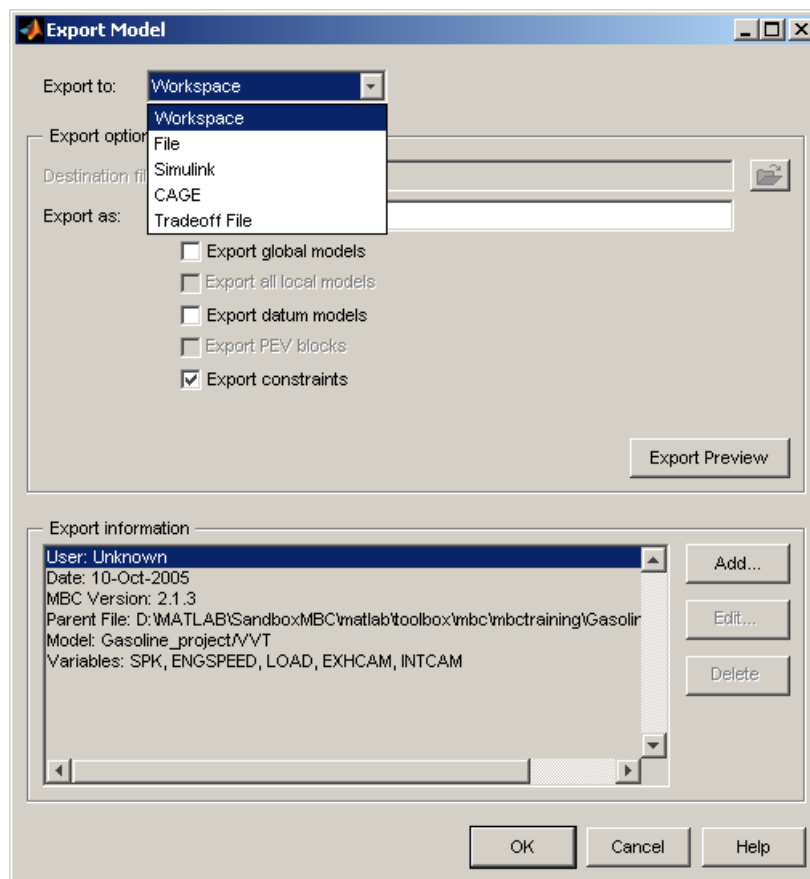
| Symbol | Definition |
|---------------|--|
| s_i | i^{th} R-Student or Externally Scaled Studentized Residual.
$s_i = \frac{e_i}{\sqrt{MSE_{(i)}(1 - H_{ii})}}$ |
| r_i | i^{th} Standardized or Internally Scaled Studentized Residual.
$r_i = \frac{e_i}{\sqrt{MSE(1 - H_{ii})}}$ |
| D | Cook's D Influence Diagnostic.
$D_i = \frac{r_i^2 H_{ii}}{p(1 - H_{ii})}$ |
| SEBETA | (px1) vector of model coefficient standard errors.
$SEBETA = MSE [\sqrt{c_{11}}, \sqrt{c_{22}}, \dots, \sqrt{c_{pp}}]'$ where
$c = (X^T X)^{-1}$ |
| PRESS | Predicted Error Sum of Squares. $PRESS = e'_{(i)} e_{(i)}$ |

For more on PRESS and other displayed statistics, see "PRESS statistic" on page 4-174, Summary Table on page 4-97, and "Pooled Statistics" on page 4-79.

Exporting Models

You can export all models created in the Model Browser using the menu item **File > Export Models**.

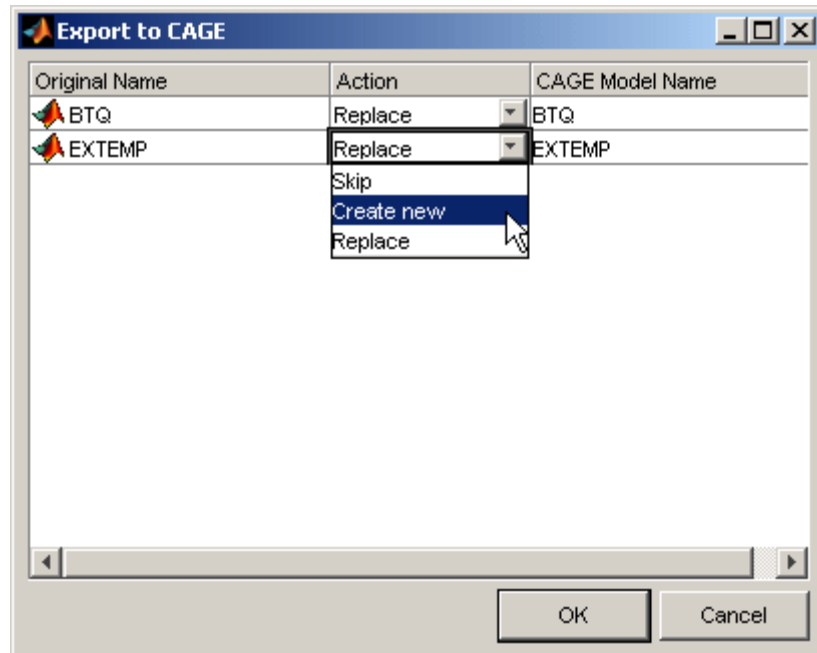
The Export Model dialog appears.



Choose the export format from the **Export to** drop-down menu:

- **Workspace** — to the MATLAB workspace. See “Evaluating Models in the Workspace” on page 4-200.

- File — to an EXM file for CAGE. The EXM file format is specifically designed for loading into the CAGE part of MBC Toolbox, for example, to create calibrations.
- Simulink — to a Simulink model (an MDL file)
- CAGE — export directly to CAGE (only if CAGE is open). If you export to CAGE you will see the Export to CAGE dialog after you click **OK** to export.



In the dialog you can see a list of the models to export. For each model you can edit the name, replace a suitable existing CAGE model or skip the selected model.

- Tradeoff — You can export to Tradeoff File only when you are exporting from the test plan node. This is only possible when the current test plan has at least one two-stage model available for export, and when this model has exactly two global input factors. Exporting to Tradeoff File exports all local models in a .mat file that can then be loaded into the Multimodel Tradeoff tool within CAGE.

If a file format is chosen (export to file or to Simulink), the **Destination file** controls are enabled, and the browse button allows you to locate a destination file using a file browser.

Export As — Defines the name that the model has when loaded, or the name of the workspace variable.

Export global models — When a two-stage model is being exported (from the response node) the constituent response feature models can also be exported. Multiple models are exported to the workspace as a cell array.

Export all local models — When exporting at the local node, the single local model for the current test is exported. Selecting this control exports the local models for all tests (to the workspace as a cell array).

Export datum models— When exporting a two-stage model that has a datum defined, this control allows you to export the datum global model (without exporting all other response feature models).

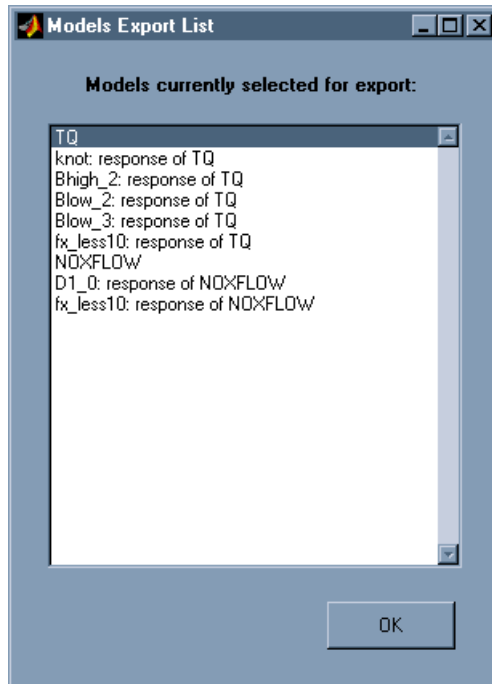
Export PEV blocks — When exporting to Simulink, you can create a PEV block as part of the Simulink diagram so that the prediction error variance can be evaluated along with the model. This is not available for models where PEV cannot be calculated.

