


System Identification Toolbox 7

User's Guide

Lennart Ljung



MATLAB[®]
& **SIMULINK[®]**

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.

System Identification Toolbox

© COPYRIGHT 1988–2007 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, SimBiology, SimHydraulics, SimEvents, and xPC TargetBox are registered trademarks and The MathWorks, the L-shaped membrane logo, Embedded MATLAB, and PolySpace are 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

April 1988	First printing	
July 1991	Second printing	
May 1995	Third printing	
November 2000	Fourth printing	Revised for Version 5.0 (Release 12)
April 2001	Fifth printing	
July 2002	Online only	Revised for Version 5.0.2 (Release 13)
June 2004	Sixth printing	Revised for Version 6.0.1 (Release 14)
March 2005	Online only	Revised for Version 6.1.1 (Release 14SP2)
September 2005	Seventh printing	Revised for Version 6.1.2 (Release 14SP3)
March 2006	Online only	Revised for Version 6.1.3 (Release 2006a)
September 2006	Online only	Revised for Version 6.2 (Release 2006b)
March 2007	Online only	Revised for Version 7.0 (Release 2007a)
September 2007	Online only	Revised for Version 7.1 (Release 2007b)

About the Developers

System Identification Toolbox is developed in association with the following leading researchers in the system identification field:

Lennart Ljung. Professor Lennart Ljung is with the Department of Electrical Engineering at Linköping University in Sweden. He is a recognized leader in system identification and has published numerous papers and books in this area.

Qinghua Zhang. Dr. Qinghua Zhang is a researcher at Institut National de Recherche en Informatique et en Automatique (INRIA) and at Institut de Recherche en Informatique et Systèmes Aléatoires (IRISA), both in Rennes, France. He conducts research in the areas of nonlinear system identification, fault diagnosis, and signal processing with applications in the fields of energy, automotive, and biomedical systems.

Peter Lindskog. Dr. Peter Lindskog is employed by NIRA Dynamics AB, Sweden. He conducts research in the areas of system identification, signal processing, and automatic control with a focus on vehicle industry applications.

Anatoli Juditsky. Professor Anatoli Juditsky is with the Laboratoire Jean Kuntzmann at the Université Joseph Fourier, Grenoble, France. He conducts research in the areas of nonparametric statistics, system identification, and stochastic optimization.

Working with the System Identification Tool GUI

1

How to Use the System Identification Tool GUI	1-2
Starting and Managing GUI Sessions	1-3
What Is a Session?	1-3
Starting a New Session	1-4
Overview of the System Identification Tool GUI	1-5
Opening a Saved Session	1-6
Saving, Merging, and Closing Sessions	1-6
Deleting a Session	1-7
Getting Help in the GUI	1-7
Exiting the System Identification Tool GUI	1-8
Managing Data Sets and Models in the GUI	1-9
Adding Data to the GUI	1-9
Adding Models to the GUI	1-10
Specifying Estimation and Validation Data	1-11
Importing Models into the System Identification Tool GUI	1-11
Viewing Data and Model Properties	1-13
Organizing Data and Model Icons	1-14
Deleting Data and Models	1-15
Exporting Data and Models to the MATLAB Workspace Browser	1-16
Plots Data and Models in the GUI	1-18
Plotting Data in the GUI	1-18
Plotting Models in the GUI	1-19
Identifying Data Sets and Models on Plots	1-20
Changing and Restoring Default Axis Limits	1-21
Selecting Measured and Noise Channels in Plots	1-23
Grid, Line Styles, and Redrawing Plots	1-24
Opening a Plot in a MATLAB Figure Window	1-25
Printing Plots	1-25

Customizing the System Identification Tool GUI	1-26
Types of GUI Customization	1-26
Displaying Warnings While You Work	1-26
Saving Session Preferences	1-26
Modifying idlayout.m	1-27

Preparing Data

2

Ways to Prepare Data for Modeling	2-3
Data Requirements	2-6
Types of Data You Can Model	2-6
Data with Uniform and Nonuniform Sampling Rates	2-7
Representing Time-Domain Data	2-7
Representing Time-Series Data	2-8
Representing Frequency-Domain Data	2-9
Representing Frequency-Response Data	2-10
Importing Data into MATLAB	2-13
Importing Data into the GUI	2-14
Different Ways to Create New Data Sets in the GUI	2-14
Importing Data into the System Identification Tool GUI ..	2-14
Specifying the Data Sampling Interval	2-19
Preprocessing Data Using Quick Start	2-20
Creating Data Sets from Channel Subsets	2-20
Creating Multiexperiment Data Sets in the GUI	2-23
Renaming Data and Changing Display Color	2-30
Creating iddata Objects	2-32
iddata Constructor	2-32
iddata Properties	2-35
Creating Multiexperiment Data at the Command Line ...	2-38
Subreferencing iddata Objects	2-40
Modifying Time and Frequency Vectors	2-44
Naming, Adding, and Removing Data Channels	2-48
Concatenating iddata Objects	2-50
See Also	2-51

Creating idfrd Objects	2-52
idfrd Constructor	2-52
idfrd Properties	2-53
Subreferencing idfrd Objects	2-55
Concatenating idfrd Objects	2-56
See Also	2-59
Getting Advice About Your Data	2-60
Plotting Data	2-62
Supported Plots for Data	2-62
Plotting Data in the System Identification Tool GUI	2-62
How to Plot Data at the Command Line	2-68
Handling Missing Data and Outliers	2-71
Working with Missing Data	2-71
Working with Outliers	2-72
Example – Extracting and Modeling Specific Data Segments	2-73
See Also	2-74
Detrending Data	2-75
What Is Detrending Data?	2-75
Detrending Data for Nonlinear Versus Linear Models	2-75
When to Subtract the Mean Values	2-76
When to Subtract Linear Trends	2-77
How to Detrend Data Using the GUI	2-77
How to Detrend Data Using detrend	2-78
See Also	2-79
Resampling Data	2-80
What Is Resampling?	2-80
Resampling Data Using the GUI	2-81
Resampling Data at the Command Line	2-81
Resampling Data Without Aliasing Effects	2-83
See Also	2-86
Filtering Data	2-87
Supported Filters in System Identification Toolbox	2-87
Deciding to Prefilter Your Data	2-87
How to Filter Data Using the GUI	2-88

How to Filter Data at the Command Line	2-91
See Also	2-94
Selecting Data	2-95
Why Select Subsets of Data	2-95
Selecting Data Using the GUI	2-96
Selecting Data at the Command Line	2-98
Creating Data Using Simulation	2-99
Commands for Generating and Simulating Data	2-99
Example – Creating Data with Periodic Inputs	2-100
Example – Simulating Model Output at the Command Line	2-101
Simulating Data Using Other Mathworks Products	2-102
Transforming Between Time- and Frequency-Domain Data	2-103
Transforming Data Domain in the GUI	2-103
Transforming Data Domain at the Command Line	2-110
Handling Complex-Valued Data	2-115
Supported Operations for Complex Data	2-115
Manipulating Complex iddata Signals	2-115

Creating and Manipulating Models

3

Creating Model Objects	3-2
Definition of a Model Object	3-2
Basic Object-Oriented Concepts	3-3
Command for Creating Model Objects	3-4
Commands for Model Estimation	3-5
Example – Estimating Models Using armax	3-7
Model Properties	3-13
See Also	3-19
Subreferencing Models	3-20
What Is Subreferencing?	3-20

Limitation on Supported Models	3-20
Subreferencing Specific Measured Channels	3-20
Subreferencing Measured and Noise Models	3-21
Treating Noise Channels as Measured Inputs	3-23
Concatenating Model Objects	3-25
About Concatenating Models	3-25
Limitation on Supported Models	3-25
Horizontal Concatenation of Model Objects	3-26
Vertical Concatenation of Model Objects	3-26
Concatenating Noise Spectral Data of idfrd Objects	3-27
See Also	3-28
Merging Model Objects	3-29

Estimating Nonparametric Models

4

Estimating Time-Domain Correlation Models	4-2
What Is Time-Domain Correlation Analysis?	4-2
Data Supported by Correlation Analysis	4-3
How to Estimate Correlation Models Using the GUI	4-3
How to Estimate Correlation Models at the Command	
Line	4-4
How to Compute Response Values	4-6
How to Identify Delay Using Transient-Response Plots ...	4-6
Algorithm for Correlation Analysis	4-8
Estimating Frequency-Response Models	4-10
What Is a Frequency-Response Model?	4-10
Data Supported by Frequency-Response Models	4-11
How to Estimate Frequency-Response Models in the	
GUI	4-11
How to Estimate Frequency-Response Models at the	
Command Line	4-13
Options for Computing Spectral Models	4-13
Options for Frequency Resolution	4-14
Spectral Analysis Algorithm	4-16
Understanding Spectrum Normalization	4-19

Estimating Process Models	5-2
What Is a Process Model?	5-2
Data Supported by a Process Model	5-3
How to Estimate Process Models Using the GUI	5-3
Estimating Process Models at the Command Line	5-9
Options for Specifying the Process-Model Structure	5-15
Options for Multiple-Input Models	5-16
Options for the Disturbance Model Structure	5-17
Options for Frequency-Weighing Focus	5-18
Options for Initial States	5-19
Estimating Polynomial Models	5-21
What Are Black-Box Polynomial Models?	5-21
Data Supported by Polynomial Models	5-28
Preliminary Step – Estimating Model Orders and Input Delays	5-29
How to Estimate Polynomial Models in the GUI	5-37
How to Estimate Polynomial Models at the Command Line	5-40
Options for Multiple-Input and Multiple-Output ARX Orders	5-44
Option for Frequency-Weighing Focus	5-45
Options for Initial States	5-46
Algorithms for Estimating Polynomial Models	5-46
Estimating State-Space Models	5-48
What Are State-Space Models?	5-48
Data Supported by State-Space Models	5-52
Supported State-Space Parameterizations	5-53
Preliminary Step – Estimating State-Space Model Orders	5-54
How to Estimate State-Space Models in the GUI	5-59
How to Estimate State-Space Models at the Command Line	5-62
How to Estimate Free-Parameterization State-Space Models	5-65
How to Estimate State-Space Models with Canonical Parameterization	5-66

How to Estimate State-Space Models with Structured Parameterization	5-68
How to Estimate the State-Space Equivalent of ARMAX and OE Models	5-75
Options for Frequency-Weighing Focus	5-75
Options for Initial States	5-76
Algorithms for Estimating State-Space Models	5-76
Refining Linear Parametric Models	5-78
When to Refine Models	5-78
What You Specify to Refine a Model	5-78
How to Refine Linear Parametric Models in the GUI	5-79
How to Refine Linear Parametric Models at the Command Line	5-80

Estimating Time-Series Models

6

What Are Time-Series Models?	6-2
Preparing Time-Series Data	6-3
Estimating Time-Series Power Spectra	6-4
How to Estimate Time-Series Power Spectra Using the GUI	6-4
How to Estimate Time-Series Power Spectra at the Command Line	6-5
Estimating AR and ARMA Models	6-7
Definition of AR and ARMA Models	6-7
Estimating Polynomial Models in the GUI	6-7
Estimating AR and ARMA Models at the Command Line	6-10
Estimating State-Space Time-Series Models	6-12
Definition of State-Space Time-Series Model	6-12
Estimating State-Space Models at the Command Line ...	6-12

Example – Estimating Time Series at the Command Line	6-14
Estimating Nonlinear Models for Time-Series Data ...	6-15

Estimating Linear Models Using Recursive Techniques

7

What Is Recursive Estimation?	7-2
Commands for Recursive Estimation	7-3
Algorithms for Recursive Estimation	7-6
Types of Recursive Estimation Algorithms	7-6
General Form of Recursive Estimation Algorithm	7-6
Kalman Filter Algorithm	7-8
Forgetting Factor Algorithm	7-10
Unnormalized and Normalized Gradient Algorithms	7-11
Data Segmentation	7-14

Estimating Nonlinear Black-Box Models

8

Supported Nonlinear Black-Box Models	8-2
Supported Data for Estimating Nonlinear Black-Box Models	8-3
Estimating Nonlinear ARX Models	8-4
Supported Data for Nonlinear ARX Models	8-4
Definition of the Nonlinear ARX Model	8-4
Using Regressors	8-5

Nonlinearity Estimators for Nonlinear ARX Models	8-8
How to Estimate Nonlinear ARX Models in the GUI	8-8
How to Estimate Nonlinear ARX Models at the Command Line	8-10
Estimating Hammerstein-Wiener Models	8-14
Supported Data for Estimating Hammerstein-Wiener Models	8-14
Definition of the Hammerstein-Wiener Model	8-14
Nonlinearity Estimators for Hammerstein-Wiener Models	8-16
How to Estimate Hammerstein-Wiener Models in the GUI	8-17
How to Estimate Hammerstein-Wiener Models at the Command Line	8-18
Supported Nonlinearity Estimators	8-22
Types of Nonlinearity Estimators	8-22
Creating Custom Nonlinearities	8-23
Refining Nonlinear Black-Box Models	8-25
How to Refine Nonlinear Black-Box Models in the GUI ...	8-25
How to Refine Nonlinear Black-Box Models at the Command Line	8-26
Next Steps After Estimating Nonlinear Black-Box Models	8-27

Estimating Grey-Box Models

9

Supported Grey-Box Models	9-2
Data Supported by Grey-Box Models	9-3
Before You Begin	9-4

Estimating Linear Grey-Box Models	9-5
Specifying the Linear Grey-Box Model Structure	9-5
Example – Representing a Grey-Box Model in an M-File ..	9-6
Example – Estimating a Continuous-Time Grey-Box Model for Heat Diffusion	9-8
Example – Estimating a Discrete-Time Grey-Box Model with Parameterized Disturbance	9-11
Estimating Nonlinear Grey-Box Models	9-13
Supported Nonlinear Grey-Box Models	9-13
Nonlinear Grey-Box Demos and Examples	9-13
Specifying the Nonlinear Grey-Box Model Structure	9-13
Constructing the idnlgrey Object	9-15
Using pem to Estimate Nonlinear Grey-Box Models	9-16
Options for the Estimation Algorithm	9-16
After Estimating Grey-Box Models	9-19

Validating and Analyzing Models

10

Overview of Model Validation and Plots	10-3
When to Validate Models	10-3
Ways to Validate Models	10-3
Data for Validating Models	10-4
Supported Model Plots	10-5
Getting Advice About Models	10-6
Using Model Output Plots to Validate Models	10-7
Supported Model Types	10-7
What Does a Model Output Plot Show?	10-7
Choosing Simulated or Predicted Output	10-8
How to Plot Model Output Using the GUI	10-10
Displaying the Confidence Interval	10-12
How to Plot Model Output at the Command Line	10-13
Using Residual Analysis Plots to Validate Models	10-15
What Is Residual Analysis?	10-15
Supported Model Types	10-16

What Does the Residuals Plot Show?	10-16
Displaying the Confidence Interval	10-17
How to Plot Residuals Using the GUI	10-18
How to Plot Residuals at the Command Line	10-20
Using Impulse- and Step-Response Plots to Validate	
Models	10-21
Supported Models	10-21
How Transient Response Helps to Validate Models	10-21
What Does a Transient Response Plot Show?	10-22
How to Plot Impulse and Step Response Using the GUI ..	10-23
Displaying the Confidence Interval	10-26
How to Plot Impulse and Step Response at the Command	
Line	10-27
Using Frequency-Response Plots to Validate Models ..	10-29
What Is Frequency Response?	10-29
How Frequency Response Helps to Validate Models	10-30
What Does a Frequency-Response Plot Show?	10-31
How to Plot Bode Plots Using the GUI	10-32
How to Plot Bode and Nyquist Plots at the Command	
Line	10-35
Creating Noise-Spectrum Plots	10-37
Supported Models	10-37
What Does a Noise Spectrum Plot Show?	10-37
Displaying the Confidence Interval	10-38
How to Plot the Noise Spectrum Using the GUI	10-39
How to Plot the Noise Spectrum at the Command Line ...	10-42
Using Pole-Zero Plots to Validate Models	10-44
Supported Models	10-44
What Does a Pole-Zero Plot Show?	10-44
How to Plot Model Poles and Zeros Using the GUI	10-45
How to Plot Poles and Zeros at the Command Line	10-47
Reducing Model Order Using Pole-Zero Plots	10-48
Using Nonlinear ARX Plots to Validate Models	10-49
About Nonlinear ARX Plots	10-49
How to Plot Nonlinear ARX Plots Using the GUI	10-49
Configuring the Nonlinear ARX Plot	10-50
Axis Limits, Legend, and 3-D Rotation	10-51

How to plot Nonlinear ARX Plots at the Command Line . .	10-52
Using Hammerstein-Wiener Plots to Validate Models . .	10-53
About Hammerstein-Wiener Plots	10-53
How to Create Hammerstein-Wiener Plots in the GUI . . .	10-53
How to Plot Hammerstein-Wiener Plots at the Command Line	10-55
Plotting Nonlinear Block Characteristics	10-55
Plotting Linear Block Characteristics	10-56
Using Akaike's Criteria to Validate Models	10-58
Definition of FPE	10-58
Computing FPE	10-59
Definition of AIC	10-59
Computing AIC	10-60
Computing Model Uncertainty	10-61
Why Analyze Model Uncertainty	10-61
What Is Model Covariance?	10-61
Viewing Model Uncertainty Information	10-62
Extracting Parameter Values from Linear Models	10-64
Extracting Dynamic Model and Noise Model Separately	10-66
Troubleshooting Models	10-68
About Troubleshooting Models	10-68
Model Order Is Too High or Too Low	10-68
Nonlinearity Estimator Produces a Poor Fit	10-69
Substantial Noise in the System	10-70
Unstable Models	10-70
Missing Input Variables	10-71
Complicated Nonlinearities	10-72
Next Steps After Getting an Accurate Model	10-73

Transforming Model Representation

11

Transforming Between Discrete-Time and Continuous-Time Representations	11-2
Why Transform Between Continuous and Discrete Time?	11-2
Using the c2d, d2c, and d2d Commands	11-2
Specifying Intersample Behavior	11-4
How d2c Handles Input Delays	11-4
Effects on the Noise Model	11-5
Transforming Between Linear Model Representations	11-7

Simulating and Predicting Model Output

12

Simulating Versus Predicting Output	12-2
Simulation and Prediction in the GUI	12-4
Example – Simulating Model Output with Noise at the Command Line	12-5
Example – Simulating a Continuous-Time State-Space Model at the Command Line	12-6
Predicting Model Output at the Command Line	12-7
Specifying Initial States	12-8
When to Specify Initial States	12-8
Setting Initial States to Zero	12-8
Setting Initial States to Equilibrium Values	12-9
Estimating Initial States from the Data	12-9

Using Models with Control System Toolbox	13-2
How Control System Toolbox Works with Estimated Models	13-2
Using balred to Reduce Model Order	13-3
Compensator Design Using Control System Toolbox	13-3
Converting Models to LTI Objects	13-4
Viewing Model Response Using the LTI Viewer	13-5
Combining Model Objects	13-6
Example – Using System Identification Toolbox and Control System Toolbox	13-6
Using Models with Simulink	13-8
What You Can Do with Simulink	13-8
Opening the System Identification Library	13-9
Working with System Identification Blocks	13-9
Blocks for Importing Data and Models into Simulink	13-9
Blocks for Model Estimation	13-10
Example – Using Simulink to Simulate a Multiexperiment Model	13-11

Working with the System Identification Tool GUI

How to Use the System Identification Tool GUI (p. 1-2)	Typical tasks in the System Identification Tool GUI and where to learn more about each task.
Starting and Managing GUI Sessions (p. 1-3)	How to start and manage sessions using the System Identification Tool GUI.
Managing Data Sets and Models in the GUI (p. 1-9)	How to manage data and models in the System Identification Tool GUI by adding and deleting data and models, specifying estimation and validation data, viewing properties, organizing icons, and exporting to the MATLAB Workspace browser.
Plots Data and Models in the GUI (p. 1-18)	How to plot data and models, identify data and models on a plot, select to plot specific input and output channels, set plot options, and print plots.
Customizing the System Identification Tool GUI (p. 1-26)	How to customize the behavior and appearance of the System Identification Tool GUI.

How to Use the System Identification Tool GUI

To work with the System Identification Tool GUI, you must have already imported your data into the MATLAB Workspace browser. For more information about importing data into MATLAB, see the MATLAB Programming documentation.

A typical workflow in the System Identification Tool GUI consists of the following tasks:

- 1 Start a new session in the System Identification Tool GUI, or open a saved session. For more information, see “Starting and Managing GUI Sessions” on page 1-3.
- 2 Import data into the GUI from the MATLAB Workspace browser. For more information, see “Importing Data into the GUI” on page 2-14.
- 3 Plot the data. For more information, see “Plotting Data” on page 2-62.
- 4 Preprocess the data. For example, you can remove constant offsets or linear trends (for linear models only), filter data, or select regions of interest. For more information, see Chapter 2, “Preparing Data”.
- 5 Specify the data to use for estimation and validation. For more information, see “Specifying Estimation and Validation Data” on page 1-11.
- 6 Estimate models using commands in the **Estimate** menu.
- 7 Validate models. For more information, see Chapter 10, “Validating and Analyzing Models”.
- 8 Export models to the MATLAB Workspace browser for further processing. For more information about exporting models, see “Exporting Data and Models to the MATLAB Workspace Browser” on page 1-16.

Starting and Managing GUI Sessions

In this section...
“What Is a Session?” on page 1-3
“Starting a New Session” on page 1-4
“Overview of the System Identification Tool GUI” on page 1-5
“Opening a Saved Session” on page 1-6
“Saving, Merging, and Closing Sessions” on page 1-6
“Deleting a Session” on page 1-7
“Getting Help in the GUI” on page 1-7
“Exiting the System Identification Tool GUI” on page 1-8

What Is a Session?

A *session* is the data and models that are currently in the System Identification Tool GUI.

You can save a session to a file that has a `.sid` extension. For example, you can save different stages of your progress as different sessions so that you can revert to any stage simply by opening the corresponding session.