Export constraints — Where design constraints and boundary constraints exist, you can choose to export them.

- To File, Workspace or CAGE: exports the boundary constraints with the models. You can also export boundary models to Simulink directly from the Boundary Editor.
- To Simulink: exports the design constraints with the models. You can also export design constraints with designs to .mvd file from the Design Editor.

Export Preview — Displays the models that are exported with the current choice of options.

For example:



See “What Is Exported?” on page 4-198 for details on what to expect here.

Export information — Traceability information is exported with the models. You can add/edit/delete further comments using the buttons on the right.

Click **OK** to export the current selection of models and close the dialog.

What Is Exported?

In the Export Model dialog you can use the **Export Preview** button to check the models that have been selected for export. This displays the Models Export List dialog.

At the test plan node:

- You cannot export from the test plan node unless all response models within that test plan have a best model chosen (that is, you used the

Select process at response level to assign the best model and calculate the two-stage model). All models within the test plan are exported.

- If the datum model check box is selected, the datum model is exported.
- If the checkbox for constraints is selected, constraints are also exported (boundary constraints to CAGE, File and workspace; design constraints to Simulink). See **Export constraints** under “Exporting Models” on page 4-195
- If the check box for global models is selected, all the response features are also exported.

At the response node:

- The response model is exported. You cannot export from a response node until it contains a best model (it is empty before that).
- If the check box for global models is selected, all the response features are also exported. Note that the datum model is not necessarily a response feature.
- There is also a datum model check box. As at the test plan node, this exports the datum model along with the two-stage model (without exporting all other response feature models).

At the local node:

- The local model *for the current test only* is exported.
- If the check box for all local models is selected, all the local models are exported for all the tests.
- If the node is purely a local model (with a house icon, that is; no two-stage model has yet been calculated), the model is exported under its own name (for example, PS2,3); if a two-stage model has been calculated (that is, the local node has a two-stage house-and-globe icon), the local model is exported under the name of the response node.

At response feature nodes and all child nodes from here:

- The current response feature (or other selected subnode) only is exported.

Note When you are exporting to Simulink, linear models support PEV, so the **Export PEV blocks** check box is active. This is only the case when you are exporting from response features that are linear functions of the local parameters. See Export PEV blocks under “Exporting Models” on page 4-195.

Evaluating Models in the Workspace

If a model is exported to the workspace as MyModel and has four input factors, it can be evaluated at a point as in this example:

```
Y = MyModel([3.7, 89.55, -0.005, 1]);
```

If column vectors p1,p2,p3,p4 (of equal length) are created for each input factor, the model can be evaluated to give a column vector output

```
Y = MyModel([p1,p2,p3,p4]);
```

Left to right argument order corresponds exactly to top-down input order in the Test Plan view in the Model Browser.

If you export a group of models, the group is exported as a cell array of models. The argument order in the curly braces {1 to n} cell array reference corresponds exactly to top-down model order in the model tree in the Model Browser.

The inputs and outputs for MATLAB model evaluation are in natural engineering units, not coded units.

It is also possible to evaluate the PEV (prediction error variance) for the model using the command

```
[pev, y] = pev(MyModel, [x1 x2 x3])
```

You can use one or two arguments, as follows:

[p]=pev(x) gives pev at x

[p,y]=pev(x) gives pev at x and model eval at x

Evaluating Boundary Models

You can use the function `ceval` to evaluate a boundary constraint exported to the Workspace. For example, if your exported model is `M`, then `ceval(M, X)` evaluates the boundary constraint attached to `M` at the points given by the matrix `X` (values less than zero are inside the boundary). See “Boundary Constraint Modeling” on page 4-153.

For example, if you have exported multiple responses from a test plan as a cell array named `modeltutorial`, entering the following at the command line evaluates the boundary model for the first response {1} at the point where all 4 inputs are zero.

```
ceval(modeltutorial{1}, [0,0,0,0])
```

Response models are in top-down order in the model tree, e.g. in this example {1} is the top model in the tree under the test plan node. [0,0,0,0] is the matrix of input values, where left to right argument order corresponds exactly to top-down input order in the Boundary Editor or the Test Plan view in the Model Browser, e.g. `spk`, `load`, `rpm`, and `afr`.

You can quickly check the number of model inputs as follows:

```
nfactors( modeltutorial{1} )
```

You can click a point in the boundary editor (in the 1-D, 2-D and 3-D views) to check the input names and get example input values to evaluate in the workspace, e.g.

```
ceval (modeltutorial{1}, [25, 0.64, 5000, 14.43] )
```

```
ans = 3.0284e-004
```

Boundary constraint distance of zero means the point is on the boundary, negative values are inside the constraint, and positive values are outside. The range is typically [-1,1] but not always, and roughly linear. Rather like information criteria it is only a comparison that is meaningful (point `x` has a greater distance than point `y`) rather than the absolute value.

Confidence Intervals

The confidence interval is given by:

$$\text{upperbound} = y + t * \text{sqrt}(\text{pev})$$

$$\text{lowerbound} = y - t * \text{sqrt}(\text{pev})$$

Where y is the model prediction, and t is the appropriate percentile of the t-statistic, with $df = n\text{Obs}-1$ degrees of freedom. This is calculated using the function `tinv` from the Statistics Toolbox as follows:

$$t = \text{tinv}(p,v)$$

p =confidence level e.g. 95%

v = degrees of freedom ($n-1$)

$$t = \text{tinv}(1-\alpha/2, df)$$

where $\alpha = 0.05$ for 95% confidence intervals.

Radial Basis Functions

This section discusses the following topics:

| | |
|---|--|
| Guide to Radial Basis Functions for Model Building (p. 5-3) | An mathematical introduction to Radial Basis Functions (RBFs). |
| Types of Radial Basis Functions (p. 5-5) | An illustrated guide to the types of Radial Basis Function available in the toolbox. |
| Fitting Routines (p. 5-11) | An introduction to the characteristics of RBFs that must be decided: weights, centers, width and lambda. |
| Center Selection Algorithms (p. 5-12) | Description and parameters for the Rols, RedErr, WiggleCenters and CenterExchange algorithms. |
| Lambda Selection Algorithms (p. 5-15) | Description and parameters for the IterateRidge, IterateRols and StepItRols algorithms. |
| Width Selection Algorithms (p. 5-19) | Description and parameters for the TrialWidths, WidPerDim and Tree Regression algorithms. |
| Prune Functionality (p. 5-25) | How to use the Prune function to reduce the number of centers in a radial basis function network. This helps you decide how many centers are needed. |
| Statistics (p. 5-28) | Mathematical description and references. |

Hybrid Radial Basis Functions
(p. 5-32)

Parameters specific to hybrid RBFs.

Tips for Modeling with Radial Basis
Functions (p. 5-34)

A set of guidelines to help you choose
and improve RBF models.

Guide to Radial Basis Functions for Model Building

A radial basis function has the form

$$z(x) = \phi(\|x - \mu\|)$$

where x is a n -dimensional vector, μ is an n -dimensional vector called the center of the radial basis function, $\| \cdot \|$ denotes Euclidean distance, and ϕ is a univariate function, defined for positive input values, that we shall refer to as the profile function.

The model is built up as a linear combination of N radial basis functions with N distinct centers. Given an input vector x , the output of the RBF network is the activity vector \hat{y} given by

$$\hat{y}(x) = z \sum_{j=1}^N \beta_j z_j(x)$$

where β_j is the weight associated with the j th radial basis function, centered at μ_j , and $z_j = \Phi(\|x - \mu_j\|)$. The output \hat{y} approximates a target set of values denoted by y .

A variety of radial basis functions are available in MBC, each characterized by the form of Φ . All of the radial basis functions also have an associated width parameter σ , which is related to the spread of the function around its center. Selecting the box in the model setup provides a default setting for the width. The default width is the average over the centers of the distance of each center to its nearest neighbor. This is a heuristic given in Hassoun (see “References” on page 5-31) for Gaussians, but it is only a rough guide that provides a starting point for the width selection algorithm.

Another parameter associated with the radial basis functions is the regularization parameter λ . This (usually small) positive parameter is used in most of the fitting algorithms. The parameter λ penalizes large weights, which tends to produce smoother approximations of y and to reduce the tendency of the network to overfit (that is, to fit the target values y well, but to have poor predictive capability).

The following sections explain the different parameters for the radial basis functions available in the Model-Based Calibration Toolbox, and how to use them for modeling.

Types of Radial Basis Functions

Within the model setup, you can choose which RBF kernel to use. Kernels are the types of RBF (multiquadric, gaussian, thinplate, and so on).

Gaussian


This is the radial basis function most commonly used in the neural network community. Its profile function is

$$\Phi(r) = e^{(-r^2/\sigma^2)}$$

This leads to the radial basis function


$$z(x) = \exp\left(-\frac{\|x - \mu\|^2}{\sigma^2}\right)$$

In this case, the width parameter is the same as the standard deviation of the gaussian function.



Thin-Plate Spline


This radial basis function is an example of a smoothing spline, as popularized by Grace Wahba (<http://www.stat.wisc.edu/~wahba/>). They are usually supplemented by low-order polynomial terms. Its profile function is



Logistic Basis Function

These radial basis functions are mentioned in Hassoun (see “References” on page 5-31). They have the profile function

$$\Phi(r) = \frac{1}{1 + \exp\left(\frac{r}{\sigma}\right)}$$



Wendland's Compactly Supported Function

These form a family of radial basis functions that have a piecewise polynomial profile function and compact support [Wendland, see “References” on page 5-31]. The member of the family to choose depends on the dimension of the space (n) from which the data is drawn and the desired amount of continuity of the polynomials.

| Dimension | Continuity | Profile |
|-----------|------------|------------------------------------|
| n=1 | 0 | $\Phi(r) = (1-r)_+$ |
| | 2 | $\Phi(r) = (1-r)_+^3(3r+1)$ |
| | 4 | $\Phi(r) = (1-r)_+^5(8r^2+5r+1)$ |
| n=3 | 0 | $\Phi(r) = (1-r)_+^2$ |
| | 2 | $\Phi(r) = (1-r)_+^4 + (4r+1)$ |
| | 4 | $\Phi(r) = (1-r)_+^6(35r^2+18r+3)$ |
| n=5 | 0 | $\Phi(r) = (1-r)_+^3$ |
| | 2 | $\Phi(r) = (1-r)_+^5(5r+1)$ |
| | 4 | $\Phi(r) = (1-r)_+^7(16r^2+7r+1)$ |

We have used the notation $a_+ := \begin{cases} a, & a > 0 \\ 0, & a \leq 0 \end{cases}$ for the positive part of a.


When n is even, the radial basis function corresponding to dimension n+1 is used.

Note that each of the radial basis functions is nonzero when r is in [0,1]. It is possible to change the support to be [0,σ] by replacing r by r/σ in the preceding formula. The parameter σ is still referred to as the width of the radial basis function.

Similar formulas for the profile functions exist for n>5, and for even continuity > 4. Wendland's functions are available up to an even continuity of 6, and in any space dimension n.

Notes on Use

- Better approximation properties are usually associated with higher continuity.
- For a given data set the width parameter for Wendland's functions should be larger than the width chosen for the Gaussian.



Multiquadrics

These are a popular tool for scattered data fitting. They have the profile function


$$\Phi(r) = \sqrt{r^2 + \sigma^2} \lambda$$

Reciprocal Multiquadrics

These have the profile function


$$\Phi(r) = 1/\sqrt{r^2 + \sigma^2}$$

Note that a width σ of zero is invalid.



Linear

These have the profile function $\phi = -r$.



Fitting Routines

There are four characteristics of the RBF that need to be decided: weights, centers, width, and λ . Each of these can have significant impact on the quality of the resulting fit, and good values for each of them need to be determined. The weights are always determined by specifying the centers, width, and λ , and then solving an appropriate linear system of equations. However, the problem of determining good centers, width, and λ in the first place is far from simple, and is complicated by the strong dependencies among the parameters. For example, the optimal λ varies considerably as the width parameter changes. A global search over all possible center locations, width, and λ is computationally prohibitive in all but the simplest of situations.

To try to combat this problem, the fitting routines come in three different levels.

At the lowest level are the algorithms that choose appropriate centers for given values of width and λ . The centers are chosen one at a time from a candidate set (usually the set of data points or a subset of them). The resulting centers are therefore ranked in a rough order of importance.

At the middle level are the algorithms that choose appropriate values for λ and the centers, given a specified width.

At the top level are the algorithms that aim to find good values for each of the centers, width, and λ . These top-level algorithms test different width values. For each value of width, one of the middle-level algorithms is called that determines good centers and values for λ .

These algorithms and their fit parameters are described in the following sections:

- 1 “Center Selection Algorithms” on page 5-12
- 2 “Lambda Selection Algorithms” on page 5-15
- 3 “Width Selection Algorithms” on page 5-19

Center Selection Algorithms

Rols

This is the basic algorithm as described in Chen, Chng, and Alkadhimi [See “References” on page 5-31]. In Rols (Regularized Orthogonal Least Squares) the centers are chosen one at a time from a candidate set consisting of all the data points or a subset thereof. It picks new centers in a forward selection procedure. Starting from zero centers, at each step the center that reduces the regularized error the most is selected. At each step the regression matrix X is decomposed using the Gram-Schmidt algorithm into a product $X = WB$ where W has orthogonal columns and B is upper triangular with ones on the diagonal. This is similar in nature to a QR decomposition. Regularized error is given by $e'e + \lambda g'g$ where $g = Bw$ and e is the residual, given by $e = y - \hat{y}$. Minimizing regularized error makes the sum square error $e'e$ small, while at the same time not letting $g'g$ get too large. As g is related to the weights by $g = Bw$, this has the effect of keeping the weights under control and reducing overfit. The term $g'g$ rather than the sum of the squares of the weights $w'w$ is used to improve efficiency.

The algorithm terminates either when the maximum number of centers is reached, or adding new centers does not decrease the regularized error ratio significantly (controlled by a user-defined tolerance).

Fit Parameters

Maximum number of centers — The maximum number of centers that the algorithm can select. The default is the smaller of 25 centers or π of the number of data points. The format is $\min(n_{\text{obs}}/4, 25)$. You can enter a value (for example, entering 10 produces ten centers) or edit the existing formula (for example, $(n_{\text{obs}}/2, 25)$ produces half the number of data points or 25, whichever is smaller).

Percentage of data to be candidate centers — The percentage of the data points that should be used as candidate centers. This determines the subset of the data points that form the pool to select the centers from. The default is 100%, that is, to consider all the data points as possible new centers. This can be reduced to speed up the execution time.

Regularized error tolerance — Controls how many centers are selected before the algorithm stops. See Chen, Chng, and Alkadhimi [“References” on page 5-31] for details. This parameter should be a positive number between 0 and 1. Larger tolerances mean that fewer centers are selected. The default is 0.0001. If less than the maximum number of centers is being chosen, and you want to force the selection of the maximum number, then reduce the tolerance to epsilon (eps).