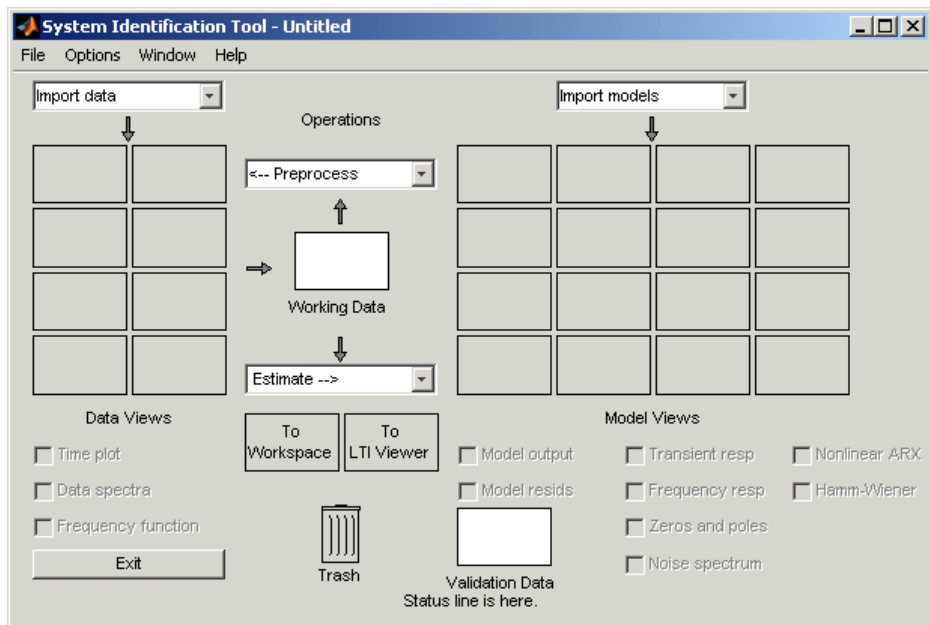
Note Only one session can be open at a time.

Starting a New Session

To start a new session in the System Identification Tool GUI, type the following command at the MATLAB prompt:

```
ident
```

Alternatively, you can start a new session by selecting **Start > Toolboxes > System Identification > System Identification Tool GUI** in the MATLAB Command Window.




Note Only one session can be open at a time.

You can also start a new session by closing the current session using **File > Close session**. This toolbox prompts you to save your current session if it is not yet saved.

Overview of the System Identification Tool GUI

The following figure describes the different areas in the System Identification Tool GUI.



Select check boxes to display data plots.

Select check boxes to display model plots.

The layout of the window organizes tasks and information from left to right. This organization follows a typical workflow, where you start in the top-left corner by importing data into the System Identification Tool GUI using

the **Import data** menu and end in the bottom-right corner by plotting the characteristics of your estimated model on model plots.

The **Data Board** area, located below the **Import data** menu in the System Identification Tool GUI, contains rectangular icons that represent the data you imported into the GUI.

The Model Board, located to the right of the **<--Preprocess** menu in the System Identification Tool GUI, contains rectangular icons that represent the models you estimated or imported into the GUI. You can drag and drop model icons in the Model Board and into open dialog box boxes.

For more information about System Identification Tool GUI workflow, see “How to Use the System Identification Tool GUI” on page 1-2. To learn how to manage data and models in the GUI, see “Managing Data Sets and Models in the GUI” on page 1-9.

Opening a Saved Session

You can open a previously saved session using the following syntax:

```
ident(session,path)
```

In the preceding syntax, `session` is the file name of the session you want to open, and `path` is the location of the session file. Session files have the extension `.sid`. When the session file is on the `matlabpath`, you can omit the `path` argument.

If the System Identification Tool GUI is already open, you can open a session by selecting **File > Open session**.

Note If there is data in the System Identification Tool GUI, you must close the current session before you can open a new session by selecting **File > Close session**.

Saving, Merging, and Closing Sessions

The following table summarizes the menu commands for saving, merging, and closing sessions in the System Identification Tool GUI.

Task	Command	Description
Close the current session and start a new session.	File > Close session	You are prompted to save the current session before closing it.
Merge the current session with a previously saved session.	File > Merge session	You must start a new session and import data or models before you can select to merge it with a previously saved session. You are prompted to select the session file to merge with the current. This operation combines the data and the models of both sessions in the current session.
Save the current session.	File > Save	Useful for saving the session repeatedly after you have already saved the session for the first time.
Save the current session under a new name.	File > Save As	Useful when you want to save your work incrementally. This command lets you revert to a previous stage, if necessary.

Deleting a Session

To delete a saved session, you must delete the corresponding session file.

Getting Help in the GUI

System Identification Tool GUI provides online help topics that you can access from the **Help** menu.

Contextual help is available from each dialog box by clicking the **Help** button in the dialog box.

Exiting the System Identification Tool GUI

To exit System Identification Tool GUI, click **Exit** in the bottom-left corner of the window. Alternatively, select **File > Exit System Identification Tool GUI**.

Managing Data Sets and Models in the GUI

In this section...

“Adding Data to the GUI” on page 1-9

“Adding Models to the GUI” on page 1-10

“Specifying Estimation and Validation Data” on page 1-11

“Importing Models into the System Identification Tool GUI” on page 1-11

“Viewing Data and Model Properties” on page 1-13

“Organizing Data and Model Icons” on page 1-14

“Deleting Data and Models” on page 1-15

“Exporting Data and Models to the MATLAB Workspace Browser” on page 1-16

Adding Data to the GUI

There are three ways to add data to the System Identification Tool GUI:

- Import data into the current session.

For more information about importing data, see “Importing Data into the System Identification Tool GUI” on page 2-14.

- Merge a saved session that contains data into the current session.

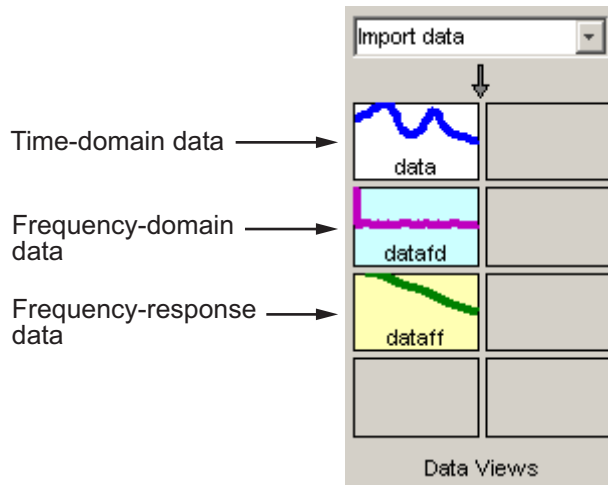
For more information about merging sessions, see “Saving, Merging, and Closing Sessions” on page 1-6.

- Preprocess data.

When you detrend, transform, filter, or select a portion of a data set, you create a new data set in the System Identification Tool GUI. For more information about preprocessing data, see Chapter 2, “Preparing Data”.

Each data set has a unique color to help you distinguish it from the other data sets. In addition, the background color of a data icon is color-coded, as follows:

- White background represents time-domain data.
- Blue background represents frequency-domain data.
- Yellow background represents frequency-response data.



Colors Representing Data Domains

Adding Models to the GUI

There are three ways to add models to the System Identification Tool GUI:

- Import the model into the current session.

For more information about importing models, see “Importing Models into the System Identification Tool GUI” on page 1-11.

- Merge a saved session containing models into the current session.

For more information about merging sessions, see “Saving, Merging, and Closing Sessions” on page 1-6.

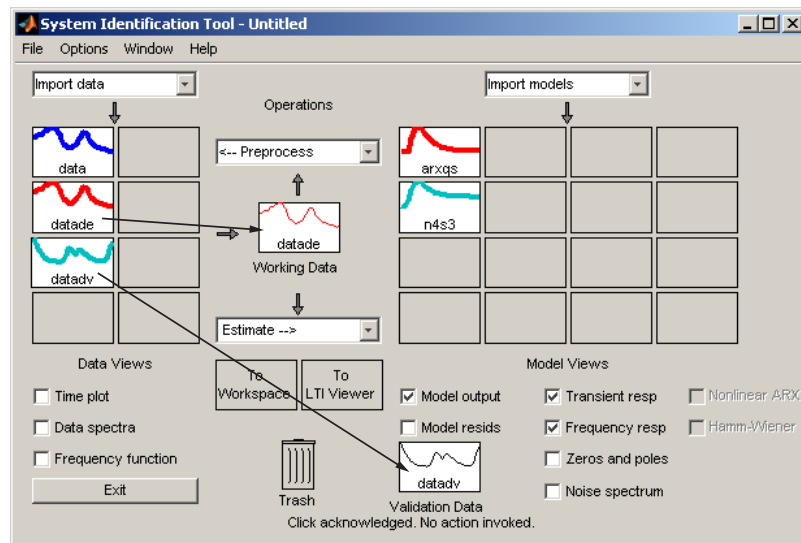
- Estimate models from the data.

Specifying Estimation and Validation Data

You should use independent data sets to estimate and validate your model.

In the System Identification Tool GUI, **Working Data** refers to the estimation data. Similarly, **Validation Data** refers to the data set you use to validate a model. For example, when you plot the model output and residual-analysis plots, the input to the model is the input signal from the validation data set. These plots compare model output to the measured output in the validation data set.

To specify **Working Data**, drag and drop the corresponding data icon into the **Working Data** rectangle, as shown in the following figure.



Similarly, to specify **Validation Data**, drag and drop the corresponding data icon into the **Validation Data** rectangle.

Importing Models into the System Identification Tool GUI

You can import System Identification Toolbox model objects from the MATLAB Workspace browser into the System Identification Tool GUI. If you

have Control System Toolbox, you can also import the LTI objects created using Control System Toolbox.

The following procedure assumes that you begin with the System Identification Tool GUI already open. If this window is not open, type the following command at the MATLAB prompt:

```
ident
```

To import models into the System Identification Tool GUI:

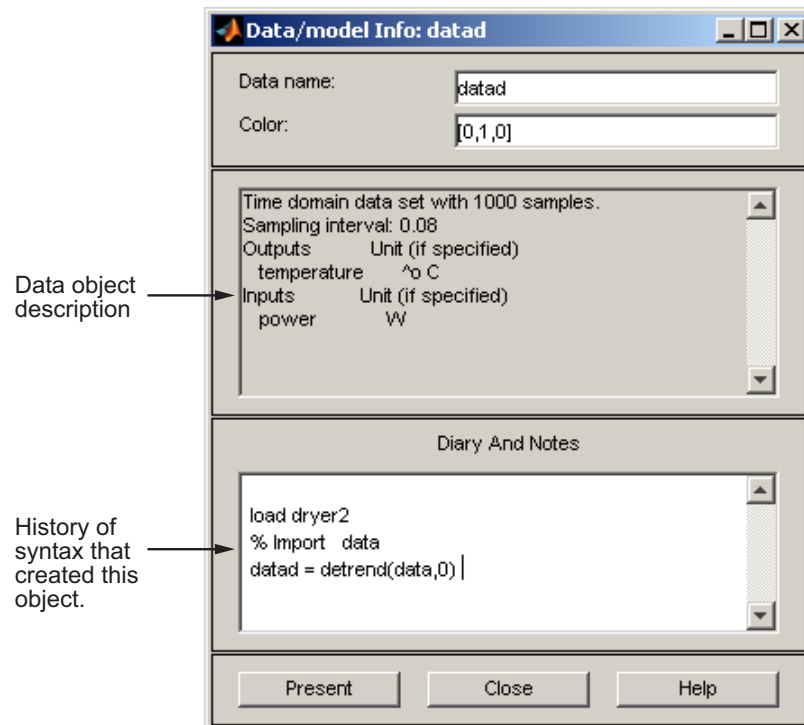
- 1** In the System Identification Tool GUI, select **Import** from the **Import models** list to open the Import Model Object dialog box.
- 2** In the **Enter the name** field, type the name of a model object in the MATLAB Workspace browser. Press **Enter**.
- 3** (Optional) In the **Notes** field, type any notes you want to store with this model.
- 4** Click **Import** to add the model to the Model Board in the System Identification Tool GUI.
- 5** To close the Import Model Object dialog box, click **Close**.

Viewing Data and Model Properties

You can get information about each data set or model in the System Identification Tool GUI by right-clicking the data icon in the System Identification Tool GUI.

The Data/model Info dialog box opens. This dialog box describes the contents and the properties of the corresponding data or model. It also displays any associated notes and the command-line equivalent of the operations you used to create this data or model.

Tip To view or modify properties for several data sets or models, keep this window open and right-click each data set in the System Identification Tool GUI. The Data/model Info dialog box updates as you select each data set.



In the Data/model Info dialog box, you can rename the data and models. To rename data, enter a new name in the **Data name** field. To rename models, enter a new name in the **Model name** field.

You can also specify a new display color using three RGB values in the **Color** field.

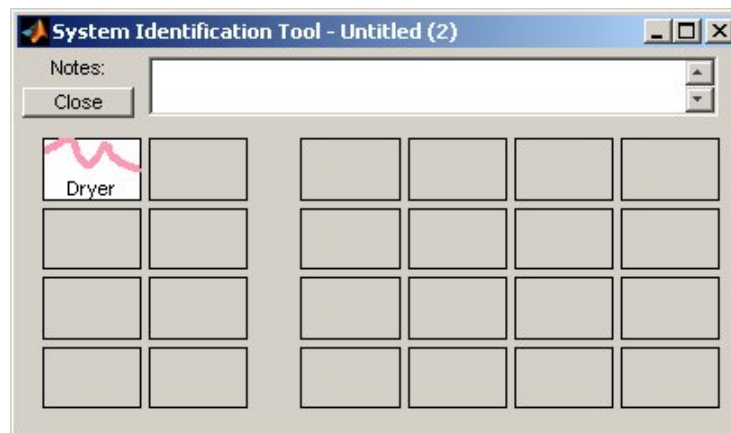
Clicking **Present** displays the basic characteristics of this data object in the MATLAB Command Window.

Organizing Data and Model Icons

You can rearrange data or model icons in the System Identification Tool GUI by dragging and dropping the icons to empty rectangles.

Note You cannot drag and drop a data icon into the Model Board or a model icon into the Data Board.

When you need additional space for organizing data or model icons, select **Options > Extra model/data board** in the System Identification Tool GUI. This action opens an extra session window with blank rectangles for data and models. The new window is an extension of the current session, and does not represent a new session.



Tip When you import or create data sets or estimate models and there is insufficient space for the icons, an additional session window opens automatically.

You can drag and drop data and model icons between the main System Identification Tool GUI and any extra session windows.

Type comments in the **Notes** field to describe the data and models. When you save a session, as described in “Saving, Merging, and Closing Sessions” on page 1-6, all additional windows and notes are also saved.


Deleting Data and Models

To delete data and models in the System Identification Tool GUI, drag and drop the corresponding icon into **Trash**. Moving items to **Trash** does not permanently delete these items.

Note You cannot delete a data set that is currently designated as **Working Data** or **Validation Data**. You must first specify a different data set in the System Identification Tool GUI to be **Working Data** or **Validation Data**, as described in “Specifying Estimation and Validation Data” on page 1-11.

To restore a data set or a model from **Trash**, drag its icon from **Trash** to the Data or Model Board in System Identification Toolbox window. You can view the **Trash** contents by double-clicking the **Trash** icon in the System Identification Tool GUI.

Note You must restore data to the Data Board; you cannot drag data icons to the Model Board. Similarly, you must restore a model to the Model Board in System Identification Toolbox window.



To permanently delete all data sets and models in **Trash**, select **Options > Empty trash**.

Exiting a session empties the **Trash** automatically.

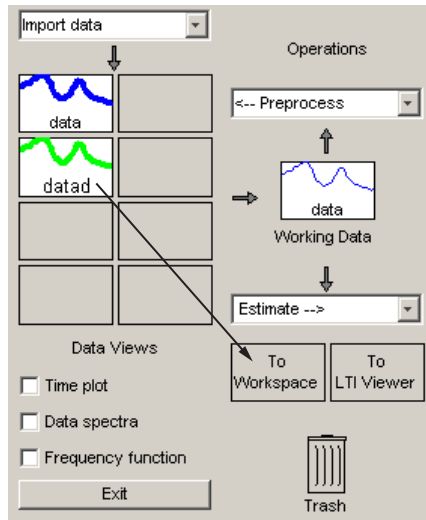
Exporting Data and Models to the MATLAB Workspace Browser

The data and models you create in the System Identification Tool GUI are not available in the MATLAB Workspace browser until you export them. Exporting to MATLAB is necessary when you need to perform an operation on the data or model that is only available using command-line syntax.

Exporting models to the MATLAB Workspace browser also makes them available to Simulink[®] or another toolbox, such as Control System Toolbox.

To export a data set or model to the MATLAB Workspace browser, drag and drop the corresponding icon to the **To Workspace** rectangle.

When you export data and model to the MATLAB Workspace browser, the resulting MATLAB variables have the same name as in the System Identification Tool GUI. For example, the following figure shows how to export the time-domain data object `data` to the MATLAB Workspace browser.



Exporting Data to the MATLAB Workspace Browser

The MATLAB Workspace browser now contains a variable named `data`, which is an `iddata` object.

Plots Data and Models in the GUI

In this section...

“Plotting Data in the GUI” on page 1-18

“Plotting Models in the GUI” on page 1-19

“Identifying Data Sets and Models on Plots” on page 1-20

“Changing and Restoring Default Axis Limits” on page 1-21

“Selecting Measured and Noise Channels in Plots” on page 1-23

“Grid, Line Styles, and Redrawing Plots” on page 1-24

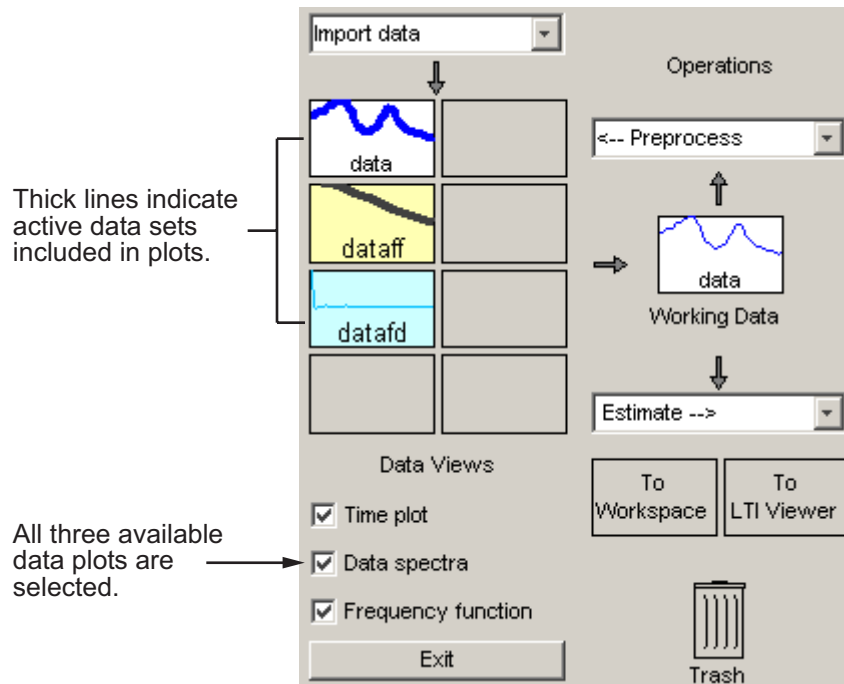
“Opening a Plot in a MATLAB Figure Window” on page 1-25

“Printing Plots” on page 1-25

Plotting Data in the GUI

To create one or more plots of your data, select the corresponding check box in the **Data Views** area of the System Identification Tool GUI. An *active* data icon has a thick line in the icon, while an *inactive* data set has a thin line.

Only active data sets appear on the selected plots. To toggle including and excluding data or models on a plot, click the corresponding icon in the System Identification Tool GUI. Clicking the data icon updates any plots that are currently open.



To close a plot, clear the corresponding check box in the System Identification Tool GUI.

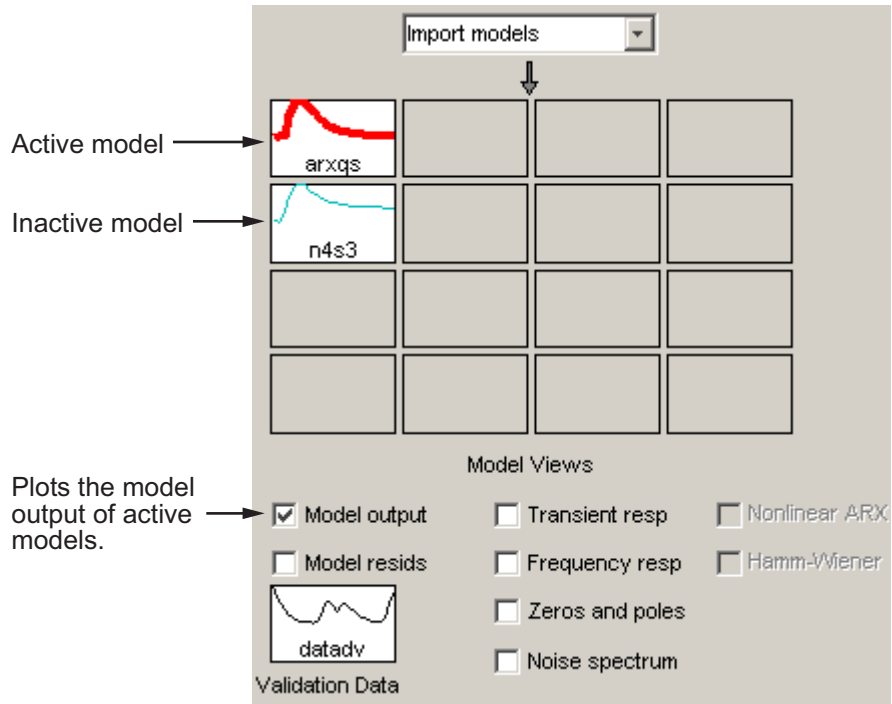
Tip To get information about a specific plot, select a help topic from the **Help** menu in the plot window.

Plotting Models in the GUI

To create one or more plots of your models, select the corresponding check box in the **Model Views** area of the System Identification Tool GUI. An *active* model icon has a thick line in the icon, while an *inactive* model has a thin line. Only active models appear on the selected plots.

To include or exclude a model on a plot, click the corresponding icon in the System Identification Tool GUI. Clicking the model icon updates any plots that are currently open.

For example, in the following figure, **Model output** is selected. In this case, the models n4s4 is not included on the plot because only arx441 is active.



Plots Include Only Active Models

To close a plot, clear the corresponding check box in the System Identification Tool GUI.

Tip To get information about a specific plot, select a help topic from the **Help** menu in the plot window.

Identifying Data Sets and Models on Plots


You can identify data sets and models on a plot by color: the color of the line in the data or model icon in the System Identification Tool GUI matches the line color on the plots.