RedErr

RedErr stands for Reduced Error. This algorithm also starts from zero centers, and selects centers in a forward selection procedure. The algorithm finds (among the data points not yet selected) the data point with the largest residual, and chooses that data point as the next center. This process is repeated until the maximum number of centers is reached.

Fit Parameters

Only has **Number of centers**.

WiggleCenters

This algorithm is based on a heuristic that you should put more centers in a region where there is more variation in the residual. For each data point, a set of neighbors is identified as the data points within a distance of \sqrt{nf} divided by the maximum number of centers, where nf is the number of factors. The average residuals within the set of neighbors is computed, then the amount of wiggle of the residual in the region of that data point is defined to be the sum of the squares of the differences between the residual at each neighbor and the average residuals of the neighbors. The data point with the most wiggle is selected to be the next center.

Fit Parameters

Almost as in the Rols algorithm, except no **Regularized error**.

CenterExchange

This algorithm takes a concept from optimal Design of Experiments and applies it to the center selection problem in radial basis functions. A candidate

set of centers is generated by a Latin hypercube, a method that provides a quasi-uniform distribution of points. From this candidate set, n centers are chosen at random. This set is augmented by p new centers, then this set of $n+p$ centers is reduced to n by iteratively removing the center that yields the best PRESS statistic (as in stepwise). This process is repeated the number of times specified in **Number of augment/reduce cycles**.

CentreExchange and Tree Regression (see “Tree Regression” on page 5-21) are the only algorithms that permit centers that are not located at the data points. This means that you do not see centers on model plots. The CentreExchange algorithm has the potential to be more flexible than the other center selection algorithms that choose the centers to be a subset of the data points; however, it is significantly more time consuming and not recommended on larger problems.

Fit Parameters

Number of centers — The number of centers that will be chosen

Number of augment/reduce cycles — The number of times that the center set is augmented, then reduced

Number of centers to augment by — How many centers to augment by

Lambda Selection Algorithms

Lambda is the regularization parameter.

IterateRidge

For a specified width, this algorithm optimizes the regularization parameter with respect to the GCV criterion (generalized cross-validation; see the discussion under GCV criterion).

The initial centers either are selected by one of the low-level center selection algorithms or the previous choice of centers is used (see discussion under the parameter Do not reselect centers). You can select an initial start value for λ by testing an initial number of values for lambda (set by the user) that are equally spaced on a logarithmic scale between 10^{-10} and 10 and choosing the one with the best GCV score. This helps avoid falling into local minima on the GCV - λ curve. The parameter λ is then iterated to try to minimize GCV using the formulas given in the GCV criterion section. The iteration stops when either the maximum number of updates is reached or the $\log_{10}(\text{GCV})$ value changes by less than the tolerance.

Fit Parameters

Center selection algorithm — The center selection algorithm to use.

Maximum number of updates — Maximum number of times that the update of λ is made. The default is 10.

Minimum change in $\log_{10}(\text{GCV})$ — Tolerance. This defines the stopping criterion for iterating λ ; the update stops when the difference in the $\log_{10}(\text{GCV})$ value is less than the tolerance. The default is 0.005.

Number of initial test values for lambda — Number of test values of λ to determine a starting value for λ . Setting this parameter to 0 means that the best λ so far is used.

Do not reselect centers for new width — This check box determines whether the centers are reselected for the new width value, and after each lambda update, or if the best centers to date are to be used. It is cheaper to

keep the best centers found so far, and often this is sufficient, but it can cause premature convergence to a particular set of centers.

Display — When you select this check box, this algorithm plots the results of the algorithm. The starting point for λ is marked with a black circle. As λ is updated, the new values are plotted as red crosses connected with red lines. The best λ found is marked with a green asterisk.

If too many graphs are likely to be produced, because of the **Display** check box being activated here, a warning is generated, and you have the option to stop execution.

A lower bound of 10^{-12} is placed on λ , and an upper bound of 10.

IterateRols

For a specified width, this algorithm optimizes the regularization parameter in the Rols algorithm with respect to the GCV criterion. An initial fit and the centers are selected by Rols using the user-supplied λ . As in IterateRidge, you select an initial start value for λ by testing an initial number of start values for lambda that are equally spaced on a logarithmic scale between 10^{-10} and 10, and choosing the one with the best GCV score.

λ is then iterated to improve GCV. Each time that λ is updated, the center selection process is repeated. This means that IterateRols is much more computationally expensive than IterateRidge.

A lower bound of 10^{-12} is placed on λ , and an upper bound of 10.

Fit Parameters

Center selection algorithm — The center selection algorithm to use. For IterateRols the only center selection algorithm available is Rols.

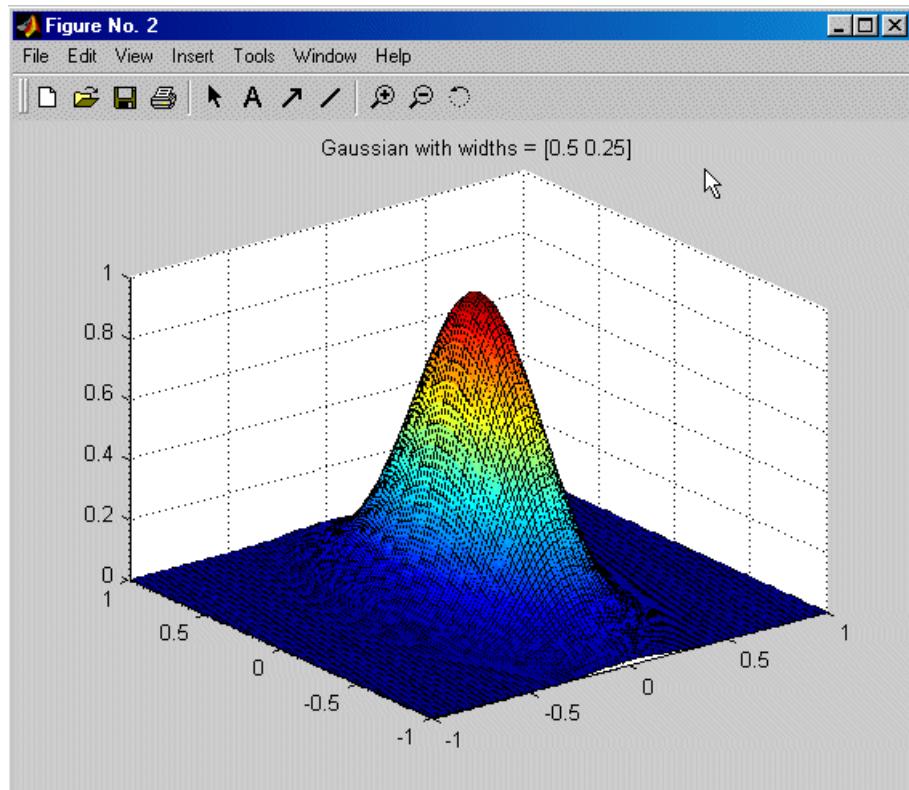
Maximum number of updates — The same as for IterateRidge.


Minimum change in log10(GCV) — The same as for IterateRidge.

Number of initial test values for lambda — The same as for IterateRidge.

Do not reselect centers for new width — This check box determines whether the centers are reselected for the new width value or if the best centers to date are to be used.

Display — When you select this check box, this algorithm plots the results of the algorithm. The starting point for λ is marked with a black circle.




As the above figure is updated, the new values are plotted as red crosses connected with red lines. The best  found is marked with a green asterisk.

If too many graphs are likely to be produced, because of the **Display** check box being activated here, a warning is generated, and you have the option to stop execution.

StepItRols

This algorithm combines the center-selection and lambda-selection processes.

Rather than waiting until all centers are selected before  is updated (as with the other lambda-selection algorithms), this algorithm offers the ability to update λ after each center is selected. It is a forward selection algorithm that, like Rols, selects centers on the basis of regularized error reduction. The stopping criterion for StepItRols is on the basis of $\log_{10}(\text{GCV})$ changing by less than the tolerance more than a specified number of times in a row (given in the parameter **Maximum number of times $\log_{10}(\text{GCV})$ change is minimal**). Once the addition of centers has stopped, the intermediate fit with the smallest $\log_{10}(\text{GCV})$ is selected. This can involve removing some of the centers that entered late in the algorithm.

Fit Parameters

Maximum number of centers — As in the Rols algorithm.

Percentage of data to candidate centers — As in the Rols algorithm.

Number of centers to add before updating — How many centers are selected before iterating λ begins.

Minimum change in $\log_{10}(\text{GCV})$ — Tolerance. It should be a positive number between 0 and 1. The default is 0.005.

Maximum number of times $\log_{10}(\text{GCV})$ change is minimal — Controls how many centers are selected before the algorithm stops. The default is 5. Left at the default, the center selection stops when the $\log_{10}(\text{GCV})$ values change by less than the tolerance five times in a row.

Width Selection Algorithms

TrialWidths

This routine tests several width values by trying different widths. A set of trial widths equally spaced between specified initial upper and lower bounds is selected. The width with the lowest value of $\log_{10}(\text{GCV})$ is selected. The area around the best width is then tested in more detail — this is referred to as a zoom. Specifically, the new range of trial widths is centered on the best width found at the previous range, and the length of the interval from which the widths are selected is reduced to $2/5$ of the length of the interval at the previous zoom. Before the new set of trial widths is tested, the center selection is updated to reflect the best width and λ found so far. This can mean that the location of the optimum width changes between zooms because of the new center locations.

Fit Parameters

Lambda selection algorithm — Midlevel fit algorithm that you test with the various trial values of λ . The default is `IterateRidge`.

Number of trial widths in each zoom — Number of trials made at each zoom. The widths tested are equally spaced between the initial upper and lower bounds. Default is 10.

Number of zooms — Number of times you zoom in. Default is 5.

Initial lower bound on width — Lower bound on the width for the first zoom. Default is 0.01.

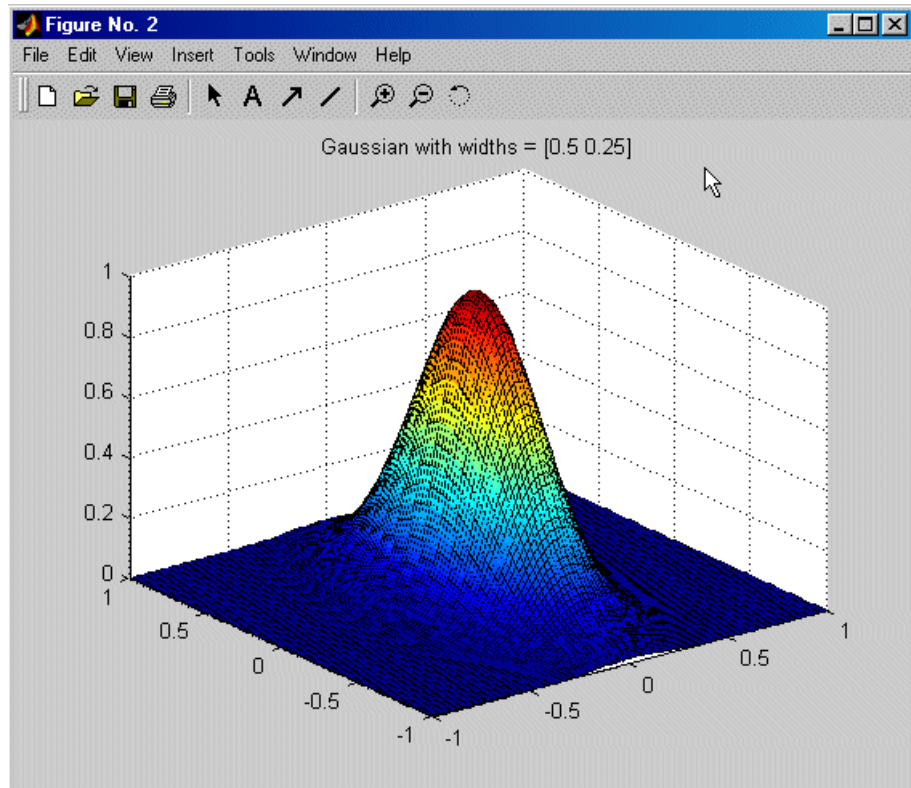
Initial upper bound on width — Upper bound on the width for the first zoom. Default is 20.

Display — If you select this check box, a stem plot of $\log_{10}(\text{GCV})$ against width is plotted. The best width is marked by a green asterisk.

WidPerDim

In the `WidPerDim` algorithm (Width Per Dimension), the radial basis functions are generalized. Rather than having a single width parameter, a different

width in each input factor can be used; that is, the level curves are elliptical rather than circular (or spherical, with more factors). The basis functions are no longer radially symmetric.



This can be especially helpful when the amount of variability varies considerably in each input direction. This algorithm offers more flexibility than `TrialWidths`, but is more computationally expensive.

You can set **Initial width** in the RBF controls on the Global Model Setup dialog. For most algorithms the **Initial width** is a single value. However, for `WidPerDim` (available in the **Width selection algorithm** pull down), you can specify a vector of widths to use as starting widths.

If supplying a vector of widths, there should be the same number as the number of global variables, and they must be in the same order as specified in the test plan. If you provide a single width, then all dimensions start off from the same initial width, but are likely to move from there to a vector of widths during model fitting.

An estimation of the time for the width per dimension algorithm is computed. This is given as a number of time units (as it depends on the machine). A time estimate of over 10 but less than 100 generates a warning. A time estimate of over 100 might take a prohibitively long amount of time (probably over five minutes on most machines). You have the option to stop execution and change some of the parameters to reduce the run time.

Fit Parameters

As for the `TrialWidths` algorithm.

Tree Regression

There are three parts to the tree regression algorithm for RBFs:

- Tree building
- Alpha selection
- Center selection.

Tree Building

The tree regression algorithm builds a regression tree from the data and uses the nodes (or panels) of this tree to infer candidate centers and widths for the RBF. The root panel of the tree corresponds to a hypercube that contains all of the data points. This panel is divided into two child panels such that each child contains the same amount of variation, as much as is possible. The child panel with the most variation is then split in a similar way. This process continues until there are no panels left to split, i.e., no childless panel has more than the minimum number of data points, or until the maximum number of panels has been reached. Each panel in the tree corresponds to a candidate center and the size of the panel determines the width that goes with that vector.

The size of the child panels can be based solely on the size of the parent panel or can be determined by shrinking the child panel onto the data that it contains.

Once you have selected Radial Basis Function in the Global Model Setup dialog, you can choose Tree Regression from the **Width Selection Algorithm** drop-down menu.

Click **Advanced** to open the Radial Basis Functions Options dialog to reach settings such as maximum number of panels and minimum number of data points per panel. To shrink child panels to fit the data, select the check box **Shrink panels to data**.

Alpha Selection Algorithm

The size for the candidate widths are not taken directly from the panel sizes: we need to scale the panel sizes to get the corresponding widths. This scaling factor is called alpha. The same scaling factor needs to be applied to every panel in the tree and to determine the optimal value of alpha we use an alpha selection algorithm.