You can also display data tips for each line on the plot. How you display the data tip depends on whether the zoom feature is enabled:

- When zoom is enabled, press and hold down **Shift**, and click the desired curve to display a data tip on a plot.
- When zoom is disabled, click a plot curve and hold down the mouse button to display the data tip.

For more information about enabling zoom, see “Magnifying Plots” on page 1-22.

The following figure shows an example of a data tip, which contains the name of the data set and the coordinates of the location on the curve you clicked.



Data Tip on a Plot

Changing and Restoring Default Axis Limits

There are two ways to change which portion of the plot is currently in view:

- “Magnifying Plots” on page 1-22
- “Setting Axis Limits” on page 1-22

Magnifying Plots

Enable zoom by selecting **Style > Zoom** in the plot window. To disable zoom, select **Style > Zoom** again.

Tip To verify that zoom is active, click the **Style** menu. A check mark should appear next to **Zoom**.

You can adjust magnification in the following ways:

- To zoom in default increments, left-click the portion of the plot you want to center in the plot window.
- To zoom in on a specific region, click and drag a rectangle that identifies the region for magnification. When you release the mouse button, the selected region is displayed.
- To zoom out, right-click on the plot.

Note To restore the full range of the data in view, select **Options > Autorange** in the plot window.

Setting Axis Limits

You can change axis limits for the vertical and the horizontal axes of the input and output channels that are currently displayed on the plot.

- 1 Select **Options > Set axes limits** to open the Limits dialog box.
- 2 Specify a new range for each axis by editing its lower and upper limits. The limits must be entered using the format *[LowerLimit UpperLimit]*. Click **Apply**. For example:

[0.1 100]

Note To restore full axis limits, select the **Auto** check box to the right of the axis name, and click **Apply**.

- 3** To plot data on a linear scale, clear the **Log** check box to the right of the axis name, and click **Apply**.

Note To revert to base-10 logarithmic scale, select the **Log** check box to the right of the axis name, and click **Apply**.

- 4** Click **Close**.

Note To restore the full range of the data in view, select **Options > Autorange** in the plot window.

Selecting Measured and Noise Channels in Plots

Input and output variables are called *channels*. When you create a plot of a multivariate input or output data set or model, the plot only shows one input/output channel pair at a time. The selected channel names are displayed in the title bar of the plot window.

Note When you select to plot multiple data sets, and each data set contains several input and output channels, the **Channel** menu lists channel pairs from all data sets.

You can select a different input/output channel pair from the **Channel** menu in the plot window.

The **Channel** menu uses the following notation to list channels: $u1 \rightarrow y2$ means that the transfer command from input channel $u1$ to output channel $y2$ is displayed.

The symbol e represents a noise source and might appear as $e@y3 \rightarrow y1$, which means that the transfer command from the noise channel (associated with $y3$) to output channel $y2$ is displayed. System identification estimates as many noise sources as there are output channels. In general, $e@ynam$ indicates that the noise source corresponds to the output with name $ynam$. For more information about noise channels, see “Subreferencing Measured and Noise Models” on page 3-21.

Tip When you import data into the System Identification Tool GUI, it is helpful to assign meaningful channel names in the Import Data dialog box. For more information about importing data, see “Importing Data into the GUI” on page 2-14.

Grid, Line Styles, and Redrawing Plots

Although different plot types contain different options in the **Style** menu that are specific to each plot type, there are several **Style** options that are common to all plot types. These include the following:

- “Grid Lines” on page 1-24
- “Solid or Dashed Lines” on page 1-24
- “Plot Redrawing” on page 1-25

Grid Lines

To toggle showing or hiding grid lines, select **Style > Grid**.

Solid or Dashed Lines

To display currently visible lines as a combination of solid, dashed, dotted, and dash-dotted line style, select **Style > Separate linestyles**.

To display all solid lines, select **Style > All solid lines**. This choice is the default.

All line styles match the color of the corresponding data or model icon in the System Identification Tool GUI.

Plot Redrawing

To specify that the plot be redrawn when you add another data set or model to the plot, select **Style > Erase mode normal**. This choice is the default setting.

To avoid redrawing the entire plot when you add another data set or model to the plot, select **Style > Erase mode xor**. Although this selection results in faster response, it might also produce poor plot quality.

Opening a Plot in a MATLAB Figure Window

The MATLAB Figure window provides editing and printing controls for plots that are not available in the System Identification Toolbox plot window. Sometimes it is useful to create a plot in the System Identification Tool GUI, and then open it in a MATLAB Figure window to fine-tune its appearance.

After you create the plot, as described in “Plotting Models in the GUI” on page 1-19, select **File > Copy figure** in the plot window. This command opens the plot in a MATLAB Figure window.

Printing Plots

To print a System Identification Toolbox plot, select **File > Print** in the plot window. In the Print dialog box, select the printing options and click **OK**.

Customizing the System Identification Tool GUI

In this section...

“Types of GUI Customization” on page 1-26

“Displaying Warnings While You Work” on page 1-26

“Saving Session Preferences” on page 1-26

“Modifying idlayout.m” on page 1-27

Types of GUI Customization

The System Identification Tool GUI lets you customize the window behavior and appearance and save the session.

For example, you can set the size and position of specific dialog box boxes and modify the appearance of plots. Next, you can save the session to save the GUI state.

Advanced users might choose to edit the M-file that controls default settings, as described in “Modifying idlayout.m” on page 1-27.

Displaying Warnings While You Work

In the System Identification Tool GUI, select **Options > Warnings** to display informational dialog box boxes while you work. Verify that a check mark appears to the right of **Warnings** in the menu.

To stop warnings from being displayed during your session, select **Options > Warnings** and clear the check mark.

Saving Session Preferences

Use **Options > Save preferences** to save the current state of the System Identification Tool GUI. This command saves the following settings to System Identification Toolbox preferences file, `idprefs.mat`:

- Size and position of the System Identification Tool GUI
- Sizes and positions of dialog box boxes

- Four recently used sessions
- Plot options, such as line styles, zoom, grid, and information about whether the input is plotted using zero-order hold or first-order hold

Because `idprefs.mat` is a data file, you cannot edit it directly. The `idprefs.mat` file is located in the same directory as `startup.m`, by default. To change the location where your preferences are saved, use the `midprefs` command with the new path as the argument.

For example:

```
midprefs('c:\matlab\toolbox\local\')
```

You can also type `midprefs`, and browse to the desired directory.

To restore the default preferences, select **Options > Default preferences**.

Modifying `idlayout.m`

Advanced users might want to customize the default plot options by editing `idlayout.m`. By editing `idlayout.m`, you can customize more layout properties than the ones you can set by saving preferences to `idprefs.mat`.

To customize `idlayout.m` defaults, save a copy of `idlayout.m` to a folder in your `matlabpath` just above the `ident` directory level.

Caution Do not edit the original file to avoid overwriting the `idlayout.m` defaults shipped with the product.

You can customize the following plot options in `idlayout.m`:

- The order in which colors are assigned to data and model icons. You can later edit the color for a specific data or model in the Data/model Info dialog box. To open the Data/model Info dialog box, right-click the corresponding data or model icon in the System Identification Tool GUI.
- Line colors on plots.
- Axis limits and tick marks.

- Plot options, set in the plot menus.
- Font size.

Note When you save preferences using **Options > Save preferences** to `idprefs.mat`, these preferences override the defaults in `idlayout.m`. To give `idlayout.m` precedence every time you start a new session, select **Options > Default preferences**.

Preparing Data

Ways to Prepare Data for Modeling
(p. 2-3)

Summarizes how to prepare your data for system identification.

Data Requirements (p. 2-6)

MATLAB variables that must result from importing your time-domain, frequency-domain, and frequency-response data into MATLAB so that you can represent this data in System Identification Toolbox.

Importing Data into MATLAB
(p. 2-13)

Resources for learning more about importing data into MATLAB.

Importing Data into the GUI (p. 2-14)

Importing data from the MATLAB Workspace browser into the System Identification Tool GUI, creating new data sets by segmenting and combining data sets, renaming data sets, and changing the color of data icons.

Creating iddata Objects (p. 2-32)

Using the iddata constructor to represent time-domain and frequency-domain data and working with iddata objects.

Creating idfrd Objects (p. 2-52)

Using the idfrd constructor to represent frequency-response data and working with idfrd objects.

Getting Advice About Your Data (p. 2-60)	Using the advice command to identify constant offsets and linear trends, delays, feedback, and signal excitation levels in the data.
Plotting Data (p. 2-62)	Plotting time-domain, frequency-domain, and frequency-response data.
Handling Missing Data and Outliers (p. 2-71)	Handling missing or erroneous data values.
Detrending Data (p. 2-75)	Removing constant offsets and linear trends.
Resampling Data (p. 2-80)	Decimating and interpolating data.
Filtering Data (p. 2-87)	Deciding whether to filter data before model estimation, and how to prefilter data.
Selecting Data (p. 2-95)	Selecting portions of your data for identification.
Creating Data Using Simulation (p. 2-99)	Creating input data with specific characteristics and simulating the output data using a model.
Transforming Between Time- and Frequency-Domain Data (p. 2-103)	Transforming between time-domain, frequency domain, and frequency-response data.
Handling Complex-Valued Data (p. 2-115)	Summarizes supported operations and limitations for handling complex data, and lists commands for manipulating complex iddata signals.

Ways to Prepare Data for Modeling

The following tasks help to prepare your data for system identification or modeling:

Import data into MATLAB

Before you can estimate models in System Identification Toolbox, you must first import your data into the MATLAB Workspace browser. You can import the data from external data files, create data by simulation, or manually create data arrays at the command line.

For more information about requirements on MATLAB variables, see “Data Requirements” on page 2-6.

Represent data in the GUI or at the command line

You can representing the data variables in the MATLAB Workspace browser for system identification by doing one of the following:

- For working in the GUI, import data into the System Identification Tool GUI. See “Importing Data into the GUI” on page 2-14.
- For working at the command line, create an `iddata` or `idfrd` object.

For time-domain or frequency-domain data, see “Creating `iddata` Objects” on page 2-32.

For frequency-response data, see “Creating `idfrd` Objects” on page 2-52.

Simulate data

As an alternative to using measured data, you can simulate your data instead. You can choose to simulate data with and without noise.

To learn how to create data sets using simulation, see “Creating Data Using Simulation” on page 2-99.

Plot and analyze data

You can analyze your data by doing any of the following:

- Plotting data to examine both time- and frequency-domain behavior. See “Plotting Data” on page 2-62.
- Using the `advise` command to analyze the data for the presence of constant offsets and trends, delay, feedback, and signal excitation levels. See “Getting Advice About Your Data” on page 2-60.

Preprocess data

Review the data characteristics for any of the following features to determine if there is a need for preprocessing:

- Missing or faulty values (also known as *outliers*). For example, you might see gaps that indicate missing data, values that do not fit with the rest of the data, or noninformative values. See “Handling Missing Data and Outliers” on page 2-71.
- Offsets and drifts in signal levels (low-frequency disturbances).
See “Detrending Data” on page 2-75 for information about subtracting means and linear trends, and “Filtering Data” on page 2-87 for information about filtering.
- High-frequency disturbances above the frequency interval of interest for the system dynamics.
See “Resampling Data” on page 2-80 for information about decimating and interpolating values, and “Filtering Data” on page 2-87 for information about filtering.

Selecting a subset of your data

You can use data selection as a way to clean the data and exclude parts with noisy or missing information. You can also use data selection to create independent data sets for estimation and validation.

To learn more about selecting data, see “Selecting Data” on page 2-95.

Combining data from multiple experiments

You can combine data from multiple experiments performed under the same condition into a single data set. The model you estimate from a multiple-experiment data set is an average model.

To learn more about creating multi-experiment data sets, see “Creating Multiexperiment Data Sets in the GUI” on page 2-23 or “Creating Multiexperiment Data at the Command Line” on page 2-38.

Data Requirements

In this section...

“Types of Data You Can Model” on page 2-6

“Data with Uniform and Nonuniform Sampling Rates” on page 2-7

“Representing Time-Domain Data” on page 2-7

“Representing Time-Series Data” on page 2-8

“Representing Frequency-Domain Data” on page 2-9

“Representing Frequency-Response Data” on page 2-10

Types of Data You Can Model

For linear models, System Identification Toolbox supports both time- and frequency-domain data with single or multiple inputs and outputs. Time-domain data can be either real or complex.

For nonlinear models, this toolbox supports only time-domain data.

Time-domain data is one or more input variables $u(t)$ and one or more output variables $y(t)$, sampled as a command of time. A special case of time-domain data is time-series data, which is one or more outputs $y(t)$ and no measured input. *Frequency-domain data* is the Fourier transform of the input and output time-domain signals. *Frequency-response data*, also called *frequency-command data*, represents complex frequency-response values for a linear system characterized by its transfer command G .

You can measure frequency-response data values directly using a spectrum analyzer, for example. In this section, *frequency-domain* and *frequency-response* are both called frequency-domain data for the sake of brevity.

For time-series data, you can estimate both linear and nonlinear models.

Note If your data is complex valued, see “Handling Complex-Valued Data” on page 2-115 for information about supported operations in System Identification Toolbox.

Data with Uniform and Nonuniform Sampling Rates

A *sampling interval* is the time between successive data samples.

System Identification Toolbox interface provides limited support for nonuniformly sampled data. For more information about specifying uniform and nonuniform time vectors, see “Constructing `iddata` for Time-Domain Data” on page 2-33.

Note The System Identification Tool GUI only supports uniformly sampled data.

Representing Time-Domain Data

Time-domain data consists of one or more input variables $u(t)$ and one or more output variables $y(t)$, sampled as a function of time. If there is no output data, see “Representing Time-Series Data” on page 2-8.

The following variables must exist in the MATLAB Workspace browser before you can represent time-domain data in System Identification Toolbox:

- Input data

For single-input/single-output (SISO) data, the input must be a column-wise vector.

For a data set with N_u inputs and N_T samples (measurements), the input is an N_T -by- N_u matrix.

- Output data

For single-input/single-output (SISO) data, the output must be a column-wise vector.

For a data set with N_y outputs and N_T samples (measurements), the output is an N_T -by- N_y matrix.

- Sampling time interval

If you are working with uniformly sampled data, use the actual sampling interval in your experiment. Each data value is assigned a sample time, which is calculated from the start time and the sampling interval. If you are working with nonuniformly sampled data at the command line, you can specify a vector of time instants using the `iddata` `TimeInstants` property, as described in “Constructing `iddata` for Time-Domain Data” on page 2-33.

Representing Time-Series Data

A special case of time-domain data is time-series data, which consist of one or more outputs $y(t)$ with no corresponding input.

The following variables must exist in the MATLAB Workspace browser before you can represent time-series data in System Identification Toolbox:

- Output data

- For single-input/single-output (SISO) data, the output must be a column-wise vector.
- For a data set with N_y outputs and N_T samples (measurements), the output is an N_T -by- N_y matrix.

- Sampling time interval

- If you are working with uniformly sampled data, use the actual sampling interval in your experiment. Each data value is assigned a sample time, which is calculated from the start time and the sampling interval. If you are working with nonuniformly sampled data at the command line, you can specify a vector of time instants using the `iddata` `TimeInstants` property, as described in “Constructing `iddata` for Time-Domain Data” on page 2-33.

For information about estimating time-series models, see Chapter 6, “Estimating Time-Series Models”.

Representing Frequency-Domain Data

- “About Frequency-Domain Data” on page 2-9
- “Representing Frequency-Domain Data” on page 2-10

About Frequency-Domain Data

Frequency-domain data is the Fourier transform of the input and output time-domain signals. For continuous-time signals, the Fourier transform over the entire time axis is defined as follows:

$$Y(i\omega) = \int_{-\infty}^{\infty} y(t)e^{-i\omega t} dt$$

$$U(i\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} u(t)e^{-i\omega t} dt$$

In the context of numerical computations, continuous equations must be replaced by their discretized equivalents to handle discrete data values. For a discrete-time system with a sampling interval T , the frequency-domain output $Y(e^{i\omega})$ and input $U(e^{i\omega})$ is the time-discrete Fourier transform (TDFT):

$$Y(e^{i\omega T}) = T \sum_{k=1}^N y(kT)e^{-i\omega kT}$$

In this example, $k = 1, 2, \dots, N$, where N is the number of samples in the sequence.

Note This form only discretizes the time. The frequency is continuous.

When the frequencies are not equally spaced, it is useful to also discretize the frequencies in the Fourier transform. The resulting discrete Fourier transform (DFT) of time-domain data is:

$$Y(e^{iw_n T}) = \sum_{k=1}^N y(kT)e^{-iw_n kT}$$
$$w_n = \frac{2\pi n}{T} \quad n = 0, 1, 2, \dots, N-1$$

The DFT is useful because it can be calculated very efficiently using the fast Fourier transform (FFT) method. Fourier transforms of the input and output data are complex values.

Representing Frequency-Domain Data

The following variables must exist in the MATLAB Workspace browser before you can represent frequency-domain data in System Identification Toolbox:

- Input data
 - For single-input/single-output (SISO) data, the input must be a column-wise vector.
 - For a data set with N_u inputs and N_f frequencies, the input is an N_f -by- N_u matrix.
- Output data
 - For single-input/single-output (SISO) data, the output must be a column-wise vector.
 - For a data set with N_y outputs and N_f frequencies, the output is an N_f -by- N_y matrix.
- Frequency values
 - Must be a column-wise vector.

Representing Frequency-Response Data

- “About Frequency-Response Data” on page 2-11

- “Representing Frequency-Response Data” on page 2-12

About Frequency-Response Data

Frequency-response data, also called *frequency-command data*, consists of complex frequency-response values for a linear system characterized by its transfer command G . You can measure frequency-response data values directly using a spectrum analyzer, for example, which provides a compact representation of the input and the output (compared to storing input and output independently).

The transfer command G is essentially an operator that takes the input u of a linear system to the output y :

$$y = Gu$$

For a continuous-time system, the transfer command relates the Laplace transforms of the input $U(s)$ and output $Y(s)$:

$$Y(s) = G(s)U(s)$$

In this case, the frequency command $G(i\omega)$ is the transfer command evaluated on the imaginary axis $s=i\omega$.

For a discrete-time system sampled with a time interval T , the transfer command relates the Z-transforms of the input $U(z)$ and output $Y(z)$:

$$Y(z) = G(z)U(z)$$

In this case, the frequency command $G(e^{i\omega T})$ is the transfer command $G(z)$ evaluated on the unit circle. The argument of the frequency command $G(e^{i\omega T})$ is scaled by the sampling interval T to make the frequency command periodic with the sampling frequency $2\pi/T$.

For a sinusoidal input to the system, the output is also a sinusoid with the same frequency. The frequency-response data magnifies the amplitude of the input by $|G|$ and shifts its phase by $\phi = \arg G$. Because the frequency command is evaluated at the sinusoid frequency, the values of the frequency

command at a specific frequency describe the response of the linear system to an input at that frequency.

Frequency-response data represents a (nonparametric) model of the relationship between the input and the outputs as a command of frequency. You might use such a model, which consists of a table of values, to study the system frequency response. However, you cannot use this model for simulation and prediction and must create a parametric model from the frequency-response data.

Representing Frequency-Response Data

There are two ways to represent frequency-response data in System Identification Toolbox. The first approach lets you manipulate the data using both System Identification Tool GUI and the command line, and the second approach is only used for working with data in the System Identification Tool GUI.

The following variables must exist in the MATLAB Workspace browser before you can represent frequency-response data in System Identification Toolbox:

- In System Identification Tool GUI or MATLAB Command Window, represent complex-valued $G(e^{iw})$.

For single-input single-output (SISO) systems, the frequency command is a column-wise vector.

For a data set with N_u inputs, N_y outputs, and N_f frequencies, the frequency command is an N_y -by- N_u -by- N_f array.

- In System Identification Tool GUI only, represent amplitude $|G|$ and phase shift $\varphi = \arg G$.

For single-input single-output (SISO) systems, the amplitude and the phase must each be a column-wise vector.

For a data set with N_u inputs, N_y outputs, and N_f frequencies, the amplitude and the phase must each be an N_y -by- N_u -by- N_f array.

- Frequency values must be a column-wise vector.

Importing Data into MATLAB

If you are modeling data from an external data file, you must first import this data into the MATLAB Workspace browser.

For more information about the types of variables required in the MATLAB Workspace browser before you can represent your data for identification, see the following topics:

- “Representing Time-Domain Data” on page 2-7
- “Representing Time-Series Data” on page 2-8
- “Representing Frequency-Domain Data” on page 2-9
- “Representing Frequency-Response Data” on page 2-10

The simplest way to import data into MATLAB is to use the MATLAB Import Wizard, which supports the following data formats:

- MAT-files containing MATLAB arrays or System Identification Toolbox `iddata` and `idfrd` data structures
- Text files, such as `.txt` and `.dat`
- Spreadsheet files, such as `.xls`
- Graphics files, such as `.gif` and `.jpg`
- Audio and video files, such as `.avi` and `.wav`

For more information about importing data into MATLAB, see the MATLAB Programming documentation.

After you import data into MATLAB, you must represent it for system identification.

Importing Data into the GUI

In this section...
“Different Ways to Create New Data Sets in the GUI” on page 2-14
“Importing Data into the System Identification Tool GUI” on page 2-14
“Specifying the Data Sampling Interval” on page 2-19
“Preprocessing Data Using Quick Start” on page 2-20
“Creating Data Sets from Channel Subsets” on page 2-20
“Creating Multiexperiment Data Sets in the GUI” on page 2-23
“Renaming Data and Changing Display Color” on page 2-30

Different Ways to Create New Data Sets in the GUI

There are several ways to create new data sets in the System Identification Tool GUI. For example, you can import data into the System Identification Tool GUI and combine selected subsets of input and output channels to create a new data set. You can also merge several data sets into a single multiexperiment data set, or extract selected experiments from a multiexperiment data set into a single data set.

Note Data preprocessing operations also create new data sets.

For information about working with the System Identification Tool GUI, see Chapter 1, “Working with the System Identification Tool GUI”.