You can choose the parameter Specify Alpha to specify the exact value of alpha to use, or you can select Trial Alpha. Trial Alpha is very similar to the Trial Widths algorithm. The only difference is that the trial alpha algorithm can specify how to space the values to search. Linear is the same as used by trial widths but Logarithmic searches more values near the lower range.

Click **Advanced** to open the Radial Basis Functions Options dialog to reach further settings such as bounds on alpha, number of zooms and number of trial alphas. Here you can select the **Display** check box to see the progress of the algorithm and the values of alpha trailed.

Center Selection Algorithm

The tree building generates candidate centers and the alpha selection generates candidate widths for these centers. The center selection chooses which of those centers to use.

Generic Center Selection is a center selection algorithm that knows nothing about the tree structure to be used. It uses Rols, which is very

fast way to choose centers and works in this case as well as the usual RBF cases. However, in this case the candidates for centers are not the data by the centers from the regression tree.

Tree-based center selection uses the regression tree. It is natural to use the regression tree to select centers because of the way it is built. In particular, the panel corresponding to the root node should be considered for selection before any of its children as it captures coarse detail while nodes at the leaves of the tree capture fine detail. This is what the Tree-based center selection does. You can also set the maximum number of centers.

Click **Advanced** to open the Radial Basis Functions Options dialog to reach the **Model selection criteria** setting. The **Model selection criteria** determines what function should be used as a measure of how good a model is. This can be BIC (Bayesian information criterion) or GCV (generalized cross-validation). BIC is usually less susceptible to over fitting.

The Stepwise menu is the same for all RBFs, see “Global Model Class: Radial Basis Function” on page 4-14.

Tree Regression and CentreExchange are the only algorithms that permit centers that are not located at the data points. This means that you do not see centers on model plots.

If you leave the **Alpha selection algorithm** at the default, Trial Alpha, you will see a progress dialog when you click **OK** to begin modeling. An example is shown.

| Alpha | Cost | Num Centers |
|-----------------|----------|-------------|
| 0.100000 | 1.294212 | 26 |
| 0.265915 | 1.245689 | 25 |
| 0.707107 | 1.249623 | 26 |
| 1.880302 | 1.312370 | 25 |
| 5.000000 | 0.957709 | 25 |
| zoom in | | |
| 3.440151 | 1.042787 | 25 |
| 4.042561 | 0.823649 | 25 |
| 4.750460 | 1.132577 | 25 |
| 5.582320 | 0.904029 | 25 |
| 6.559849 | 0.950359 | 25 |
| zoom in | | |
| 3.560633 | 0.894569 | 26 |
| 3.797906 | 0.952282 | 25 |
| 4.050991 | 0.823333 | 25 |

This is an example progress dialog of a Tree Regression RBF model fitting in progress. Here you can see each trial value of alpha with its calculated cost and the best number of centers with that value of alpha. The alpha value in red is the best so far. Alpha values no longer red, but in bold, are previous best values. You can then refine your model by zooming in on the best values for alpha and number of centers.


Reference: M. Orr, J. Hallam, K. Takezawa, A. Murray, S. Ninomiya, M. Oide, T. Leonard, "Combining Regression Trees and Radial Basis Function Networks," *International Journal of Neural Systems*, Vol. 10, No. 6 (2000) 453-465.

<http://www.anc.ed.ac.uk/~mjo/rbf.html>

Prune Functionality

You can use the Prune function to reduce the number of centers in a radial basis function network. This helps you decide how many centers are needed.


To use the Prune facility:

- 1 Select an RBF global model in the model tree.
- 2 Either click the  toolbar button or select the menu item **Model > Utilities > Prune**.

The Number of Centers Selector dialog appears.

The graphs show how the fit quality of the network builds up as more RBFs are added. It makes use of the fact that most of the center selection algorithms are greedy in nature, and so the order in which the centers were selected roughly reflects the order of importance of the basis functions.

The default fit criteria are the logarithms of PRESS, GCV, RMSE, and Weighted PRESS. Additional options are determined by your selections in Summary Statistics. See “Model Menu (Global Level)” on page 4-104. Weighted PRESS penalizes having more centers, and choosing number of centers to minimize weighted PRESS is often a good option.



All four criteria in this typical example indicate the same minimum at eight centers.

If the graphs all decrease, as in the preceding example, this suggests that the maximum number of centers is too small, and the number of centers should be increased.

Clicking the **Minimize** button selects the number of centers that minimizes the criterion selected in the drop-down menu on the left. It is good if this value also minimizes all the other criteria. The **Clear** button returns to the previous selection.

Note that reducing the number of centers using Prune only refits the linear parameters (RBF weights). The nonlinear parameters (center locations, width, and lambda) are not adjusted. You can perform a cheap width refit on exiting the dialog by selecting the **Refit widths on close** check box. If a network has been pruned significantly, you should use the Update Model Fit toolbar button. This performs a full refit of all the parameters.

Statistics

Let A be the matrix such that the weights are given by $\beta = A^{-1}X'y$ where X is the regression matrix. The form of A varies depending on the basic fit algorithm employed.

In the case of ordinary least squares, we have $A = X'X$.

For ridge regression (with regularization parameter λ), A is given by $A = X'X + \lambda I$

Next is the Rols algorithm. During the Rols algorithm X is decomposed using the Gram-Schmidt algorithm to give $X = WB$, where W has orthogonal columns and B is upper triangular. The corresponding matrix A for Rols is then $A = X'X + \lambda B'B$.

The matrix $H := XA^{-1}X'$ is called the hat matrix, and the leverage of the i th data point h_i is given by the i th diagonal element of H . All the statistics derived from the hat matrix, for example, PRESS, studentized residuals, confidence intervals, and Cook's distance, are computed using the hat matrix appropriate to the particular fit algorithm.

Similarly PEV, given in the "Toolbox Terms and Statistics Definitions" on page 4-192 as

$$PEV(x) = \text{var}y(\hat{y}) = x(X'X)^{-1}x'TMSE$$

becomes

$$PEV(x) = \text{var}y(\hat{y}) = xA^{-1}x'TMSE$$

PEV is computed using the form of A appropriate to the particular fit algorithm (ordinary least squares, ridge or rols).

GCV Criterion

Generalized cross-validation (GCV) is a measure of the goodness of fit of a model to the data that is minimized when the residuals are small, but not so small that the network has overfitted the data. It is easy to compute, and

networks with small GCV values should have good predictive capability. It is related to the PRESS statistic.

The definition of GCV is given by Orr (4, see “References” on page 5-31).

$$GCV = \frac{N(y'P^2y)}{(\text{trace}(P))^2}$$

where y is the target vector, N is the number of observations, and P is the projection matrix, given by $I - XA^{-1}X^T$. See “Statistics” on page 5-28 for definition of A .

An important feature of using GCV as a criterion for determining the optimal network in our fit algorithms is the existence of update formulas for the regularization parameter λ . These update formulas are obtained by differentiating GCV with respect to λ and setting the result to zero. That is, they are based on gradient-descent.

This gives the general equation (from Orr, 6, “References” on page 5-31)

$$y'P\frac{\partial(Py)}{\partial\lambda}\text{trace}(P) = (y'P^2y)\frac{\partial(\text{trace}(P))}{\partial\lambda}$$

We now specialize these formulas to the case of ridge regression and to the Rols algorithm.

GCV for Ridge Regression

It is shown in Orr (4), and stated in Orr (5, see “References” on page 5-31) that for the case of ridge regression GCV can be written as

$$GCV = \frac{N(e'e)}{(N-p)^2}$$

where p is the “effective number of parameters” that is given by

$$p = \text{NumTerms} - \lambda\text{trace}(A^{-1})$$

where NumTerms is the number of terms included in the model.

For RBFs, 'p' is the effective number of parameters, that is, the number of terms minus an adjustment to take into account the smoothing effect of lambda in the fitting algorithm. When lambda = 0, the effective number of parameters is the same as the number of terms.

$$\lambda = \frac{\eta}{N-p} \frac{(e'e)}{(\beta'A^{-1}\beta)}$$

The formula for updating λ is given by $\eta = tr(A^{-1} - \lambda A^{-2})$ where

In practice, the preceding formulas are not used explicitly in Orr (5, see "References" on page 5-31). Instead a singular value decomposition of X is made, and the formulas are rewritten in terms of the eigenvalues and eigenvectors of the matrix XX'. This avoids taking the inverse of the matrix A, and it can be used to cheaply compute GCV for many values of λ . See "Statistics" on page 5-28 for definition of A.

GCV for Rols

In the case of Rols, the components for the formula

$$GCV = \frac{N(y'P^2y)}{(trace(P))^2}$$

are computed using the formulas given in Orr [6; see "References" on page 5-31]. Recall that the regression matrix is factored during the Rols algorithm into the product X = WB. Let w_j denote the jth column of W, then we have

$$P^2y = y'y - \sum_{j=1}^N \frac{(2\lambda + w_j'w_j)(y'w_j)^2}{(\lambda + w_j'w_j)^2}$$

and the "effective number of parameters" is given by

$$Trace(P) = NumTerms - \sum_{j=1}^N \frac{w_j'w_j}{(\lambda + w_j'w_j)}$$

This is equivalent to 'p' (the effective number of parameters) defined in "GCV for Ridge Regression" on page 5-29.

The reestimation formula for λ is given by

$$\lambda = \frac{\eta}{\text{Trace}(P)} \frac{y'P^2y}{(\beta'A^{-1}\beta)^2} \text{ where additionally}$$

$$\eta = \sum_{j=1}^N \frac{w_j'w_j}{(\lambda + w_j'w_j)^2} \quad \text{and} \quad 3'A^{-1}\beta = \sum_{j=1}^N \frac{(y'w_j)^2}{(\lambda + w_j'w_j)^3}$$

Note that these formulas for Rols do not require the explicit inversion of A. See "Statistics" on page 5-28 for definition of A.

References

- 1 Chen, S., Chng, E.S., Alkadhimi, Regularized Orthogonal Least Squares Algorithm for Constructing Radial Basis Function Networks, Int J. Control, 1996, Vol. 64, No. 5, pp. 829-837.
- 2 Hassoun, M., Fundamentals of Artificial Neural Networks, MIT, 1995.
- 3 Orr, M., Introduction to Radial Basis Function Networks, available from <http://www.anc.ed.ac.uk/~mjo/rbf.html>.
- 4 Orr, M., Optimizing the Widths of Radial Basis Functions, available from <http://www.anc.ed.ac.uk/~mjo/rbf.html>.
- 5 Orr, M., Regularisation in the Selection of Radial Basis Function Centers, available from <http://www.anc.ed.ac.uk/~mjo/rbf.html>.
- 6 Wendland, H., Piecewise Polynomials, Positive Definite and Compactly Supported Radial Basis Functions of Minimal Degree, Advances in Computational Mathematics 4 (1995), pp. 389-396.

Hybrid Radial Basis Functions

Hybrid RBFs combine a radial basis function model with more standard linear models such as polynomials or hybrid splines. The two parts are added together to form the overall model. This approach offers the ability to combine a priori knowledge, such as the expectation of quadratic behavior in one of the variables, with the nonparametric nature of RBFs.

The model setup GUI for hybrid RBFs has a top **Set Up** button, where you can set the fitting algorithm and options. The interface also has two tabs, one to specify the radial basis function part, and one for the linear model part.

Width Selection Algorithm: TrialWidths

This is the same algorithm as is used in ordinary RBFs, that is, a guided search for the best width parameter.

Lambda and Term Selection Algorithms: Interlace

This algorithm is a generalization of StepItRols for RBFs. The algorithm chooses radial basis functions and linear model terms in an interlaced way, rather than in two steps. At each step a forward search procedure is performed to select the radial basis function (with a center chosen from within the set of data points) or the linear model term (chosen from the ones specified in the linear model setup pane) that decreases the regularized error the most. This process continues until the maximum number of terms is chosen. The first few terms are added using the stored value of lambda, until the **Number of terms to add before updating** has been reached. Subsequently lambda is iterated after each center is added to improve GCV.

The fit options for this algorithm are as follows:

- **Maximum number of terms:** Maximum number of terms that will be chosen. The default is a quarter of the data points, or 25, whichever is smaller.
- **Maximum number of centers:** Maximum number of terms that can be radial basis functions. The default is the same as the maximum number of terms.

- **Percentage of data to be candidate centers:** Percentage of the data points that are available to be chosen as centers. The default is 100% when the number of data points is ≤ 200 .
- **Number of terms to add before updating:** How many terms to add before updating lambda begins.
- **Minimum change in $\log_{10}(\text{GCV})$:** Tolerance.
- **Maximum no. times $\log_{10}(\text{GCV})$ change is minimal:** Number of steps in a row that the change in $\log_{10}(\text{GCV})$ can be less than the tolerance before the algorithm terminates.


Lambda and Term Selection Algorithms: Two-Step

This algorithm starts by fitting the linear model specified in the linear model pane, and then fits a radial basis function network to the residual. You can specify the linear model terms to include in the usual way using the term selector. If desired, you can activate the stepwise options. In this case, after the linear model part is fitted, some of the terms are automatically added or removed before the RBF part is fitted. You can choose the algorithm and options that are used to fit the nonlinear parameters of the RBF by clicking the **Set Up** button in the RBF training options.

Tips for Modeling with Radial Basis Functions

Plan of Attack


Determine which parameters have the most impact on the fit by following these steps:

- 1 Fit the default RBF. Remove any obvious outliers.
- 2 Get a rough idea of how many RBFs are going to be needed. If a center coincides with a data point, it is marked with a magenta asterisk on the Predicted/Observed plot. You can view the location of the centers in graphical and table format by using the **View Centers** toolbar button . If you remove an outlier which coincided with a center (marked with an asterisk), refit by clicking **Update Fit** in the toolbar.
- 3 Try with more than one kernel. You can alter the parameters in the fit by clicking the **Set Up** button in the Model Selection dialog.
- 4 Decide on the main width selection algorithm. Try with both TrialWidths and WidPerDim algorithms.
- 5 Determine which types of kernel look most hopeful.
- 6 Narrow the corresponding width range to search over.
- 7 Decide on the center selection algorithm.
- 8 Decide on the lambda-selection algorithm.
- 9 Try changing the parameters in the algorithms.
- 10 If any points appear to be possible outliers, try fitting the model both with and without those points.

If at any stage you decide on a change that has a big impact (such as removal of an outlier), then you should repeat the previous steps to determine if this would affect the path you have chosen.


See “Fitting Routines” on page 5-11 for details on all the fit parameters.

The Model Browser has a quick option for comparing all the different RBF kernels.

- 1 After fitting the default RBF, select the RBF global model in the model tree.
- 2 Click the **Build Models** toolbar icon .
- 3 Select the **RBF Kernels** icon in the Build Models dialog that appears and click **OK**.

One of each kernel is built as a selection of child nodes of the current RBF model.