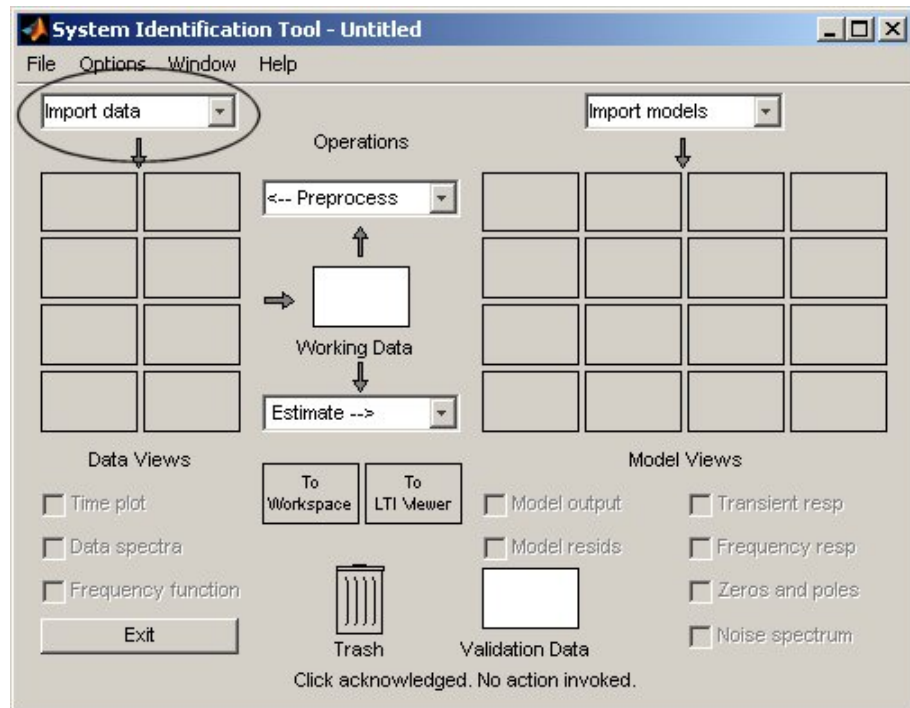
Importing Data into the System Identification Tool GUI

To open the GUI, type the following command at the MATLAB prompt:

```
ident
```

In the **Import data** list, select the type of data to import from the MATLAB Workspace browser, as shown in the following figure.

Note Your data must be sampled at equal time intervals.



The following table summarizes the commands for importing different types of data into the System Identification Tool GUI. Each command opens the Import Data dialog box. For detailed information about the fields in the Import Data dialog box, click **Help**.

This table also specifies which MATLAB variables that must be present in the MATLAB Workspace browser. For information about the dimensions of the required MATLAB variables, see “Data Requirements” on page 2-6.

For an example of importing data into the System Identification Tool GUI, see the Getting Started documentation.

Commands for Importing Data

Type of Data	Import Command	Required Information
<p>Time-domain data sampled as a command of time.</p>	<p>Select Import data > Time domain data to open the Import Data dialog box.</p> <p>Time-Domain Signals is already selected in the Data Format for Signals list.</p>	<p>Specify the following to store data as an iddata object:</p> <ul style="list-style-type: none"> • Input signal as MATLAB vector, matrix, or expression. • Output signal as MATLAB vector, matrix, or expression. • Sampling time interval. See “Specifying the Data Sampling Interval” on page 2-19. <hr/> <p>Note For time series, only import the output signal and enter [] for the input.</p> <hr/>

Commands for Importing Data (Continued)

Type of Data	Import Command	Required Information
<p>Frequency-domain data consists of Fourier transforms of time-domain data (a command of frequency).</p>	<p>Select Import data > Freq. domain data to open the Import Data dialog box.</p> <p>Frequency Domain Signals is already selected in the Data Format for Signals list.</p>	<p>Specify the following to store as an iddata object:</p> <ul style="list-style-type: none"> • Input signal as MATLAB vector, matrix, or expression. • Output signal as MATLAB vector, matrix, or expression. • Frequency vector or MATLAB expression. • Sampling time interval. See “Specifying the Data Sampling Interval” on page 2-19.

Commands for Importing Data (Continued)

Type of Data	Import Command	Required Information
<p>Frequency-response data consists of complex-valued frequency responses at specified frequencies.</p>	<p>Select Import data > Freq. domain data to open the Import Data dialog box.</p> <p>In the Data Format for Signals list, select Freq. Function (complex).</p>	<p>Specify the following to store data as an idfrd object:</p> <ul style="list-style-type: none"> • MATLAB variable or expression representing the complex frequency-response data $G(e^{i\omega})$. • Frequency MATLAB variable or expression representing a vector of frequency values.
<p>Frequency-response data consists of frequency-response amplitude and phase values at specified frequencies.</p>	<p>Select Import data > Freq. domain data to open the Import Data dialog box.</p> <p>In the Data Format for Signals list, select Freq. Function (Amp/Phase).</p>	<p>Specify the following to store data as an idfrd object:</p> <ul style="list-style-type: none"> • Amplitude variable or MATLAB expression representing G. • Phase variable or MATLAB expression representing $\varphi = \arg G$. • Frequency variable or MATLAB expression representing a vector of frequency values.

Commands for Importing Data (Continued)

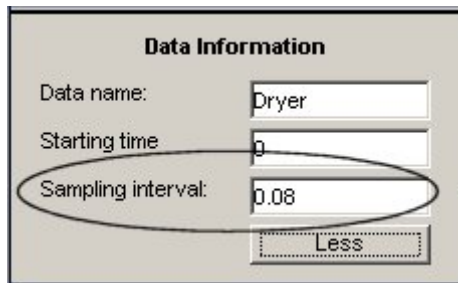
Type of Data	Import Command	Required Information
System Identification Toolbox data object (iddata, idfrd) or Control System Toolbox frd object.	Select Import data > Data object to open the Import Data dialog box. IDDATA or IDFRD/FRD is already selected in the Data Format for Signals list.	Import iddata, idfrd, or frd data object in the MATLAB Workspace browser. Note Importing an frd object from Control System Toolbox converts it to an idfrd object.
Example data from dryer2.mat, consisting of 1000 single-input and single output (SISO) power and temperature values, respectively.	Select Import data > Example to open the Import Data dialog box with the sample data information already entered in the dialog box.	Import sample data and store as an iddata object.

Specifying the Data Sampling Interval

When you specify the input and output data in the Import Data dialog box, you must also enter the sampling interval in the units of your experiment. The *sampling interval* is the time between successive data samples in your experiment and must be the actual time interval at which your data is sampled in any units. For example, enter 0.5 if your data was sampled every 0.5 ms, and enter 1 if your data was sampled every 1 s.

You can also use the sampling interval as a flag to specify continuous-time data. When importing continuous-time frequency domain or frequency-response data, set the **Sampling interval** to 0.

The sampling interval is used during model estimation. For time-domain data, the sampling interval is used together with the start time to calculate the sampling time instants. When you transform time-domain signals to frequency-domain signals (see the `fft` reference page), the Fourier transforms are computed as discrete Fourier transforms (DFTs) for this sampling interval. In addition, the sampling instants are used to set the horizontal axis on time plots.



Sampling Interval in the Import Data dialog box

Preprocessing Data Using Quick Start

As a preprocessing shortcut, select **Preprocess > Quick start** to simultaneously perform the following four actions:

- Subtract the mean value from each channel.
- Split data into two parts.
- Specify the first half as estimation data for models (or **Working Data**).
- Specify the second half as **Validation Data**

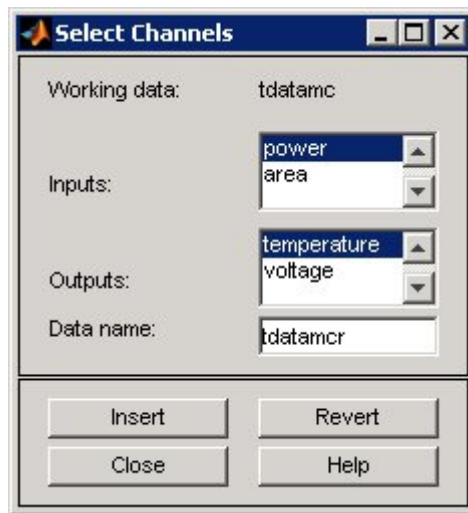
Note Perform these preprocessing operations only when you want to estimate linear models.

Creating Data Sets from Channel Subsets

You can create a new data set in the System Identification Tool GUI by extracting subsets of input and output channels from an existing data set.

To create a new data set from selected channels:

- 1 In the System Identification Tool GUI, drag the icon of the data from which you want to select channels to the **Working Data** rectangle.
- 2 Select **Preprocess > Select channels** to open the Select Channels dialog box.



The **Inputs** list displays the input channels and the **Outputs** list displays the output channels in the selected data set.

3 In the **Inputs** list, select one or more channels in any of following ways:

- Select one channel by clicking its name.
- Select adjacent channels by pressing the **Shift** key while clicking the first and last channel names.
- Select nonadjacent channels by pressing the **Ctrl** key while clicking each channel name.

Tip To exclude input channels and create time-series data, clear all selections by holding down the **Ctrl** key and clicking each selection. To reset selections, click **Revert**.

4 In the **Outputs** list, select one or more channels in any of following ways:

- Select one channel by clicking its name.
- Select adjacent channels by pressing the **Shift** key while clicking the first and last channel names.
- Select nonadjacent channels by pressing the **Ctrl** key while clicking each channel name.

Tip To reset selections, click **Revert**.

5 In the **Data name** field, type the name of the new data set. Use a name that is unique in the Data Board.

6 Click **Insert** to add the new data set to the Data Board in the System Identification Tool GUI.

7 Click **Close**.

Creating Multiexperiment Data Sets in the GUI

- “Why Create Multiexperiment Data?” on page 2-23
- “Limitations on Data Sets” on page 2-23
- “Merging Data Sets” on page 2-23
- “Extracting Experiments into a New Data Set” on page 2-27

Why Create Multiexperiment Data?

You can create a time-domain or frequency-domain data set in the System Identification Tool GUI that includes several experiments. Identifying models for multiexperiment data results in an *average* model.

Experiments can mean data that was collected during different sessions, or portions of the data collected during a single session. In the latter situation, you can create multiexperiment data by splitting a single data set into multiple segments that exclude corrupt data, and then merge the data segments.

Limitations on Data Sets

You can only merge data sets that have *all* of the following characteristics:

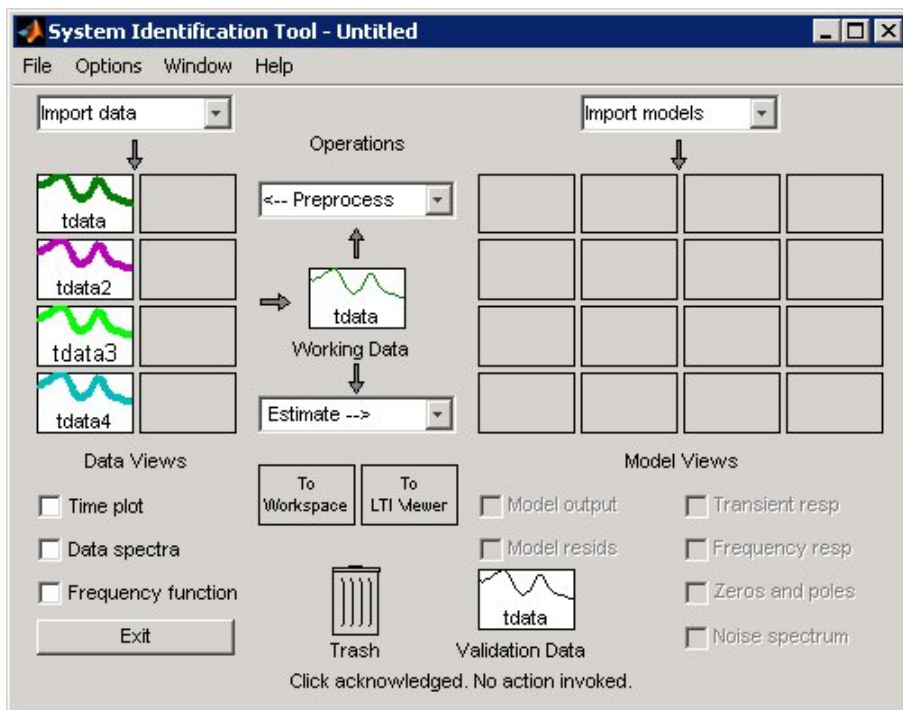
- Same number of input and output channels.
- Different names. The name of each data set becomes the experiment name in the merged data set.
- Same input and output channel names.
- Same type (that is, time-domain data or frequency-domain data only).

Merging Data Sets

You can merge data sets using the System Identification Tool GUI.

Note Before merging several segments of the same data set, verify that the time vector of each data starts at the time when that data segment was actually measured (relative to the other data sets).

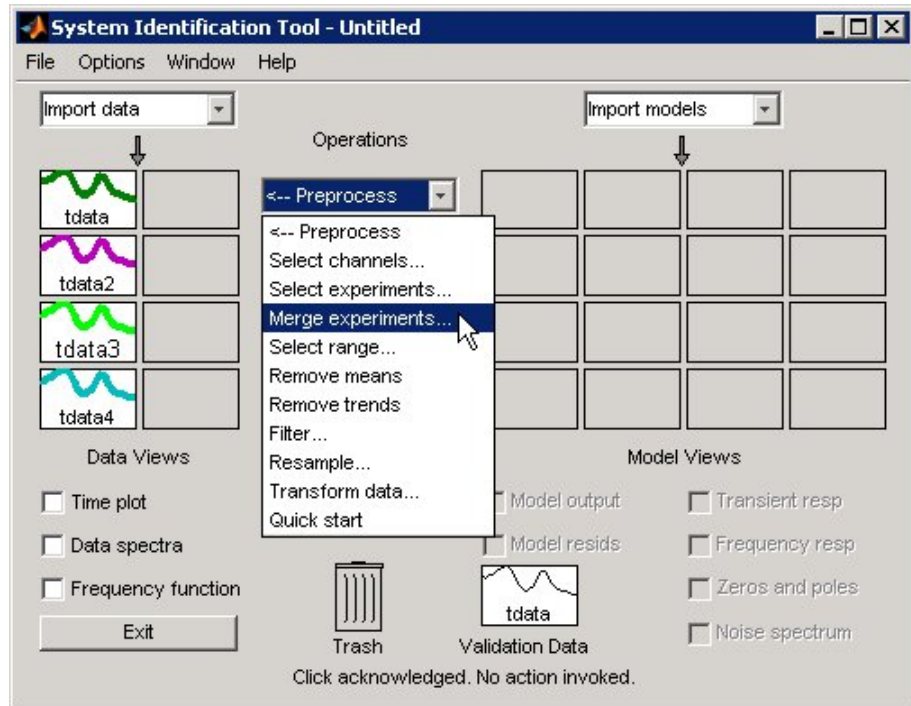
For example, suppose that you want to combine the data sets tdata, tdata2, tdata3, tdata4 shown in the following figure.



GUI Contains Four Data Sets to Merge

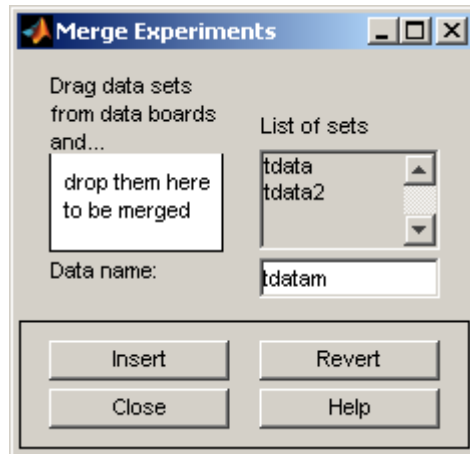
To merge data sets in the GUI:

- 1 In the **Operations** area, select **Preprocess > Merge experiments** from the drop-down menu to open the Merge Experiments dialog box.



- 2 Drag a data set from the Data Board in System Identification Toolbox window to the Merge Experiments dialog box, into the **drop them here to be merged** rectangle.

The name of the data set is added to the **List of sets**.




tdata and tdata2 to Be Merged

Tip To empty the list, click **Revert**.

- 3 Repeat step 2 for each data set you want to merge. Go to the next step after adding data sets.
- 4 In the **Data name** field, type the name of the new data set. This name must be unique in the Data Board.

- 5 Click **Insert** to add the new data set to the Data Board in the System Identification Tool GUI.



Data Board Now Contains tdatam with Merged Experiments

- 6 Click **Close** to close the Merge Experiments dialog box.

Tip To get information about the state of creation of any data set in the System Identification Tool GUI, right-click the data icon to open the Data/model Info dialog box.


Extracting Experiments into a New Data Set

When a data set already consists of several experiments, you can extract one or more of these experiments into a new data set, using the System Identification Tool GUI.

For example, suppose that `tdatam` consists of four experiments.


To create a new data set that includes only the first and third experiments in this data set:

- 1 In the System Identification Tool GUI, drag and drop the `tdata` data icon to the **Working Data** rectangle.



`tdata` Is Set to Working Data

- 2 In the **Operations** area, select **Preprocess** > **Select experiments** from the drop-down menu to open the Select Experiment dialog box.



- 3** In the **Experiments** list, select one or more data sets in either of the following ways:
 - Select one data set by clicking its name.
 - Select adjacent data sets by pressing the **Shift** key while clicking the first and last names.
 - Select nonadjacent data sets by pressing the **Ctrl** key while clicking each name.
- 4** In the **Data name** field, type the name of the new data set. This name must be unique in the Data Board.
- 5** Click **Insert** to add the new data set to the Data Board in the System Identification Tool GUI.
- 6** Click **Close** to close the Select Experiment dialog box.

Renaming Data and Changing Display Color

You can rename data and change its display color by double-clicking the data icon in the System Identification Tool GUI.

The Data/model Info dialog box opens. This dialog box describes both the contents and the properties of the data. The object description area displays the syntax of the operations you used to create the data in the GUI.

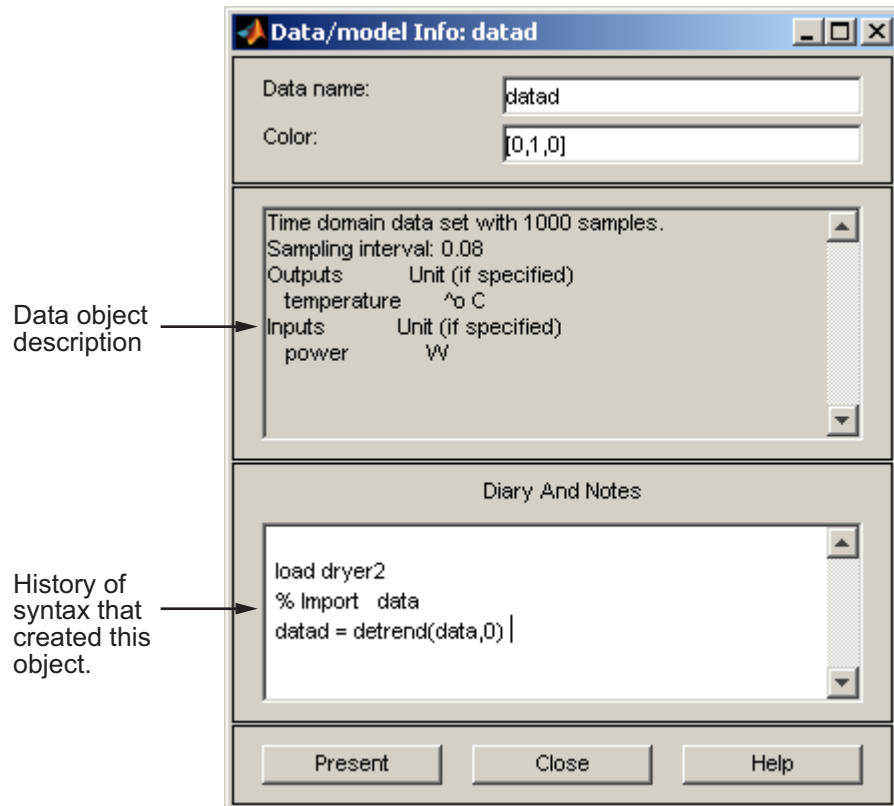
The Data/model Info dialog box also lets you rename the data by entering a new name in the **Data name** field.

You can also specify a new display color using three RGB values in the **Color** field. Each value is between 0 to 1 and indicates the relative presence of red, green, and blue, respectively. For more information about specifying default data color, see “Customizing the System Identification Tool GUI” on page 1-26.

Tip As an alternative to using three RGB values, you can enter any *one* of the following letters in single quotes:

'y' 'r' 'b' 'c' 'g' 'm' 'k'

These strings represent yellow, red, blue, cyan, green, magenta, and black, respectively.



Information About the Data

Finally, you can enter comments about the origin and state of the data in the **Diary And Notes** area. For example, you might want to include the experiment name, date, and the description of experimental conditions. When you estimate models from this data, these notes are associated with the models.

Clicking **Present** display the portions of this information in the MATLAB Command Window.

Creating iddata Objects

In this section...
“iddata Constructor” on page 2-32
“iddata Properties” on page 2-35
“Creating Multiexperiment Data at the Command Line” on page 2-38
“Subreferencing iddata Objects” on page 2-40
“Modifying Time and Frequency Vectors” on page 2-44
“Naming, Adding, and Removing Data Channels” on page 2-48
“Concatenating iddata Objects” on page 2-50
“See Also” on page 2-51

iddata Constructor

- “Before You Begin” on page 2-32
- “Constructing iddata for Time-Domain Data” on page 2-33
- “Constructing iddata for Frequency-Domain Data” on page 2-34

Before You Begin

To construct an iddata object, you must have the input and output data variables already in the MATLAB Workspace browser. Before you begin, check the requirements on data dimensions in “Data Requirements” on page 2-6.

Constructing iddata for Time-Domain Data

Use the following syntax to create a time-domain iddata object data:

```
data = iddata(y,u,Ts)
```

You can specify additional properties when you create the iddata object using the constructor syntax:

```
data = iddata(y,u,Ts,'Property1',Value1,...,'PropertyN',ValueN)
```

For more information about accessing object properties, see “iddata Properties” on page 2-35.

In this example, Ts is the sampling time, or the time interval, between successive data samples. For uniformly sampled data, Ts is a scalar value equal to the sampling interval of your experiment. The default time unit is seconds, but you can specify any unit string using the `TimeUnit` property. For more information about iddata time properties, see “Modifying Time and Frequency Vectors” on page 2-44.

For nonuniformly sampled data, specify Ts as [], and set the value of the `SamplingInstants` property as a column vector containing individual time values. For example:

```
data = iddata(y,u,Ts,[],'SamplingInstants',TimeVector)
```

Where `TimeVector` represents a vector of time values.

Note You can modify the property `SamplingInstants` by setting it to a new vector with the length equal to the number of data samples.

To represent time-series data, use the following syntax:

```
ts_data = iddata(y,[],Ts)
```

where y is the output data, [] indicates empty input data, and Ts is the sampling interval.

The following example shows how to create an `iddata` object using single-input/single-output (SISO) data from `dryer2.mat`. The input and output each contain 1000 samples with the sampling interval of 0.08 second.

```
load dryer2                % Load input u2 and output y2
data = iddata(y2,u2,0.08)  % Create iddata object
```

MATLAB returns the following output:

```
Time domain data set with 1000 samples.
Sampling interval: 0.08
```

```
Outputs      Unit (if specified)
  y1
```

```
Inputs      Unit (if specified)
  u1
```

The default channel name 'y1' is assigned to the first and only output channel. When `y2` contains several channels, the channels are assigned default names 'y1', 'y2', 'y2', ..., 'yn'. Similarly, the default channel name 'u1' is assigned to the first and only output channel. For more information about naming channels, see “Naming, Adding, and Removing Data Channels” on page 2-48.

Constructing `iddata` for Frequency-Domain Data

Frequency-domain data is the Fourier transform of the input and output signals at specific frequency values. To store frequency-domain data, use the following syntax to create the `iddata` object:

```
data = iddata(y,u,Ts,'Frequency',w)
```

'Frequency' is an `iddata` property that specifies the frequency value `w`, respectively. `Ts` is the time interval between successive data samples in seconds, and `w` is the frequency column vector that defines the frequencies at which System Identification Toolbox calculates the Fourier transform values of `y` and `u`. `w`, `y`, and `u` have the same number of rows.

Note You must specify the frequency vector.

For more information about iddata time and frequency properties, see “Modifying Time and Frequency Vectors” on page 2-44.

To specify a continuous-time system, set Ts to 0.

You can specify additional properties when you create the iddata object using the constructor syntax:

```
data = iddata(y,u,Ts,'Property1',Value1,...,'PropertyN',ValueN)
```

For more information about accessing object properties, see “iddata Properties” on page 2-35.

iddata Properties

To view the properties of the iddata object, use the get command. For example, type the following commands at the MATLAB prompt:

```
load dryer2                % Load input u2 and output y2
data = iddata(y2,u2,0.08); % Create iddata object
get(data)                   % Get property values of data
```

MATLAB returns the following object properties and values:

```
        Domain: 'Time'
          Name: []
    OutputData: [1000x1 double]
             y: 'Same as OutputData'
    OutputName: {'y1'}
    OutputUnit: {''}
     InputData: [1000x1 double]
             u: 'Same as InputData'
    InputName: {'u1'}
    InputUnit: {''}
         Period: Inf
    InterSample: 'zoh'
             Ts: 0.0800
          Tstart: []
    SamplingInstants: [1000x0 double]
          TimeUnit: ''
    ExperimentName: 'Exp1'
             Notes: []
          UserData: []
```

For a complete description of all properties, see the `iddata` reference page or type `idprops iddata` at the MATLAB prompt.

You can specify properties when you create an `iddata` object using the constructor syntax:

```
data = iddata(y,u,Ts,'Property1',Value1,...,'PropertyN',ValueN)
```

To change property values for an existing `iddata` object, use the `set` command or dot notation. For example, to change the sampling interval to 0.05, type the following at the MATLAB prompt:

```
set(data,'Ts',0.05)
```

or equivalently:

```
data.ts = 0.05
```

Property names are not case sensitive. You do not need to type the entire property name if the portion of the name you do enter uniquely identifies the property.

Tip You can use `data.y` as an alternative to `data.OutputData` to access the output values, or use `data.u` as an alternative to `data.InputData` to access the input values.

An `iddata` object containing frequency-domain data includes frequency-specific properties, such as `Frequency` for the frequency vector and `Units` for frequency units (instead of `Tstart` and `SamplingIntervals`).

To view the property list, type the following command sequence at the MATLAB prompt:

```
% Load input u2 and output y2
load dryer2;
% Create iddata object
data = iddata(y2,u2,0.08);
% Take the Fourier transform of the data
% transforming it to frequency domain
data = fft(data)
% Get property values of data
get(data)
```

MATLAB returns the following object properties and values:

```
Domain: 'Frequency'  
Name: []  
OutputData: [501x1 double]  
    y: 'Same as OutputData'  
OutputName: {'y1'}  
OutputUnit: {''}  
InputData: [501x1 double]  
    u: 'Same as InputData'  
InputName: {'u1'}  
InputUnit: {''}  
Period: Inf  
InterSample: 'zoh'  
    Ts: 0.0800  
Units: 'rad/s'  
Frequency: [501x1 double]  
TimeUnit: ''  
ExperimentName: 'Exp1'  
Notes: []  
UserData: []
```

Creating Multiexperiment Data at the Command Line

- “Why Create Multiexperiment Data Sets?” on page 2-38
- “Limitations on Data Sets” on page 2-39
- “Merging Data Sets” on page 2-39
- “Adding Experiments to an Existing iddata Object” on page 2-39

Why Create Multiexperiment Data Sets?

You can create `iddata` objects that contain several experiments. Identifying models for an `iddata` object with multiple experiments results in an *average* model that takes into account all data sets stored in the object.

In the context of System Identification Toolbox, *experiments* can mean data collected during different sessions, or portions of the data collected during a single session. In the latter situation, you can create a multiexperiment

iddata object by splitting the data from a single session into multiple segments to exclude bad data, and merge the good data portions.

Note The `idfrd` object does not support the `iddata` equivalent of multiexperiment data.

Limitations on Data Sets

You can only merge data sets that have all of the following characteristics:

- The same number of input and output channels.
- The same input and output channel names.
- The same type (that is, time-domain data or frequency-domain data)

Merging Data Sets

Create a multiexperiment `iddata` object by merging `iddata` objects, where each contains data from a single experiment or is a multiexperiment data set. For example, you can use the following syntax to merge data:

```
load iddata1      % Loads iddata object z1
load iddata3      % Loads iddata object z3
z = merge(z1,z3) % Merges experiments z1 and z3 into
                  % the iddata object z
```

This merge results in an `iddata` object with two experiments, where the experiments are assigned default names 'Exp1' and 'Exp2', respectively.

Adding Experiments to an Existing iddata Object

You can add experiments individually to an `iddata` object as an alternative approach to merging data sets, as described in “Merging Data Sets” on page 2-23.

For example, to add the experiments in the `iddata` object `dat4` to `data`, use the following syntax:

```
data(:,:,, 'Run4') = dat4
```

This syntax explicitly assigns the experiment name 'Run4' to the new experiment. The ExperimentName property of the iddata object stores experiment names.

For more information about subreferencing experiments in a multiexperiment data set, see “Subreferencing Experiments” on page 2-43.

Subreferencing iddata Objects

- “Subreferencing Input and Output Data” on page 2-40
- “Subreferencing Data Channels” on page 2-41
- “Subreferencing Experiments” on page 2-43

Subreferencing Input and Output Data

Subreferencing data and its properties lets you select data and assign new data and property values.

Use the following general syntax to subreference specific data values in iddata objects:

```
data(samples,outputchannels,inputchannels,experimentname)
```

In this syntax, samples specify one or more sample indexes, outputchannels and inputchannels specify channel indexes or channel names, and experimentname specifies experiment indexes or names.

For example, to retrieve samples 5 through 30 in the iddata object data and store them in a new iddata object data_sub, use the following syntax:

```
data_sub = data([5:30])
```

You can also use logical expressions to subreference data. For example, to retrieve all data values that fall between sample instants 1.27 and 9.3 in the iddata object data and assign them to data_sub, use the following syntax:

```
data_sub = data(data.sa>1.27&data.sa<9.3)
```

Note You do not need to type the entire property name. In this example, `sa` in `data.sa` uniquely identifies the `SamplingInstants` property.

You can retrieve the input signal from an `iddata` object using the following commands:

```
u = get(data, 'InputData')
```

or

```
data.InputData
```

or

```
data.u    % u is the abbreviation for InputData
```

Similarly, you can retrieve the output data using

```
data.OutputData
```

or

```
data.y    % y is the abbreviation for OutputData
```

Subreferencing Data Channels

Use the following general syntax to subreference specific data values in `iddata` objects:

```
data(samples,outputchannels,inputchannels,experimentname)
```

In this syntax, `samples` specify one or more sample indexes, `outputchannels` and `inputchannels` specify channel indexes or channel names, and `experimentname` specifies experiment indexes or names.

To specify several channel names, you must use a cell array of name strings.

For example, suppose the `iddata` object `data` contains three output channels (named `y1`, `y2`, and `y3`), and four input channels (named `u1`, `u2`, `u3`, and `u4`). To retrieve all data samples in `y3`, `u1`, and `u4`, type the following command at the MATLAB prompt:

```
% Uses a cell array to reference  
% input channels 'u1' and 'u4'  
data_sub = data(:, 'y3', {'u1', 'u4'})
```

or equivalently

```
% Uses channel indexes 1 and 4  
% to reference the input channels  
data_sub = data(:, 3, [1 4])
```

Tip Use a colon (`:`) to specify all samples or all channels, and the empty matrix (`[]`) to specify no samples or channels.

If you want to create a time-series object by extracting only the output data from an `iddata` object, type the following command:

```
data_ts = data(:, :, [])
```

You can assign new values to subreferenced variables. For example, the following command assigns the first 10 values of output channel 1 of `data` to values in samples 101 through 110 in the output channel 2 of `data1`. It also assigns the first 10 values of input channel 1 of `data` to values in samples 101 through 110 in the input channel 3 of `data1`.

```
data(1:10, 1, 1) = data1(101:110, 2, 3)
```

Subreferencing Experiments

Use the following general syntax to subreference specific experiments in iddata objects:

```
data(samples,outputchannels,inputchannels,experimentname)
```

In this syntax, `samples` specify one or more sample indexes, `outputchannels` and `inputchannels` specify channel indexes or channel names, and `experimentname` specifies experiment indexes or names.

When specifying several experiment names, you must use a cell array of name strings. The iddata object stores experiments name in the `ExperimentName` property.

For example, suppose the iddata object `data` contains five experiments with default names, `Exp1`, `Exp2`, `Exp3`, `Exp4`, and `Exp5`. Use the following syntax to subreference the first and fifth experiment in `data`:

```
data_sub = data(:, :, :, {'Exp1', 'Exp5'}) % Using experiment name
```

or

```
data_sub = data(:, :, :, [1 5]) % Using experiment index
```

Tip Use a colon (`:`) to denote all samples and all channels, and the empty matrix (`[]`) to specify no samples and no channels.

Alternatively, you can use the `getexp` command. The following example shows how to subreference the first and fifth experiment in `data`:

```
data_sub = getexp(data, {'Exp1', 'Exp5'}) % Using experiment name
```

or

```
data_sub = getexp(data, [1 5]) % Using experiment index
```

The following example shows how to retrieve the first 100 samples of output channels 2 and 3 and input channels 4 to 8 of Experiment 3:

```
dat(1:100, [2, 3], [4:8], 3)
```

Modifying Time and Frequency Vectors

The `iddata` object stores time-domain data or frequency-domain data and has several properties that specify the time or frequency values. To modify the time or frequency values, you must change the corresponding property values.

Note You can modify the property `SamplingInstants` by setting it to a new vector with the length equal to the number of data samples. For more information, see “Constructing `iddata` for Time-Domain Data” on page 2-33.

The following two tables summarize these time-vector and frequency-vector properties and gives an example of setting each value. In each of the syntax examples, `data` is an `iddata` object.

Note Property names are not case sensitive. You do not need to type the entire property name if the first few letters uniquely identify the property.

iddata Time-Vector Properties

Property	Description	Syntax Example
Ts	<p>Sampling time interval.</p> <ul style="list-style-type: none"> • For a single experiment, Ts is a scalar value. • For multiexperiment data with Ne experiments, Ts is a 1-by-Ne cell array, and each cell contains the sampling interval of the corresponding experiment. 	<p>To set the sampling interval to 0.05:</p> <pre>set(data, 'ts', 0.05)</pre> <p>or</p> <pre>data.ts = 0.05</pre>
Tstart	<p>Starting time of the experiment.</p> <ul style="list-style-type: none"> • For a single experiment, Ts is a scalar value. • For multiexperiment data with Ne experiments, Ts is a 1-by-Ne cell array, and each cell contains the sampling interval of the corresponding experiment. 	<p>To change starting time of the first data sample to 24:</p> <pre>data.Tstart = 24</pre> <p>Time units are set by the property TimeUnit.</p>

iddata Time-Vector Properties (Continued)

Property	Description	Syntax Example
SamplingInstants	<p>Time values in the time vector, computed from the properties Tstart and Ts.</p> <ul style="list-style-type: none"> • For a single experiment, SamplingInstants is an N-by-1 vector. • For multiexperiment data with Ne experiments, this property is a 1-by-Ne cell array, and each cell contains the sampling instants of the corresponding experiment. 	<p>To retrieve the time vector for iddata object data, use: <code>get(data, 'sa')</code> To plot the input data as a command of time: <code>plot(data.sa,data.u)</code></p> <hr/> <p>Note sa uniquely identifies the SamplingInstants property.</p>
TimeUnit	Unit of time.	<p>To change the unit of the time vector to msec: <code>data.ti = 'msec'</code></p>

iddata Frequency-Vector Properties

Property	Description	Syntax Example
Frequency	<p>Frequency values at which the Fourier transforms of the signals are defined.</p> <ul style="list-style-type: none"> • For a single experiment, Frequency is a scalar value. • For multiexperiment data with N_e experiments, Frequency is a 1-by-N_e cell array, and each cell contains the frequencies of the corresponding experiment. 	<p>To specify 100 frequency values in log space, ranging between 0.1 and 100, use the following syntax:</p> <pre>data.freq = logspace(-1,2,100)</pre>
Units	<p>Frequency unit must have the following values:</p> <ul style="list-style-type: none"> • If the TimeUnit is empty or s (seconds), enter rad/s or Hz • If the TimeUnit is any string <i>unit</i> (other than s), enter rad/<i>unit</i>. <p>For multiexperiment data with N_e experiments, Units is a 1-by-N_e cell array, and each cell contains the frequency unit for each experiment.</p>	<p>If you specified the TimeUnit as msec, your frequency units must be:</p> <pre>data.unit= 'rad/msec'</pre>

Naming, Adding, and Removing Data Channels

- “What Are Input and Output Channels?” on page 2-48
- “Naming Channels” on page 2-48
- “Adding Channels” on page 2-49
- “Modifying Channel Data” on page 2-49

What Are Input and Output Channels?

A multivariate system might contain several input variables or several output variables, or both. When an input or output signal includes several measured variables, these variables are called *channels*.

Naming Channels

The `iddata` properties `InputName` and `OutputName` store one or more channel names for the input and output signals. When you plot the data, you use channel names to select the variable displayed on the plot. If you have multivariate data, you should assign a name to each channel that describes the measured variable. For more information about selecting channels on a plot, see “Selecting Measured and Noise Channels in Plots” on page 1-23.

You can use the `set` command to specify the names of individual channels. For example, suppose `data` contains two input channels (voltage and current) and one output channel (temperature). To set these channel names, use the following syntax:

```
set(data, 'InputName', {'Voltage', 'Current'},  
      'OutputName', 'Temperature')
```

Tip You can also specify channel names as follows:

```
data.una = {'Voltage', 'Current'}  
data.yna = 'Temperature'
```

`una` is equivalent to the property `InputName`, and `yna` is equivalent to `OutputName`.

If you do not specify channel names when you create the `iddata` object, System Identification Toolbox assigns default names. By default, the output channels are named `'y1'`, `'y2'`, ..., `'yn'`, and the input channels are named `'u1'`, `'u2'`, ..., `'un'`.

Adding Channels

You can add channels to input and output data in an `iddata` object.

For example, consider an `iddata` object named `data` that contains an input signal with four channels. To add a fifth input channel, stored in the MATLAB vector `Input5`, use the following syntax:

```
data.u(:,5) = Input5;
```

In this example, `data.u(:,5)` references all samples as (indicated by `:`) of the input signal `u` and sets the values of the fifth channel. This channel is created when assigning its value to `Input5`.

You can also combine input channels and output channels of several `iddata` objects into one `iddata` object using concatenation. For more information, see “Concatenating `iddata` Objects” on page 2-50.

Modifying Channel Data

After you create an `iddata` object, you can modify or remove specific input and output channels, if needed. You can accomplish this by subreferencing the input and output matrices and assigning new values.

For example, suppose the `iddata` object `data` contains three output channels (named `y1`, `y2`, and `y3`), and four input channels (named `u1`, `u2`, `u3`, and `u4`). To replace data such that it only contains samples in `y3`, `u1`, and `u4`, type the following at the MATLAB prompt:

```
data = data(:,3,[1 4])
```

The resulting `data` object contains one output channel and two input channels.

Concatenating iddata Objects

- “iddata Properties Storing Input and Output Data” on page 2-50
- “Horizontal Concatenation” on page 2-50
- “Vertical Concatenation” on page 2-51

iddata Properties Storing Input and Output Data

The InputData iddata property stores column-wise input data, and the OutputData property stores column-wise output data. For more information about accessing iddata properties, see “iddata Properties” on page 2-35.

Horizontal Concatenation

Horizontal concatenation of iddata objects creates a new iddata object that appends all InputData information and all OutputData. This type of concatenation produces a single object with more inputs and more outputs. For example, the following syntax performs horizontal concatenation on the iddata objects data1, data2, . . . , dataN:

```
data = [data1,data2,...,dataN]
```

This syntax is equivalent to the following longer syntax:

```
data.InputData =  
    [data1.InputData,data2.InputData,...,dataN.InputData]  
data.OutputData =  
    [data1.OutputData,data2.OutputData,...,dataN.OutputData]
```

For horizontal concatenation, data1, data2, . . . , dataN must have the same number of samples and experiments and equal Ts and Tstart values.

The channels in the concatenated iddata object are named according to the following rules:

- **Combining default channel names.** If you concatenate iddata objects with default channel names, such as u1 and y1, channels in the new iddata object are automatically renamed to avoid name duplication.
- **Combining duplicate input channels.** If data1, data2, . . . , dataN have input channels with duplicate user-defined names, such that dataK

contains channel names that are already present in `dataJ` with $J < K$, the `dataK` channels are ignored.

- **Combining duplicate output channels.** If `data1`, `data2`, ..., `dataN` have input channels with duplicate user-defined names, only the output channels with unique names are added during the concatenation.

Vertical Concatenation

Vertical concatenation of `iddata` objects creates a new `iddata` object that vertically stacks the input and output data values in the corresponding data channels. The resulting object has the same number of channels, but each channel contains more data points. For example, the following syntax creates a data object such that its total number of samples is the sum of the samples in `data1`, `data2`, ..., `dataN`.

```
data = [data1;data2;... ;dataN]
```

This syntax is equivalent to the following longer syntax:

```
data.InputData =
    [data1.InputData;data2.InputData;... ;dataN.InputData]
data.OutputData =
    [data1.OutputData;data2.OutputData;... ;dataN.OutputData]
```

For vertical concatenation, `data1`, `data2`, ..., `dataN` must have the same number of input channels, output channels, and experiments.

See Also

To represent data in the System Identification Tool GUI instead, see “Importing Data into the GUI” on page 2-14.

For more information about working with frequency-response data, see “Creating `idfrd` Objects” on page 2-52.

You can create multiexperiment `iddata` objects by combining existing `iddata` objects. For more information, see “Creating Multiexperiment Data at the Command Line” on page 2-38.

Creating idfrd Objects

In this section...

“idfrd Constructor” on page 2-52

“idfrd Properties” on page 2-53

“Subreferencing idfrd Objects” on page 2-55

“Concatenating idfrd Objects” on page 2-56

“See Also” on page 2-59

idfrd Constructor

The `idfrd` represents complex frequency-response data. Before you can create an `idfrd` object, see “Representing Frequency-Response Data” on page 2-10.

Note The `idfrd` object can only encapsulate one frequency-response data set. It does not support the `iddata` equivalent of multiexperiment data.

Use the following syntax to create the data object `fr_data`:

```
fr_data = idfrd(response,f,Ts)
```

Suppose that `ny` is the number of output channels, `nu` is the number of input channels, and `nf` is a vector of frequency values. `response` is an `ny-by-nu-by-nf` 3-D array. `f` is the frequency vector that contains the frequencies of the response. `Ts` is the sampling time, which is used when measuring or computing the frequency response. If you are working with a continuous-time system, set `Ts` to 0.

`response(ky,ku,kf)`, where `ky`, `ku`, and `kf` reference the `k`th output, input, and frequency value, respectively, is interpreted as the complex-valued frequency response from input `ku` to output `ky` at frequency `f(kf)`.

Note When you work at the command line, you can only create idfrd objects from complex values of $G(e^{iw})$. For a SISO system, response can be a vector.

You can specify object properties when you create the idfrd object using the constructor syntax:

```
fr_data = idfrd(response,f,Ts,  
              'Property1',Value1,...,'PropertyN',ValueN)
```

idfrd Properties

To view the properties of the idfrd object, you can use the get command. The following example shows how to create an idfrd object that contains 100 frequency-response values with a sampling time interval of 0.08 s and get its properties:

```
% Create the idfrd data object  
fr_data = idfrd(response,f,0.08)  
% Get property values of data  
get(fr_data)
```

response and `f` are variables in the MATLAB Workspace browser, representing the frequency-response data and frequency values, respectively.

MATLAB returns the following object properties and values:

```
ans =  
  
    Name: ''  
    Frequency: [100x1 double]  
    ResponseData: [1x1x100 double]  
    SpectrumData: []  
    CovarianceData: []  
    NoiseCovariance: []  
    Units: 'rad/s'  
    Ts: 0.0800  
    InputDelay: 0  
    EstimationInfo: [1x1 struct]  
    InputName: {'u1'}  
    OutputName: {'y1'}  
    InputUnit: {''}  
    OutputUnit: {''}  
    Notes: []  
    UserData: []
```

For a complete description of all `idfrd` object properties, see the `idfrd` reference page or type `idprops idfrd` at the MATLAB prompt.

To change property values for an existing `idfrd` object, use the `set` command or dot notation. For example, to change the name of the `idfrd` object, type the following command sequence at the MATLAB prompt:

```
% Set the name of the f_data object  
set(fr_data,'name','DC_Converter')  
% Get fr_data properties and values  
get(fr_data)
```

Property names are not case sensitive. You do not need to type the entire property name, but only as much of the name as uniquely identify the `idfrd` property.