How Many RBFs to Use


- The main parameter that you must adjust in order to get a good fit with an RBF is the maximum number of centers. This is a parameter of the center selection algorithm, and is the maximum number of centers/RBFs that is chosen.
- Usually the maximum number of centers is the number of RBFs that are actually selected. However, sometimes fewer RBFs are chosen because the (regularized) error has fallen below the tolerance before the maximum was reached.
- You should use a number of RBFs that is significantly less than the number of data points, otherwise there are not enough degrees of freedom in the error to estimate the predictive quality of the model. That is, you cannot tell if the model is useful if you use too many RBFs. We would recommend an upper bound of 60% on the ratio of number of RBFs to number of data points. Having 80 centers when there are only 100 data points might seem to give a good value of PRESS, but when it comes to validation, it can sometimes become clear that the data has been overfitted, and the predictive capability is not as good as PRESS would suggest.
- One strategy for choosing the number of RBFs is to fit more centers than you think is needed (say 70 out of 100), then use the Prune toolbar button  to reduce the number of centers in the model. After pruning the network, make a note of the reduced number of RBFs. Try fitting the model again with the maximum number of centers set to this reduced

number. This recalculates the values of the nonlinear parameters (width and lambda) to be optimal for the reduced number of RBFs.

- One strategy for the use of Stepwise is to use it to minimize PRESS as a final fine-tuning for the network, once pruning has been done. Whereas Prune only allows the last RBF introduced to be removed, Stepwise allows any RBF to be taken out.
- Do not focus solely on PRESS as a measure of goodness of fit, especially at large ratios of RBFs to data points. Take $\log_{10}(\text{GCV})$ into account also.

Width Selection Algorithms

- Try both TrialWidths and WidPerDim. The second algorithm offers more flexibility, but is more computationally expensive. View the width values in each direction to see if there is significant difference, to see whether it is worth focusing effort on elliptical basis functions (use the View Model toolbar button ).



- If with a variety of basis functions the widths do not vary significantly between the dimensions, and the PRESS/GCV values are not significantly

improved using `WidPerDim` over `TrialWidths`, then focus on `TrialWidths`, and just return to `WidPerDim` to fine-tune in the final stages.

- Turn the **Display** option on in `TrialWidths` to see the progress of the algorithm. Watch for alternative regions within the width range that have been prematurely neglected. The output $\log_{10}(\text{GCV})$ in the final zoom should be similar for each of the widths tried; that is, the output should be approximately flat. If this is not the case, try increasing the number of zooms.
- In `TrialWidths`, for each type of RBF, try to narrow the initial range of widths to search over. This might allow the number of zooms to be reduced.

Which RBF to Use

- It is hard to give rules of thumb on how to select the best RBF, as the best choice is highly data-dependent. The best guideline is to try all of them with both top-level algorithms (`TrialWidths` and `WidPerDim`) and with a sensible number of centers, compare the PRESS/GCV values, then focus on the ones that look most hopeful.
- If multiquadrics and thin-plate splines give poor results, it is worth trying them in combination with low-order polynomials as a hybrid spline. Try supplementing multiquadrics with a constant term and thin-plate splines with linear (order 1) terms. See “Hybrid Radial Basis Functions” on page 5-32.
- Watch out for conditioning problems with Gaussian kernels (say condition number $> 10^8$).
- Watch out for strange results with Wendland’s functions when the ratio of the number of parameters to the number of observations is high. When these functions have a very small width, each basis function only contributes to the fit at one data point. This is because its support only encompasses the one basis function that is its center. The residuals will be zero at each of the data points chosen as a center, and large at the other data points. This scenario can indicate good RMSE values, but the predictive quality of the network will be poor.

Lambda Selection Algorithms

Lambda is the regularization parameter.

- `IterateRols` updates the centers after each update of lambda. This makes it more computationally intensive, but potentially leads to a better combination of lambda and centers.
- `StepItRols` is sensitive to the setting of **Number of centers to add before updating**. Switch the **Display** option on to view how $\log_{10}(\text{GCV})$ reduces as the number of centers builds up.
- Examine the plots produced from the lambda selection algorithm, ignoring the warning “An excessive number of plots will be produced.” Would increasing the tolerance or the number of initial test values for lambda lead to a better choice of lambda?

Center Selection Algorithms

- On most problems, `Rols` seems to be the most effective.
- If less than the maximum number of centers are being chosen, and you want to force the selection of the maximum number, reduce the tolerance to epsilon (`eps`).
- `CenterExchange` is very expensive, and you should not use this on large problems. In this case, the other center selection algorithms that restrict the centers to be a subset of the data points might not offer sufficient flexibility.

General Parameter Fine-Tuning

- Try `Stepwise` after pruning, then update the model fit with the new maximum number of centers set to the number of terms left after `Stepwise`.
- Update the model fit after removal of outliers; use the toolbar button.

Hybrid RBFs

- Go to the linear part pane and specify the polynomial or spline terms that you expect to see in the model.

Fitting too many non-RBF terms is made evident by a large value of lambda, indicating that the underlying trends are being taken care of by the linear part. In this case, you should reset the starting value of lambda (to say 0.001) before the next fit.

How to Find RBF Model Formula

With any model you can use the **View Model** toolbar button or **View > Model Definition** (or keyboard shortcut **CTRL+V**) to see the details of the current model. The Model Viewer dialog appears. Here for any RBF model you can see the kernel type, number of centers, width and regularization parameter.

However to specify the formula of an RBF model completely, you also need to give the locations of the centers, and the height of each basis function. The center location information is available in the “View Centers” dialog and the coefficients can be found in the “Stepwise” window. Note these values are all in coded units.

B

Box-Cox transformation 4-176

C

constraints manager 2-43

correlation models 4-67

covariance modeling 4-66

D

data 3-3

import wizard 3-14

loading from file 3-13

loading from the Workspace 3-15

loading user-defined functions 3-49

merging 3-14 3-17

wizard 3-36

data editor 3-6

data selection window 3-25

datum models 4-72

definitions 4-192

design

adding points 2-35

classical 2-16

constraints 2-43

deleting points 2-39

fixing points 2-39

importing 2-42

optimal 2-25

saving 2-42

sorting points 2-39

space filling 2-19

styles 2-4

tree 2-7

Design Editor 2-3

display options 2-9

displays 2-5

Menus 2-9

designing experiments 4-69

F

filter editor 3-28

G

global level 4-93

global models

free knot spline 4-22

hybrid RBF 4-20

linear

hybrid spline 4-9

multiple 4-21

polynomial 4-8

neural network 4-24

radial basis functions 4-14

transient 4-26

user-defined 4-26

global regression 4-191

H

hierarchical models 4-181

I

inputs

setting up 4-4

interaction 4-12

L

linear model statistics 4-97

local level 4-74

local model setup 4-26

local models

average fit 4-44

B-spline 4-34

growth models 4-36

linear 4-42

multiple models 4-45

polynomial 4-28

- polynomial spline 4-29
- transient 4-55
- truncated power series 4-32
- user-defined 4-49

M

- MLE 4-141
- model building process overview 4-170
- model evaluation 4-148
- model selection 4-111
- model setup
 - global 4-6
 - local 4-26
 - response 4-71
- model tree 1-19
- models
 - setting up 4-3

N

- neural nets 4-24
- new test plan templates 1-26

O

- one-stage model
 - instant setup 1-6

P

- Prediction Error Variance (PEV) viewer 2-54
- project level 1-11

R

- radial basis functions

- fitting routines 5-11
 - guide 5-3
 - hybrid RBFs 5-32
 - prune functionality 5-25
 - statistics 5-28
 - tips for modeling 5-34
- response level 4-145

S

- stepwise 4-13
- stepwise regression techniques 4-166

T

- test groupings 3-33
- test plan level 1-29
- test plans 1-25
- transforms 4-68
- two-stage model
 - instant setup 1-6

V

- variable editor 3-25
- view
 - cross-section 4-136
 - likelihood 4-131
 - predicted/observed 4-125
 - residuals 4-135
 - response surface 4-127
 - RMSE 4-133
 - tests 4-122
- viewing designs 4-69