If you import `fr_data` into the System Identification Tool GUI, this data has the name `DC_Converter` in the GUI, and not the MATLAB variable name `fr_data`.

MATLAB returns the following object properties and values:

```
ans =  
  
        Name: 'DC_Converter'  
      Frequency: [100x1 double]  
    ResponseData: [1x1x100 double]  
   SpectrumData: []  
CovarianceData: []  
NoiseCovariance: []  
          Units: 'rad/s'  
           Ts: 0.0800  
   InputDelay: 0  
EstimationInfo: [1x1 struct]  
   InputName: {'u1'}  
  OutputName: {'y1'}  
   InputUnit: {''}  
  OutputUnit: {''}  
         Notes: []  
       UserData: []
```

Subreferencing idfrd Objects

You can reference specific data values in the `idfrd` object using the following syntax:

```
fr_data(outputchannels,inputchannels)
```

Reference specific channels by name or by channel index.

Tip Use a colon (`:`) to specify all channels, and use the empty matrix (`[]`) to specify no channels.

For example, the following command references frequency-response data from input channel 3 to output channel 2:

```
fr_data(2,3)
```

You can also access the data in specific channels using channel names. To list multiple channel names, use a cell array. For example, to retrieve the power output, and the voltage and speed inputs, use the following syntax:

```
fr_data('power',{'voltage','speed'})
```

To retrieve only the responses corresponding to frequency values between 200 and 300, use the following command:

```
fr_data_sub = fselect(fr_data,[200:300])
```

You can also use logical expressions to subreference data. For example, to retrieve all frequency-response values between frequencies 1.27 and 9.3 in the `idfrd` object `fr_data`, use the following syntax:

```
fr_data_sub = fselect(fr_data,fr_data.f>1.27&fr_data.f<9.3)
```

Note You do not need to type the entire property name. In this example, `f` in `fr_data.f` uniquely identifies the Frequency property of the `idfrd` object.

Concatenating `idfrd` Objects

- “About Concatenating `idfrd` Models” on page 2-56
- “Horizontal Concatenation of `idfrd` Objects” on page 2-57
- “Vertical Concatenation of `idfrd` Objects” on page 2-57
- “Concatenating Noise Spectral Data of `idfrd` Objects” on page 2-58

About Concatenating `idfrd` Models

The horizontal and vertical concatenation of `idfrd` objects combine information in the `ResponseData` properties of these objects. `ResponseData` is an `ny-by-nu-by-nf` array that stores the response of the system, where `ny` is

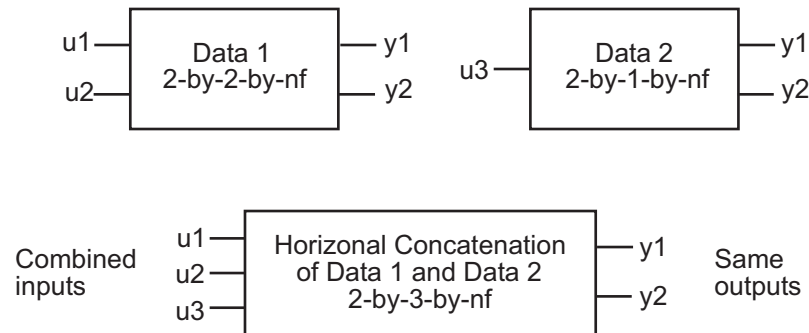
the number of output channels, nu is the number of input channels, and nf is a vector of frequency values (see “idfrd Properties” on page 2-53).

Horizontal Concatenation of idfrd Objects

The following syntax creates a new `idfrd` object `data` that contains the horizontal concatenation of `data1`, `data2`, ..., `dataN`:

```
data = [data1,data2,...,dataN]
```

`data` contains the frequency responses from all of the inputs in `data1`, `data2`, ..., `dataN` to the same outputs. The following diagram is a graphical representation of horizontal concatenation of frequency-response data. The $(j, i, :)$ vector of the resulting response data represents the frequency response from the i th input to the j th output at all frequencies.



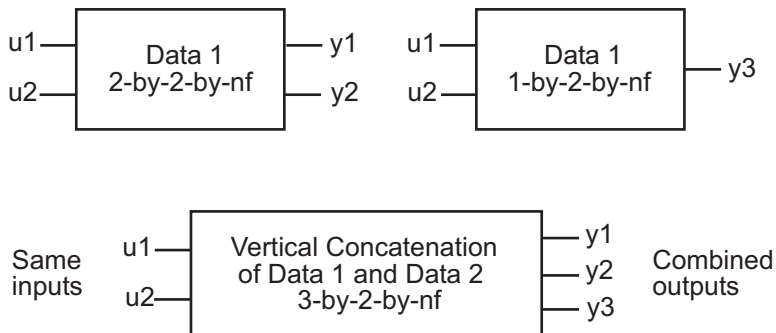
Note Horizontal concatenation of `idfrd` objects requires that they have the same outputs and frequency vectors. If the output channel names are different and their dimensions are the same, the concatenation operation uses the names of output channels in the first `idfrd` object. Input channels must have unique names.

Vertical Concatenation of idfrd Objects

The following syntax creates a new `idfrd` object `data` that contains the vertical concatenation of `data1`, `data2`, ..., `dataN`:

```
data = [data1;data2;... ;dataN]
```

The resulting `idfrd` object data contains the frequency responses from the same inputs in `data1`, `data2`, \dots , `dataN` to all the outputs. The following diagram is a graphical representation of vertical concatenation of frequency-response data. The $(j, i, :)$ vector of the resulting response data represents the frequency response from the i th input to the j th output at all frequencies.



Note Vertical concatenation of `idfrd` objects requires that they have the same inputs and frequency vectors. If the input channel names are different and their dimensions are the same, the concatenation operation uses the names of input channels in the first `idfrd` object you listed. Output channels must have unique names.

Concatenating Noise Spectral Data of `idfrd` Objects

When the `idfrd` objects contain the frequency-response data you measured or constructed manually, the concatenation operation combines only the `ResponseData` properties. Because the noise spectral data does not exist (unless you also entered it manually), `SpectralData` is empty in both the individual `idfrd` objects and the concatenated `idfrd` object.

However, when the `idfrd` objects are spectral models that you estimated, the `SpectralData` property is not empty and contains the power spectra and cross spectra of the output noise in the system. For each output channel, System Identification Toolbox estimates one noise channel to explain the difference between the output of the model and the measured output.

When the `SpectralData` property of individual `idfrd` objects is not empty, horizontal and vertical concatenation handle `SpectralData`, as follows.

In case of horizontal concatenation, there is no meaningful way to combine the `SpectralData` of individual `idfrd` objects and the resulting `SpectralData` property is empty. An empty property results because each `idfrd` object has its own set of noise channels, where the number of noise channels equals the number of outputs. When the resulting `idfrd` object contains the same output channels as each of the individual `idfrd` objects, it cannot accommodate the noise data from all the `idfrd` objects.

In case of vertical concatenation, System Identification Toolbox concatenates individual noise models diagonally. The following shows that `data.SpectrumData` is a block diagonal matrix of the power spectra and cross spectra of the output noise in the system:

$$data.s = \begin{pmatrix} data1.s & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & dataN.s \end{pmatrix}$$

`s` in `data.s` is the abbreviation for the `SpectrumData` property name.

See Also

The following System Identification Toolbox operations also create `idfrd` objects:

- Transforming `iddata` objects. For more information, see “Transforming Between Frequency-Domain and Frequency-Response Data” on page 2-112.
- Estimating nonparametric models using `etfe`, `spa`, and `spafdr`. For more information, see “Estimating Frequency-Response Models” on page 4-10.
- Converting the Control System Toolbox `frd` object. For more information, see “Using Models with Control System Toolbox” on page 13-2.

Getting Advice About Your Data

You can use the `advice` command to get information about your time-domain or frequency-domain data. This command does not support frequency-response data.

Note If you are using the System Identification Tool GUI, you must export your data to the MATLAB Workspace browser before you can use the `advice` command on this data. For more information about exporting data to the MATLAB Workspace browser, see “Exporting Data and Models to the MATLAB Workspace Browser” on page 1-16.

Suppose that `data` is an `iddata` object. `advice(data)` displays the following information about the data in the MATLAB Command Window. Ask yourself the following questions:

- Does it make sense to remove constant offsets and linear trends from the data? See also `detrend`.
- What are the excitation levels of the signals and how does this affects the model orders? See also `pexcit`.
- Is there an indication of output feedback in the data? See also `feedback`. When feedback is present in the system, only prediction-error methods work well for estimating closed-loop data.

To estimate the delay from the input to the output in the system (dead time) by using the data, use the `delayest` command. You need to know the delay when specifying a model structure for estimation.

The following example shows how to get information about your data. Consider data from a single-input/single-output system sampled at 0.08 s. Use these commands to load the data and estimate the delay in the system:

```
load dryer2                % Load the sample input
                           % and output data
data=iddata(y2,u2,0.08)    % Create iddata object
delayest(data)            % Estimate delay (dead time)
```

MATLAB responds with:

```
ans =
```

```
3
```

Type the following command to see what kind of information you receive about this data set:

```
advice(data)           % Get advice about the data
```

The results of using this command also suggest your next actions.

Plotting Data

In this section...
“Supported Plots for Data” on page 2-62
“Plotting Data in the System Identification Tool GUI” on page 2-62
“How to Plot Data at the Command Line” on page 2-68

Supported Plots for Data

System Identification Toolbox supports the following data plots:

- Time plot — Shows data values as a command of time.
- Spectral plot — Shows a *periodogram* that is computed by taking the absolute squares of the Fourier transforms of the data, dividing by the number of data points, and multiplying by the sampling interval.
- Frequency-response plot — For frequency-response data, shows the amplitude and phase of the frequency-response command. For time- and frequency-domain data, shows the empirical transfer command estimate (see `etfe`).

The plots you create using the System Identification Tool GUI provide different options that are specific to System Identification Toolbox, such as selecting specific channel pairs in a multivariate signals or converting frequency units between hertz and radians per second.

The plots you create using the plot commands, such as `plot`, `bode`, and `ffplot`, are displayed in the standard MATLAB Figure window, which provides extensive options for formatting, saving, printing, and exporting plots to a variety of file formats. For more information, see the MATLAB Graphics documentation.

Plotting Data in the System Identification Tool GUI

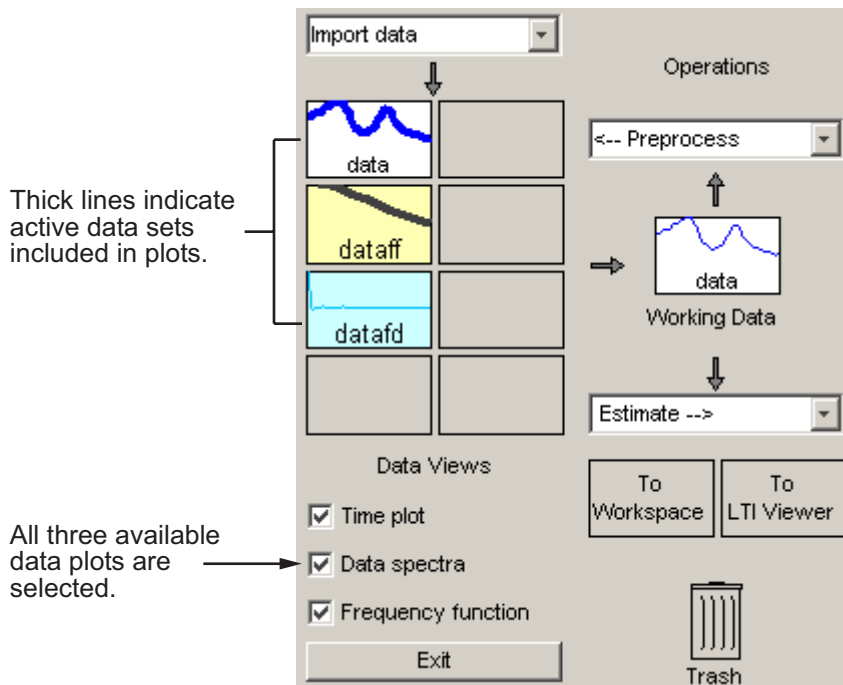
- “How to Plot Data in the GUI” on page 2-63
- “Working with a Time Plot” on page 2-64

- “Working with a Data Spectra Plot” on page 2-65
- “Working with a Frequency Function Plot” on page 2-67

How to Plot Data in the GUI

After importing data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14, you can plot the data. To create one or more plots, select the corresponding check box in the **Data Views** area of the System Identification Tool GUI.


Active data sets have a thick line in the icon. Only active data sets are displayed on the plots. To change whether to include or exclude a data set on a plot, click the data icon in the System Identification Tool GUI.



In this example, data1 and data3fd are active and appear on the three selected plots.

Working with a Time Plot

The **Time plot** only shows time-domain data. In this example, `data1` is displayed on the time plot because, of the three data sets, it is the only one that contains time-domain input and output.



Time Plot of data1

Note You can plot several data sets with the same input and output channel names. The plot displays data for all channels that have the same name. To view a different input-output channel pair, select it from the **Channel** menu. For more information about selecting different input and output pairs, see “Selecting Measured and Noise Channels in Plots” on page 1-23.

The following table summarizes options that are specific to time plots, which you can select from the plot window menus. For general information about working with System Identification Toolbox plots, see “Plots Data and Models in the GUI” on page 1-18.

Time Plot Options


Action	Command
<p>Toggle input display between piece-wise continuous (zero-order hold) and linear interpolation (first-order hold) between samples.</p> <hr/> <p>Note This option only affects the display and not the intersample behavior specified when importing the data.</p> <hr/>	<p>Select Style > Staircase input for zero-order hold or Style > Regular input for first-order hold.</p>

Working with a Data Spectra Plot

The **Data spectra** plot shows a periodogram or a spectral estimate of data1 and data3fd.

The periodogram is that is computed by taking the absolute squares of the Fourier transforms of the data, dividing by the number of data points, and multiplying by the sampling interval. The spectral estimate for time-domain data is a smoothed spectrum calculated using spa. For frequency-domain data, the **Data spectra** plot shows the absolute value of the square of the actual data.

The top axes show the input and the bottom axes show the output. The vertical axis of each plot is labeled with the corresponding channel name.



Periodograms of data1 and data3fd


Data Spectra Plot Options

Action	Command
Toggle display between periodogram and spectral estimate.	Select Options > Periodogram or Options > Spectral analysis .
Change frequency units.	Select Style > Frequency (rad/s) or Style > Frequency (Hz) .
Toggle frequency scale between linear and logarithmic.	Select Style > Linear frequency scale or Style > Log frequency scale .
Toggle amplitude scale between linear and logarithmic.	Select Style > Linear amplitude scale or Style > Log amplitude scale .

Working with a Frequency Function Plot

For time-domain data, the **Frequency** command plot shows the empirical transfer command estimate (etfe). For frequency-domain data, the plot shows the ratio of output to input data.

The frequency-response plot shows the amplitude and phase plots of the corresponding frequency response. For more information about frequency-response data, see “Representing Frequency-Response Data” on page 2-10.



Frequency Functions of data1 and data3fd

Frequency Function Plot Options

Action	Command
Change frequency units.	Select Style > Frequency (rad/s) or Style > Frequency (Hz) .

Frequency Function Plot Options (Continued)

Action	Command
Toggle frequency scale between linear and logarithmic.	Select Style > Linear frequency scale or Style > Log frequency scale .
Toggle amplitude scale between linear and logarithmic.	Select Style > Linear amplitude scale or Style > Log amplitude scale .

How to Plot Data at the Command Line

The following table summarizes the commands available for plotting time-domain, frequency-domain, and frequency-response data.

Commands for Plotting Data

Command	Description	Example
bode	For frequency-response data only. Shows the magnitude and phase of the frequency response on a logarithmic frequency scale.	To plot idfrd data: <code>bode(idfrd_data)</code>

Commands for Plotting Data (Continued)

Command	Description	Example
<code>ffplot</code>	For frequency-response data only. Shows the magnitude and phase of the frequency response on a linear frequency scale (hertz).	To plot <code>idfrd</code> data: <code>ffplot(idfrd_data)</code>
<code>plot</code>	Depending on the type of data, MATLAB generates the corresponding type of plot. For example, plotting time-domain data generates a time plot, and plotting frequency-response data generates a frequency-response plot. When plotting time- or frequency-domain inputs and outputs, the top axes show the output and the bottom axes show the input.	To plot <code>iddata</code> or <code>idfrd</code> data: <code>plot(data)</code> Note For <code>idfrd</code> data, this command is equivalent to <code>ffplot(data)</code> .

All plot commands display the data in the standard MATLAB Figure window. For more information about working with the Figure window, see the MATLAB Graphics documentation.

The following examples show use of the `plot` command.

To plot portions of the data, you can subreference specific samples (see “Subreferencing `iddata` Objects” on page 2-40 and “Subreferencing `idfrd` Objects” on page 2-55. For example:

```
plot(data(1:300))
```

For time-domain data, to plot only the input data as a command of time, use the following syntax:

```
plot(data(:,[]),:)
```

When `data.intersample = 'zoh'`, the input is piece-wise constant between sampling points on the plot. For more information about properties, see “`iddata Properties`” on page 2-35 or the `iddata` reference page.

You can generate plots of the input data in the time domain using:

```
plot(data.sa,data.u)
```

To plot frequency-domain data, you can use the following syntax:

```
semilogx(data.fr,abs(data.u))
```

In this case, `sa` is an abbreviation of the `iddata` property `SamplingInstants`. Similarly, `fr` is an abbreviation of `Frequency`. `u` is the input signal.

Note The frequencies are linearly spaced on the plot.

When you specify to plot a multivariable `iddata` object, each input-output combination is displayed one at a time in the same MATLAB Figure window. You must press **Enter** to update the Figure window and view the next channel combination. To cancel the plotting operation, press **Ctrl+C**.

Tip To plot specific input and output channels, use `plot(data(:,ky,ku))`, where `ky` and `ku` are specific output and input channel indexes or names. For more information about subreferencing channels, see “`Subreferencing Data Channels`” on page 2-41.

To plot several `iddata` sets `d1, ..., dN`, use `plot(d1, ..., dN)`. Input-output channels with the same experiment name, input name, and output name are always plotted in the same plot.

Handling Missing Data and Outliers

In this section...

“Working with Missing Data” on page 2-71

“Working with Outliers” on page 2-72

“Example – Extracting and Modeling Specific Data Segments” on page 2-73

“See Also” on page 2-74

Working with Missing Data

Data acquisition failures sometimes result in missing measurements both in the input and the output channels. When you import data that contains missing values into MATLAB using the MATLAB Import Wizard, these values are automatically set to NaN (“Not-a-Number”). NaN serves as a flag for nonexistent or undefined data. When you plot data on a time-plot that contains missing values, gaps appear on the plot where missing data exists.

You can use `misdata` to estimate missing values. This command linearly interpolates missing values to estimate the first model. Then, it uses this model to estimate the missing data as parameters by minimizing the output prediction errors obtained from the reconstructed data. You can specify the model structure you want to use in the `misdata` argument or estimate a default-order model using the `n4sid` method. For more information, see the `misdata` reference page.

Note You can only use `misdata` on time-domain data stored in an `iddata` object. For more information about creating `iddata` objects, see “Creating `iddata` Objects” on page 2-32.

For example, suppose `y` and `u` are output and input signals that contain NaNs. This data is sampled at 0.2 s. The following syntax creates a new `iddata` object with these input and output signals.

```
dat = iddata(y,u,0.2) % y and u contain NaNs
                        % representing missing data
```

Apply the `misdata` command to the new data object. For example:

```
dat1 = misdata(dat);  
plot(dat,dat1)           % Check how the missing data  
                          % was estimated on a time plot
```

Working with Outliers

Malfunctions can also produce errors in measured values, called *outliers*. Such outliers might be caused by signal spikes, or by measurement malfunctions. If you do not remove outliers from your data, this can adversely affect the estimated models. You can use System Identification Toolbox to identify and handle outliers in your data.

To identify the presence of outliers, perform one of the following tasks:

- Before estimating a model, plot the data on a time plot and identify values that appear unreliable.
- After estimating a model, plot the residuals and identify unusually large values. Then, evaluate the original data that is responsible for large residuals. For example, for the model `Model` and validation data `Data`, you can use the following commands to plot the residuals:

```
% Compute the residuals  
E = resid(Model,Data)  
% Plot the residuals  
plot(E)
```

Next, try these techniques for removing or minimizing the effects of outliers:

- Extract the informative data portions into segments and merge them into one multiexperiment data set (see “Example – Extracting and Modeling Specific Data Segments” on page 2-73). For more information about selecting and extracting data segments, see “Selecting Data” on page 2-95.

Tip The inputs in each of the data segments must be consistently exciting the system. Splitting data into meaningful segments for steady-state data results in minimum information loss. Avoid making data segments too small.

- Manually replace outliers with NaNs and then use the `misdata` command to reconstruct flagged data. This approach treats outliers as missing data and is described in “Working with Missing Data” on page 2-71. Use this method when your data contains several inputs and outputs, and when you have difficulty finding reliable data segments in all variables.
- Remove outliers by prefiltering the data for high-frequency content because outliers often result from abrupt changes. For more information about filtering, see “Filtering Data” on page 2-87.

Note The estimation algorithm handles outliers automatically by assigning a smaller weight to outlier data. A robust error criterion applies an error penalty that is quadratic for small and moderate prediction errors, and is linear for large prediction errors. Because outliers produce large prediction errors, this approach gives a smaller weight to the corresponding data points during model estimation. The value `LimitError` (see `Algorithm Properties`) quantitatively distinguishes between moderate and large outliers.

Example – Extracting and Modeling Specific Data Segments

The following example shows how to create a multiexperiment, time-domain data set by merging only the accurate-data segments and ignoring the rest. Modeling multiexperiment data sets produces an average model for the different experiments.

You cannot simply concatenate the good data segments because the transients at the connection points compromise the model. Instead, you must create a multiexperiment `iddata` object, where each experiment corresponds to a good segment of data, as follows:

```
% Plot the data in a MATLAB Figure window
plot(data)

% Create multiexperiment data set
% by merging data segments
datam = merge(data(1:340),...
              data(500:897),...
              data(1001:1200),...
              data(1550:2000));

% Model the multiexperiment data set
% using "experiments" 1, 2, and 4
m =pem(getexp(datam,[1,2,4]))

% Validate the model by comparing its output to
% the output data of experiment 3
compare(getexp(datam,3),m)
```

See Also

To learn more about the theory of handling missing data and outliers, see the chapter on preprocessing data in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

Detrending Data

In this section...

“What Is Detrending Data?” on page 2-75

“Detrending Data for Nonlinear Versus Linear Models” on page 2-75

“When to Subtract the Mean Values” on page 2-76

“When to Subtract Linear Trends” on page 2-77

“How to Detrend Data Using the GUI” on page 2-77

“How to Detrend Data Using `detrend`” on page 2-78

“See Also” on page 2-79

What Is Detrending Data?

Detrending data in System Identification Toolbox means removing the signal mean values and linear trends from the signals.

In System Identification Tool GUI, you can subtract the mean values and one linear trend. If you are working at the command line, the `detrend` command lets you subtract mean values and one or several linear trends connected at specified breakpoints. A *breakpoint* is a time value that defines the discontinuities between successive linear trends.

For a detailed discussion about handling drifts in the data, see the chapter on preprocessing data in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

Detrending Data for Nonlinear Versus Linear Models

In general, you should detrend data before estimating linear models. For more information, see “When to Subtract the Mean Values” on page 2-76 and “When to Subtract Linear Trends” on page 2-77.

In general, you should not detrend data before estimating nonlinear models. In the case of nonlinear grey-box models, do not detrend the data to make sure that the models represent the actual physical levels.

For linear models, detrending is necessary because, theoretically, a zero mean in the input results in a zero mean in the output.

Nonlinear models are more general and can include the trend as part of the model. In this case, detrending is not necessary.

However, there are some cases when detrending data for nonlinear models might be helpful. For example, to improve computation accuracy, you might detrend data with a large constant trend.

Note Whatever trend you subtract from the estimation data, you must subtract the same trend from the validation data.

When to Subtract the Mean Values

You may find it useful to subtract mean values from your data when you have steady-state and not transient data. If you have steady-state data, estimating linear models for signals measured relative to an equilibrium is usually sufficient. Thus, you can find linearized models around a physical equilibrium and avoid modeling the absolute levels in physical units. You can use the `detrend` command or System Identification Tool GUI commands to subtract mean levels from your signals.

Tip When you know the mean levels that correspond to the actual physical equilibrium, remove the equilibrium values instead of the mean value of the signals for best results.

Do not detrend your data when the physical levels are built into the underlying model or when input integrators in the system require absolute signal levels.

When you are working with transient data (such as step or impulse response), do not remove the mean from the data. With transient data, when the output at zero input is not zero, you might want to subtract the constant value corresponding to the time before the input is applied. Use core MATLAB commands to subtract a constant value from the data matrices.

When to Subtract Linear Trends

Often, the mean levels drifts during the experiment. You can eliminate drift by removing a linear trend or several piece-wise linear trends from the signals.

You can use the `detrend` command or the System Identification Tool GUI to subtract one linear trend from your signals from time-domain data. You must use the `detrend` command to remove several piece-wise-linear trends in time-domain data or to remove the mean for frequency-domain data.

Signal drift is considered a low-frequency disturbance. If you know the drift rate, you can build a custom high-pass filter and apply it as described in “Filtering Data” on page 2-87.

How to Detrend Data Using the GUI

If you are working in the System Identification Tool GUI, you can detrend time-domain data by removing the mean value and linear trend from each channel. If you need to remove piece-wise linear trends, see “How to Detrend Data Using `detrend`” on page 2-78. See “Detrending Data for Nonlinear Versus Linear Models” on page 2-75 to decide whether it is appropriate to detrend your data.

Note Select **Preprocess > Quick start** to remove the mean value from each channel, split data into two halves, specify the first half as estimation data for models (or **Working Data**), and specify the second half as **Validation Data**.

To detrend each input and output data channel:

- 1** Import time-domain data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14.
- 2** Drag the data set you want to detrend to the **Working Data** area.
- 3** Determine if you want to remove both the mean values and the linear trend from the data.
 - If yes, select **Preprocess > Remove trends**. This creates a new data set in the Data Board. You are finished.

- If no, go to step 4.

4 To only remove constant offsets from the data, select **Preprocess > Remove means**. This selection creates a new data set in the Data Board.

How to Detrend Data Using `detrend`

You can use the `detrend` command to remove the mean value and one or more linear trends from each channel of time-domain data. See “Detrending Data for Nonlinear Versus Linear Models” on page 2-75 to decide whether it is appropriate to detrend your data.

It is only possible to remove a zero-order trend from frequency-domain data. Detrending frequency-domain data is equivalent to setting the response to 0 at zero frequency.

Note You can only remove piece-wise linear trends using `detrend` and not System Identification Tool GUI.

To remove mean values from each channel in `data`, which is an `iddata` object that stores time-domain data or frequency-domain data, use the following syntax:

```
data = detrend(data);
```

To subtract one linear trend, use the following syntax:

```
data = detrend(data,1);
```

In this case, 1 indicates that a first-order trend is removed from each channel.

You can also remove several linear trends that have connections at specified “breakpoints”. For example:

```
data = detrend(data,1,[30 60 90]);
```

The vector `[30 60 90]` specifies the sample indexes of the data where breakpoints occur.

See Also

For general information about working with the System Identification Tool GUI, see Chapter 1, “Working with the System Identification Tool GUI”.

Resampling Data

In this section...
“What Is Resampling?” on page 2-80
“Resampling Data Using the GUI” on page 2-81
“Resampling Data at the Command Line” on page 2-81
“Resampling Data Without Aliasing Effects” on page 2-83
“See Also” on page 2-86

What Is Resampling?

Resampling data in System Identification Toolbox applies an antialiasing (lowpass) FIR filter to the data and changes the sampling rate of the signal by decimation or interpolation. If your data is sampled faster than needed during the experiment, you can decimate it without information loss. If your data is sampled more slowly than needed, there is a possibility that you miss important information about the dynamics at higher frequencies. Although you can resample the data at a higher rate, the resampled values occurring between measured samples do not represent measured information about your system. Instead of resampling, repeat the experiment using a higher sampling rate.

Tip You should decimate your data when it contains high-frequency noise outside the frequency range of the system dynamics.

Resampling takes into account how the data behaves between samples, which you specify when you import the data into the System Identification Tool GUI (zero-order or first-order hold).

You can resample data using the System Identification Tool GUI or the `resample` command. You can only resample time-domain data and at uniform time intervals.

Resampling Data Using the GUI

Use the System Identification Tool GUI to resample time-domain data. To specify additional options, such as the prefilter order, see “Resampling Data at the Command Line” on page 2-81.

The System Identification Tool GUI uses `idresamp` to interpolate or decimate the data. For more information about this command, type `help idresamp` at the MATLAB prompt.

To create a new data set by resampling the input and output signals:

- 1** Import time-domain data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14.
- 2** Drag the data set you want to resample to the **Working Data** area.
- 3** In the **Resampling factor** field, enter the factor by which to multiply the current sampling interval:
 - For decimation (fewer samples), enter a factor greater than 1 to increase the sampling interval by this factor.
 - For interpolation (more samples), enter a factor less than 1 to decrease the sampling interval by this factor.

Default = 1.

- 4** In the **Data name** field, type the name of the new data set. Choose a name that is unique in the Data Board.
- 5** Click **Insert** to add the new data set to the Data Board in System Identification Toolbox window.
- 6** Click **Close** to close the Resample dialog box.

Resampling Data at the Command Line

Use `resample` to decimate and interpolate time-domain `iddata` objects. You can specify the order of the antialiasing filter as an argument.

Note `resample` uses Signal Processing Toolbox command, when this toolbox is installed on your computer. If this toolbox is not installed, use `idresamp` instead. `idresamp` only lets you specify the filter order, whereas `resample` also lets you specify filter coefficients and the design parameters of the Kaiser window.

To create a new `iddata` object `datar` by resampling data, use the following syntax:

```
datar = resample(data,P,Q,filter_order)
```

In this case, P and Q are integers that specify the new sampling interval: the new sampling interval is Q/P times the original one. You can also specify the order of the resampling filter as a fourth argument `filter_order`, which is an integer (default is 10). For detailed information about `resample`, see the corresponding reference page.

For example, `resample(data,1,Q)` results in decimation with the sampling interval modified by a factor Q .

The next example shows how you can increase the sampling rate by a factor of 1.5 and compare the signals:

```
plot(u)
ur = resample(u,3,2);
plot(u,ur)
```

When Signal Processing Toolbox is not installed, using `resample` calls `idresamp` instead.

`idresamp` uses the following syntax:

```
datar = idresamp(data,R,filter_order)
```

In this case, $R=Q/P$, which means that data is interpolated by a factor P and then decimated by a factor Q . To learn more about `idresamp`, type `help idresamp`.

The data.`InterSample` property of the `iddata` object is taken into account during resampling. For more information, see “`iddata` Properties” on page 2-35.

Resampling Data Without Aliasing Effects


Typically, you decimate a signal to remove the high-frequency contributions that result from noise from the total energy. Ideally, you want to remove the energy contribution due to noise and preserve the energy density of the signal.

The command `resample` performs the decimation without aliasing effects. This command includes a factor of T to normalize the spectrum and preserve the energy density after decimation. For more information about spectrum normalization, see “Understanding Spectrum Normalization” on page 4-19.

If you use manual decimation instead of `resample`—by picking every fourth sample from the signal, for example—the energy contributions from higher frequencies are folded back into the lower frequencies. Because the total signal energy is preserved by this operation and this energy must now be squeezed into a smaller frequency range, the amplitude of the spectrum at each frequency increases. Thus, the energy density of the decimated signal is not constant.

The following example illustrates how `resample` avoids folding effects:

```
% Construct fourth-order MA-process
m0 = idpoly(1,[ ],[1 1 1 1]);
% Generate error signal
e = idinput(2000,'rgs');
e = iddata([],e,'Ts',1);
% Simulate the output using error signal
y = sim(m0,e);
% Estimate signal spectrum
g1 = spa(y);
% Estimate spectrum of modified signal including
% every fourth sample of the original signal.
% This command automatically sets Ts to 4.
g2 = spa(y(1:4:2000));
% Plot frequency response to view folding effects
ffplot(g1,g2)
% Estimate spectrum after prefiltering that does not
% introduce folding effects
g3 = spa(resample(y,1,4));
figure
ffplot(g1,g3)
```




Folding Effects with Manual Decimation

Use `resample` to decimate the signal before estimating the spectrum and plot the frequency response, as follows:

```
g3 = spa(resample(y,1,4));  
figure  
ffplot(g1,g3)
```

The following figure shows that the estimated spectrum of the resampled signal has the same amplitude as the original spectrum. Thus, there is no indication of folding effect when you use `resample` to eliminate aliasing.



No Folding Effects When Using `resample`

See Also

For general information about working with the System Identification Tool GUI, see Chapter 1, “Working with the System Identification Tool GUI”.

For a detailed discussion about handling disturbances, see the chapter on preprocessing data in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

Filtering Data

In this section...

“Supported Filters in System Identification Toolbox” on page 2-87

“Deciding to Prefilter Your Data” on page 2-87

“How to Filter Data Using the GUI” on page 2-88

“How to Filter Data at the Command Line” on page 2-91

“See Also” on page 2-94

Supported Filters in System Identification Toolbox

You can filter the input and output signals through a linear filter before estimating a model in the System Identification Tool GUI or at the command line. How you want to handle the noise in the system determines whether it is appropriate to prefilter the data.

The filter available in the System Identification Tool GUI is a fifth-order (passband) Butterworth filter. If you need to specify a custom filter, use the `idfilt` command.

Deciding to Prefilter Your Data

Prefiltering data can help remove high-frequency noise or low-frequency disturbances (drift). The latter application is an alternative to subtracting linear trends from the data, as described in “Detrending Data” on page 2-75.

In addition to minimizing noise, prefiltering lets you focus your model on specific frequency bands. The frequency range of interest often corresponds to a passband over the breakpoints on a Bode plot (see “Plotting Data” on page 2-62). For example, if you are modeling a plant for control-design applications, you might prefilter the data to specifically enhance frequencies around the desired closed-loop bandwidth.

Prefiltering the input and output data through the same filter does not change the input-output relationship for a linear system. However, prefiltering does change the noise characteristics and affects the estimated model of the system.

To get a reliable noise model, we recommend that you avoid prefiltering the data. Instead, set the `Focus` property of the estimation algorithm to `Simulation`. For more information, see the `Algorithm Properties` reference page.

Note When you prefilter during model estimation, the filtered data is used to only model the input-to-output dynamics. However, the disturbance model is calculated from the unfiltered data.

How to Filter Data Using the GUI

- “Filtering Time-Domain Data in the GUI” on page 2-88
- “Filtering Frequency-Domain or Frequency-Response Data in the GUI” on page 2-90

Filtering Time-Domain Data in the GUI

The System Identification Tool GUI lets you filter time-domain data by enhancing or selecting specific passbands.

The System Identification Tool GUI filters time-domain data using a fifth-order Butterworth filter.

To create a filtered data set:

- 1** Import time-domain data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14.
- 2** Drag the data set you want you want to filter to the **Working Data** area.
- 3** Select **Preprocess > Filter**. By default, this selection shows a periodogram of the input and output spectra (see the `etfe` reference page).

Note To display smoothed spectral estimates instead of the periodogram, select **Options > Spectral analysis**. This spectral estimate is computed using `spa` and your previous settings in the Spectral Model dialog box. To change these settings, select **Estimate > Spectral model** in the System Identification Tool GUI, and specify new model settings.

- 4** If your data contains multiple I/O channels, in the **Channel** menu, select the channel pair you want to view. Although you view only one channel pair at a time, the filter applies to all I/O channels in this data set.
- 5** Select the data of interest using one of the following ways:
 - Graphically — Draw a rectangle with the mouse on either the input-signal or the output-signal plot to select the desired frequency interval. Your selection is displayed on both plots regardless of the plot on which you draw the rectangle. The **Range** field is updated to match the selected region. If you need to clear your selection, right-click the plot.
 - By specifying the **Range** — Edit the beginning and the end frequency values. For example:
8.5 20.0 (rad/s).

Tip To change the frequency units from rad/s to Hz, select **Style > Frequency (Hz)**. To change the frequency units from Hz to rad/s, select **Style > Frequency (rad/s)**.

- 6** In the **Range is** list, select one of the following:
 - Pass band — Allows data in the selected frequency range.
 - Stop band — Removes data in the selected frequency range.
- 7** Click **Filter** to preview the filtered results. If you are satisfied, go to step 8. Otherwise, return to step 5.
- 8** In the **Data name** field, enter the name of the data set containing the selected data.

9 Click **Insert** to save the selection as a new data set and add it to the Data Board.

10 To select another range, repeat steps 5 to 9.

Filtering Frequency-Domain or Frequency-Response Data in the GUI

The System Identification Tool GUI lets you filter frequency-domain or frequency-response data by enhancing or selecting specific passbands.

For frequency-domain and frequency-response data, *filtering* is equivalent to selecting specific data ranges.

To select a range of data, which is the equivalent of filtering frequency-domain and frequency-response data:

- 1** Import data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14.
- 2** Drag the data set you want you want to filter to the **Working Data** area.
- 3** Select **Preprocess > Select range**. This selection displays one of the following plots:
 - Frequency-domain data — Plot shows the absolute of the squares of the input and output spectra.
 - Frequency-response data — Top axes show the frequency response magnitude equivalent to the ratio of the output to the input), and the bottom axes show the ratio of the input signal to itself, which has the value of 1 at all frequencies.
- 4** If your data contains multiple I/O channels, in the **Channel** menu, select the channel pair you want to view. Although you view only one channel pair at a time, the filter applies to all I/O channels in this data set.
- 5** Select the data of interest using one of the following ways:
 - Graphically — Draw a rectangle with the mouse on either the input-signal or the output-signal plot to select the desired frequency interval. Your selection is displayed on both plots regardless of the plot

on which you draw the rectangle. The **Range** field is updated to match the selected region.

If you need to clear your selection, right-click the plot.

- By specifying the **Range** — Edit the beginning and the end frequency values. For example:

8.5 20.0 (rad/s).

Tip If you need to change the frequency units from rad/s to Hz, select **Style > Frequency (Hz)**. To change the frequency units from Hz to rad/s, select **Style > Frequency (rad/s)**.

6 In the **Range is** list, select one of the following:

- Pass band — Allows data in the selected frequency range.
- Stop band — Removes data in the selected frequency range.

7 In the **Data name** field, enter the name of the data set containing the selected data.

8 Click **Insert**. This action saves the selection as a new data set and adds it to the Data Board.

9 To select another range, repeat steps 5 to 8.

How to Filter Data at the Command Line

- “Simple Passband Filter” on page 2-91
- “Defining a Custom Filter” on page 2-92
- “Causal and Noncausal Filters” on page 2-93

Simple Passband Filter

Use `idfilt` to apply passband and other custom filters to a time-domain or a frequency-domain `iddata` object.

In general, you can specify any custom filter. Use this syntax to filter an `iddata` object data using the filter called `filter`:

```
fdata = idfilt(data,filter)
```

In the simplest case, you can specify a passband filter for time-domain data using the following syntax:

```
fdata = idfilt(data,[w1 wh])
```

In this case, `w1` and `wh` represent the low and high frequencies of the passband, respectively.

You can specify several passbands, as follows:

```
filter=[ [w1l,w1h]; [ w2l,w2h]; ...; [wnl,wnh] ]
```

The filter is an `n`-by-2 matrix, where each row defines a passband in radians per second.

To define a stopband between `ws1` and `ws2`, use

```
filter = [0 ws1; ws2 Nyqf]
```

where, `Nyqf` is the Nyquist frequency.

For time-domain data, the passband filtering is cascaded Butterworth filters of specified order. The default filter order is 5. The Butterworth filter is the same as `butter` in Signal Processing Toolbox. For frequency-domain data, select the indicated portions of the data to perform passband filtering.

Defining a Custom Filter

Use `idfilt` to apply passband and other custom filters to a time-domain or a frequency-domain `iddata` object.

In general, you can specify any custom filter. Use this syntax to filter an `iddata` object data using the filter called `filter`:

```
fdata = idfilt(data,filter)
```

You can define a general single-input/single-output (SISO) system for filtering time-domain or frequency-domain data. For frequency-domain only, you can specify the (nonparametric) frequency response of the filter.

You use this syntax to filter an `iddata` object data using a custom filter specified by `filter`:

```
fdata = idfilt(data,filter)
```

`filter` can be also any of the following:

```
filter = idm  
filter = {num,den}  
filter = {A,B,C,D}
```

`idm` is a SISO `idmodel` or `LTI` object. For more information about `LTI` objects, see the Control System Toolbox documentation.

`{num,den}` defines the filter as a transfer command as a cell array of numerator and denominator filter coefficients.

`{A,B,C,D}` is a cell array of SISO state-space matrices.

Specifically for frequency-domain data, you specify the frequency response of the filter:

```
filter = Wf
```

Here, `Wf` is a vector of real or complex values that define the filter frequency response, where the inputs and outputs of `data` at frequency `data.Frequency(kf)` are multiplied by `Wf(kf)`. `Wf` is a column vector with the length equal to the number of frequencies in `data`.

When `data` contains several experiments, `Wf` is a cell array with the length equal to the number of experiments in `data`.

Causal and Noncausal Filters

For time-domain data, the filtering is causal by default. Causal filters typically introduce a phase shift in the results. To use a noncausal zero-phase

filter (corresponding to `filtfilt` in Signal Processing Toolbox), specify a third argument in `idfilt`:

```
fdata = idfilt(data,filter,'noncausal')
```

For frequency-domain data, the signals are multiplied by the frequency response of the filter. With the filters defined as passband filters, this calculation gives ideal, zero-phase filtering (“brick wall filters”). Frequencies that have been assigned zero weight by the filter (outside the passband or via frequency response) are removed.

When you apply `idfilt` to an `idfrd` data object, the data is first converted to a frequency-domain `iddata` object (see “Transforming Between Frequency-Domain and Frequency-Response Data” on page 2-112). The result is an `iddata` object.

See Also

To learn how to filter data during linear model estimation instead, you can set the `Focus` property of the estimation algorithm to `Filter` and specify the filter characteristics. For more information about model properties, see the `Algorithm Properties` reference page.

For more information about prefiltering data, see the chapter on preprocessing data in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

For practical examples of prefiltering data, see the section on posttreatment of data in *Modeling of Dynamic Systems*, by Lennart Ljung and Torkel Glad, Prentice Hall PTR, 1994.

Selecting Data

In this section...
“Why Select Subsets of Data” on page 2-95
“Selecting Data Using the GUI” on page 2-96
“Selecting Data at the Command Line” on page 2-98

Why Select Subsets of Data

You can also use data selection to create independent data sets for estimation and validation.

You can use data selection as a way to clean the data and exclude parts with noisy or missing information. For example, when your data contains missing data, outliers, level changes, and disturbances, you can select one or more portions of the data that are suitable for identification and exclude the rest.

If you only have one data set and you want to estimate linear models, you should split the data into two portions to create two independent data sets for estimation and validation, respectively. Splitting the data is selecting parts of the data set and saving each part independently.

You can merge several data segments into a single multiexperiment data set and identify an average model. For more information, see “Importing Data into the GUI” on page 2-14 or “Creating iddata Objects” on page 2-32.

Note Subsets of the data set must contain enough samples to adequately represent the system and the inputs must provide suitable excitation to the system.

Selecting portions of frequency-domain data is equivalent to filtering the data. For more information about filtering, see “Filtering Data” on page 2-87.

Selecting Data Using the GUI

- “Ways to Select Data in the GUI” on page 2-96
- “Selecting a Range for Time-Domain Data” on page 2-96
- “Selecting a Range of Frequency-Domain Data” on page 2-98

Ways to Select Data in the GUI

You can use System Identification Tool GUI to select ranges of data on a time-domain or frequency-domain plot. Selecting data in the frequency domain is equivalent to passband-filtering the data.

After you select portions of the data, you can specify to use one data segment for estimating models and use the other data segment for validating models. For more information, see “Specifying Estimation and Validation Data” on page 1-11.

Note Selecting **<--Preprocess > Quick start** performs the following actions simultaneously:

- Remove the mean value from each channels.
 - Split the data into two part.
 - Specify the first part as estimation data (or **Working Data**)
 - Specify the second part as **Validation Data**.
-

Selecting a Range for Time-Domain Data

You can select a range of data values on a time plot and save it as a new data set in the System Identification Tool GUI.

Note You cannot extract data in data sets containing multiple experiments. For more information about multiexperiment data, see “Creating Multiexperiment Data Sets in the GUI” on page 2-23.

To extract a subset of time-domain data and save it as a new data set:

- 1** Import time-domain data into the System Identification Tool GUI, as described in “Importing Data into the GUI” on page 2-14.
- 2** Drag the data set you want to subset to the **Working Data** area.
- 3** If your data contains multiple I/O channels, in the **Channel** menu, select the channel pair you want to view. The upper plot corresponds to the input signal, and the lower plot corresponds to the output signal.

Although you view only one I/O channel pair at a time, your data selection is applied to all channels in this data set.

- 4** Select the data of interest in either of the following ways:

- Graphically — Draw a rectangle on either the input-signal or the output-signal plot with the mouse to select the desired time interval. Your selection appears on both plots regardless of the plot on which you draw the rectangle. The **Time span** and **Samples** fields are updated to match the selected region.
- By specifying the **Time span** — Edit the beginning and the end times in seconds. The **Samples** field is updated to match the selected region. For example:

28.5 56.8

- By specifying the **Samples** range — Edit the beginning and the end indices of the sample range. The **Time span** field is updated to match the selected region. For example:

342 654

To clear your selection, click **Revert**.

- 5** In the **Data name** field, enter the name of the data set containing the selected data.
- 6** Click **Insert**. This action saves the selection as a new data set and adds it to the Data Board.
- 7** To select another range, repeat steps 4 to 6.

Selecting a Range of Frequency-Domain Data

Selecting a range of values in frequency domain is equivalent to filtering the data in System Identification Toolbox. For more information about data filtering, see “Filtering Frequency-Domain or Frequency-Response Data in the GUI” on page 2-90.

Selecting Data at the Command Line

Selecting ranges of data values is equivalent to subreferencing the data.

For more information about subreferencing time-domain and frequency-domain data stored in `iddata` objects, see “Subreferencing `iddata` Objects” on page 2-40.

For more information about subreferencing frequency-response data, see “Subreferencing `idfrd` Objects” on page 2-55.

Creating Data Using Simulation

In this section...

“Commands for Generating and Simulating Data” on page 2-99

“Example – Creating Data with Periodic Inputs” on page 2-100

“Example – Simulating Model Output at the Command Line” on page 2-101

“Simulating Data Using Other Mathworks Products” on page 2-102

Commands for Generating and Simulating Data

System Identification Toolbox lets you generate input data and simulate output data using a model structure.

Simulating output data requires that you have a parametric model. For more information about commands for constructing models, see “Command for Creating Model Objects” on page 3-4.

To generate input data, use `idinput` to construct a signal with the desired characteristics, such as a random Gaussian or binary signal or a sinusoid. `idinput` returns a matrix of input values.

The following table lists the commands you can use to simulate output data. For more information about these commands, see the references page.

Commands for Generating and Simulating Data

Command	Description	Example
<code>iddata</code>	Constructs an <code>iddata</code> object with input channels only.	To construct input data <code>data</code> , use the following command: <pre>data = iddata([],[u v])</pre> <p><code>u</code> is the input data, and <code>v</code> is white noise.</p>

Commands for Generating and Simulating Data (Continued)

Command	Description	Example
idinput	Constructs a signal with the desired characteristics, such as a random Gaussian or binary signal or a sinusoid, and returns a matrix of input values.	<pre>u = iddata([],... idinput(400,'rbs',[0 0.3]));</pre>
sim	Simulates response data based on existing linear or nonlinear parametric model in the MATLAB Workspace browser.	<p>To simulate the model output y for a given input, use the following command:</p> <pre>y = sim(m,data)</pre> <p>m is the model object name, and $data$ is input data matrix or <code>iddata</code> object.</p>

Example – Creating Data with Periodic Inputs

- 1 Create a periodic input for two inputs and consisting of five periods, where each period is 300 samples.

```
per_u = idinput([300 2 5])
```

- 2 Create an `iddata` object using the periodic input and leaving the output empty.

```
u = iddata([],per_u,'Period',...
           [300; 300]);
```

You can use the periodic input to simulate the output, and the use `etfe` to compute the estimated response of the model.

```
% Construct polynomial model
m0 =idpoly([1 -1.5 0.7],[0 1 0.5]);
% Construct random binary input
u = idinput([10 1 150],'rbs');
```

```

% Construct input data and noise
u = iddata([],u,'Period',10);
e = iddata([],randn(1500,1));
% Simulate model output with noise
y = sim(m0,[u e])
% Estimate frequency response
g = etfe([y u])
% Generate Bode plot
bode(g,'x',m0)

```

For periodic input, `etfe` honors the period and computes the frequency response using an appropriate frequency grid. In this case, the Bode plot shows a good fit at the five excited frequencies.

Example – Simulating Model Output at the Command Line

This example demonstrates how you can create input data and a model, and then use the data and the model to simulate output data. You create the ARMAX model and simulate output data with random binary input `u`.

- 1 Load the three-input and one-output sample data.

```
load iddata8
```

- 2 Construct an ARMAX model, using the following commands:

```

A = [1 -1.2 0.7];
B(1,:) = [0 1 0.5 0.1]; % first input
B(2,:) = [0 1.5 -0.5 0]; % second input
B(3,:) = [0 -0.1 0.5 -0.1]; % third input
C = [1 0 0 0 0];
Ts = 1;
m = idpoly(A,B,C,'Ts',1);

```

In this example, the leading zeros in the B matrix indicate the input delay (nk), which is 1 for each input channel. The trailing zero in `B(2,:)` makes the number of coefficients equal for all channels.

3 Construct pseudorandom binary data for input to the simulation.

```
u = idinput([200,3], 'prbs');
```

4 Simulate the model output.

```
sim(m,u)
```

5 Compare model output to measured data to see how well the model captures the underlying dynamics.

```
compare(z8,m)
```

Simulating Data Using Other Mathworks Products

You can also simulate data using Simulink and Signal Processing Toolbox. Data simulated outside System Identification Toolbox must be in the MATLAB Workspace browser to be available to System Identification Toolbox environment. For more information about simulating models using Simulink, see “Using Models with Simulink” on page 13-8.

Transforming Between Time- and Frequency-Domain Data

In this section...

- “Transforming Data Domain in the GUI” on page 2-103
- “Transforming Data Domain at the Command Line” on page 2-110

Transforming Data Domain in the GUI

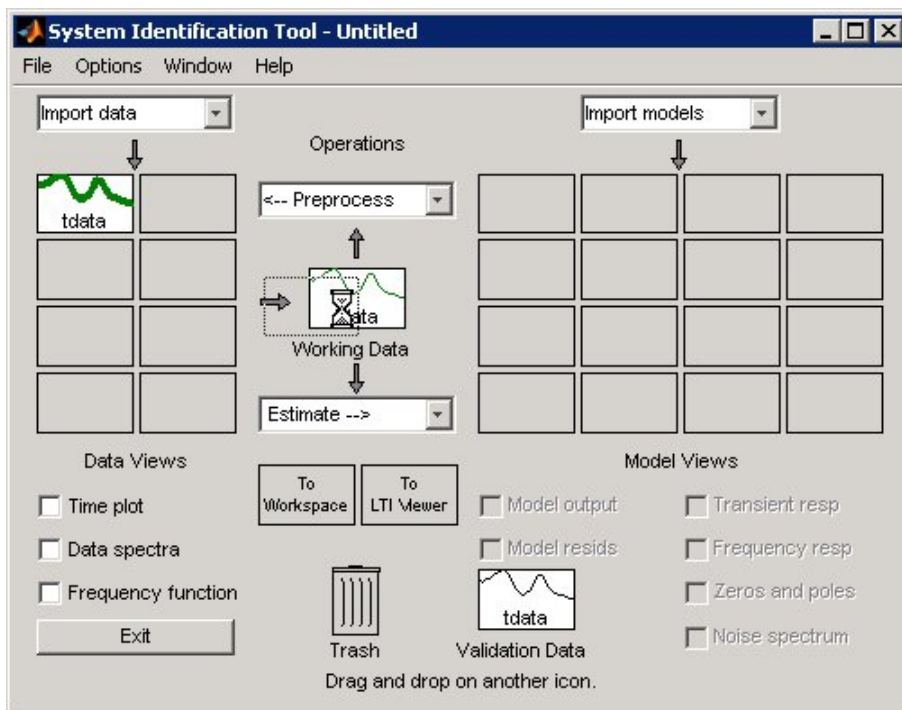
- “Transforming Time-Domain Data” on page 2-103
- “Transforming Frequency-Domain Data” on page 2-107
- “Transforming Frequency-Response Data” on page 2-108
- “See Also” on page 2-110

Transforming Time-Domain Data

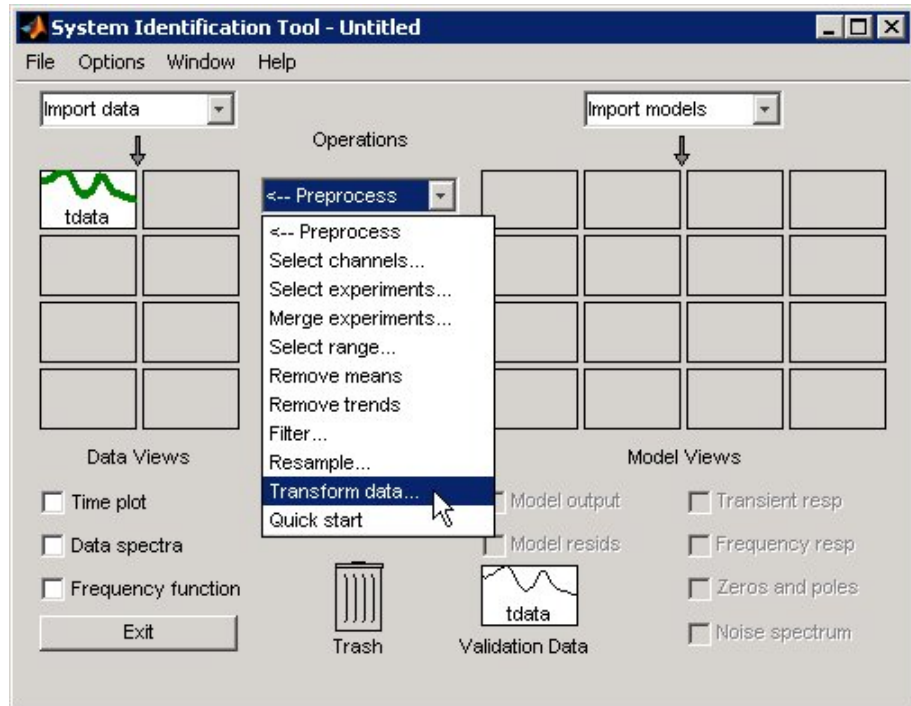
In the System Identification Tool GUI, time-domain data has an icon with a white background. You can transform time-domain data to frequency-domain or frequency-response data. The frequency values of the resulting frequency vector range from 0 to the Nyquist frequency $f_S = \pi/T_s$, where T_s is the sampling interval.

Transforming from time-domain to frequency-response data is equivalent to creating a nonparametric model of the data using the `spafdr` method.

- 1 In the System Identification Tool GUI, drag the icon of the data you want to transform to the **Working Data** rectangle, as shown in the following figure.




- 2 In the **Operations** area, select **<--Preprocess > Transform data** in the drop-down menu to open the Transform Data dialog box.



3 In the **Transform to** drop-down list, select one of the following:

- **Frequency Function** — Create a new `idfrd` object using the `spafdr` method. Go to step 4.



- **Frequency Domain Data** — Create a new `iddata` object using the `fft` method. Go to step 6.

4 In the **Frequency Spacing** list, select the spacing of the frequencies at which the frequency command is estimated:

- **linear** — Uniform spacing of frequency values between the endpoints.
- **logarithmic** — Base-10 logarithmic spacing of frequency values between the endpoints.

5 In the **Number of Frequencies** field, enter the number of frequency values.

6 In the **Name of new data** field, type the name of the new data set. This name must be unique in the Data Board.

7 Click **Transform** to add the new data set to the Data Board in the System Identification Tool GUI.

8 Click **Close** to close the Transform Data dialog box.

Transforming Frequency-Domain Data

In the System Identification Tool GUI, frequency-domain data has an icon with a green background. You can transform frequency-domain data to time-domain or frequency-response (frequency-command) data.

Transforming from time-domain or frequency-domain data to frequency-response data is equivalent to creating a nonparametric model of the data using the `spafdr` method.

- 1** In the System Identification Tool GUI, drag the icon of the data you want to transform to the **Working Data** rectangle.
- 2** Select **<--Preprocess > Transform data**.
- 3** In the **Transform to** list, select one of the following:
 - **Frequency Function** — Create a new `idfrd` object using the `spafdr` method. Go to step 4.
 - **Time Domain Data** — Create a new `iddata` object using the `ifft` (inverse fast Fourier transform) method. Go to step 6.
- 4** In the **Frequency Spacing** list, select the spacing of the frequencies at which the frequency command is estimated:
 - **linear** — Uniform spacing of frequency values between the endpoints.
 - **logarithmic** — Base-10 logarithmic spacing of frequency values between the endpoints.
- 5** In the **Number of Frequencies** field, enter the number of frequency values.
- 6** In the **Name of new data** field, type the name of the new data set. This name must be unique in the Data Board.
- 7** Click **Transform** to add the new data set to the Data Board in the System Identification Tool GUI.
- 8** Click **Close** to close the Transform Data dialog box.

Transforming Frequency-Response Data

In the System Identification Tool GUI, frequency-response data has an icon with a yellow background. You can transform frequency-response data to frequency-domain data (iddata object) or to frequency-response data with a different frequency resolution.

When you select to transform single-input/single-output (SISO) frequency-response data to frequency-domain data, System Identification Toolbox creates outputs that equal the frequency responses, and inputs equal to 1. Therefore, the ratio between the Fourier transform of the output and the Fourier transform of the input is equal to the system frequency response.

For the multiple-input case, System Identification Toolbox transforms the frequency-response data to frequency-domain data as if each input contributes independently to the entire output of the system and then combines information. For example, if a system has three inputs, u_1 , u_2 , and u_3 and two frequency samples, the input matrix is set to:

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}$$

In general, for nu inputs and ns samples (the number of frequencies), the input matrix has nu columns and $(ns \cdot nu)$ rows.

Note To make the response from each input a separate experiment at the command line, see “Transforming Between Frequency-Domain and Frequency-Response Data” on page 2-112.

When you transform frequency-response data by changing its frequency resolution, you can modify the number of frequency values by changing between linear or logarithmic spacing. You might specify variable frequency spacing to increase the number of data points near the system’s resonance

frequencies, and also make the frequency vector coarser in the region outside the system dynamics. Typically, high-frequency noise dominates away from frequencies where interesting system dynamics occur. The System Identification Tool GUI lets you specify logarithmic frequency spacing, which results in a variable frequency resolution.

Note The `spafdr` command lets you specify any variable frequency resolution.

- 1** In the System Identification Tool GUI, drag the icon of the data you want to transform to the **Working Data** rectangle.
- 2** Select **<--Preprocess > Transform data**.
- 3** In the **Transform to** list, select one of the following:
 - **Frequency Domain Data** — Create a new `iddata` object. Go to step 6.
 - **Frequency Function** — Create a new `idfrd` object with different resolution (number and spacing of frequencies) using the `spafdr` method. Go to step 4.
- 4** In the **Frequency Spacing** list, select the spacing of the frequencies at which the frequency command is estimated:
 - **linear** — Uniform spacing of frequency values between the endpoints.
 - **logarithmic** — Base-10 logarithmic spacing of frequency values between the endpoints.
- 5** In the **Number of Frequencies** field, enter the number of frequency values.
- 6** In the **Name of new data** field, type the name of the new data set. This name must be unique in the Data Board.
- 7** Click **Transform** to add the new data set to the Data Board in the System Identification Tool GUI.
- 8** Click **Close** to close the Transform Data dialog box.

See Also

For a description of time-domain, frequency-domain, and frequency-response data, see “Data Requirements” on page 2-6.

To learn how to transform data at the command line instead of the GUI, see “Transforming Data Domain at the Command Line” on page 2-110.

For more information about working with the System Identification Tool GUI, see Chapter 1, “Working with the System Identification Tool GUI”.

Transforming Data Domain at the Command Line

- “Supported Data Transformations” on page 2-110
- “Transforming Between Time and Frequency Domain” on page 2-111
- “Transforming Between Frequency-Domain and Frequency-Response Data” on page 2-112
- “See Also” on page 2-114

Supported Data Transformations

The following table uses a matrix to show the different ways you can transform the data of one format (row) to another format (column). If the transformation is supported for a given row and column combination, the method used by the software is listed in the cell at their intersection.

Original Data Format	To Time Domain (iddata object)	To Frequency Domain (iddata object)	To Frequency Function (idfrd object)
Time Domain (iddata object)	No.	Yes, using fft.	Yes, using etfe, spa, or spafdr.

Original Data Format	To Time Domain (iddata object)	To Frequency Domain (iddata object)	To Frequency Function (idfdr object)
Frequency Domain (iddata object)	Yes, using <code>ifft</code> .	No.	Yes, using <code>etfe</code> , <code>spa</code> , or <code>spafdr</code> .
Frequency Function (idfdr object)	No.	Yes. Calculation creates frequency-domain iddata object that has the same ratio between output and input as the original idfdr object.	Yes. Calculates a frequency command with different resolution (number and spacing of frequencies) using <code>spafdr</code> .

Transforming Between Time and Frequency Domain

The `iddata` object stores time-domain or frequency-domain data. The following table summarizes the commands for transforming data between time and frequency domains.

Command	Description	Syntax Example
fft	Transforms time-domain data to the frequency domain. You can specify N, the number of frequency values.	To transform time-domain iddata object t_data to frequency-domain iddata object f_data with N frequency points, use: f_data = fft(t_data,N)
ifft	Transforms frequency-domain data to the time domain. Frequencies are linear and equally spaced.	To transform frequency-domain iddata object f_data to time-domain iddata object t_data, use: t_data = ifft(f_data)

Transforming Between Frequency-Domain and Frequency-Response Data

You can transform frequency-response data to frequency-domain data (iddata object). The idfrd object represents complex frequency-response of the system at different frequencies. For a description of this type of data, see “Representing Frequency-Response Data” on page 2-10.

When you select to transform single-input/single-output (SISO) frequency-response data to frequency-domain data, System Identification Toolbox creates outputs that equal the frequency responses, and inputs equal to 1. Therefore, the ratio between the Fourier transform of the output and the Fourier transform of the input is equal to the system frequency response.

For information about changing the frequency resolution of frequency-response data to a new constant or variable (frequency-dependent) resolution, see the spafdr reference page. You might use this advanced feature to increase the number of data points near the system’s resonance frequencies and also make the frequency vector coarser in the region outside the system dynamics.

Typically, high-frequency noise dominates away from frequencies where interesting system dynamics occur.

Note You cannot transform an `idfrd` object to a time-domain `iddata` object.

To transform an `idfrd` object with the name `idfrdobj` to a frequency-domain `iddata` object, use the following syntax:

```
dataf = iddata(idfrdobj)
```

The resulting frequency-domain `iddata` object contains values at the same frequencies as the original `idfrd` object.

For the multiple-input case, System Identification Toolbox represents frequency-response data as if each input contributes independently to the entire output of the system and then combines information. For example, if a system has three inputs, `u1`, `u2`, and `u3` and two frequency samples, the input matrix is set to:

$$\begin{bmatrix} 1 & 0 & 0 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \\ 0 & 0 & 1 \end{bmatrix}$$

In general, for `nu` inputs and `ns` samples, the input matrix has `nu` columns and `(ns · nu)` rows.

If you have `ny` outputs, the transformation operation produces an output matrix has `ny` columns and `(ns · nu)` rows using the values in the complex frequency response $G(iw)$ matrix (`ny-by-nu-by-ns`). In this example, `y1` is determined by unfolding $G(1,1,:)$, $G(1,2,:)$, and $G(1,3,:)$ into three column vectors and vertically concatenating these vectors into a single column. Similarly, `y2` is determined by unfolding $G(2,1,:)$, $G(2,2,:)$, and $G(2,3,:)$ into three column vectors and vertically concatenating these vectors.

If you are working with multiple inputs, you also have the option of storing the contribution by each input as an independent experiment in a multiexperiment data set. To transform an `idfrdobj` object with the name `idfrdobj` to a multiexperiment data set `datf`, where each experiment corresponds to each of the inputs in `idfrdobj`

```
datf = iddata(idfrdobj, 'me')
```

In this example, the additional argument `'me'` specifies that multiple experiments are created.

By default, transformation from frequency-response to frequency-domain data strips away frequencies where the response is `inf` or `NaN`. To preserve the entire frequency vector, use `datf = iddata(idfrdobj, 'inf')`. For more information, type `help idfrd/iddata`.

See Also

Transforming from time-domain or frequency-domain data to frequency-response data is equivalent to creating a nonparametric model of the data. For more information about creating nonparametric spectral models, see the `etfe`, `spa`, and `spafdr` reference page.

Handling Complex-Valued Data

In this section...

“Supported Operations for Complex Data” on page 2-115

“Manipulating Complex `iddata` Signals” on page 2-115

Supported Operations for Complex Data

System Identification Toolbox estimation routines support complex data. For example, the following estimation commands estimate complex models from complex data: `ar`, `armax`, `arx`, `bj`, `covf`, `ivar`, `iv4`, `oe`, `pem`, `spa`, and `n4sid`

Model transformation routines, such as `freqresp`, `zpkdata`, work for complex-valued models. However, they do not provide pole-zero confidence regions. For complex models, the parameter variance-covariance information refers to the complex-valued parameters and the accuracy of the real and imaginary is not computed separately.

The display commands `compare` and `plot` also work with complex-valued data and models, but only show the absolute values of the signals. To plot the real and imaginary parts of the data separately, use `plot(real(data))` and `plot(imag(data))`, respectively.

Manipulating Complex `iddata` Signals

If the `iddata` object `data` contains complex values, you can use the following commands to process the complex data and create a new `iddata` object:

Function	Description
<code>abs(data)</code>	Absolute value of complex signals in <code>iddata</code> object.
<code>angle(data)</code>	Phase angle (in radians) of each complex signals in <code>iddata</code> object.

Function	Description
<code>complex(data)</code>	For time-domain data, this command makes the <code>iddata</code> object complex—even when the imaginary parts are zero. For frequency-domain data that only stores the values for nonnegative frequencies, such that <code>realdata(data)=1</code> , it adds signal values for negative frequencies using complex conjugation.
<code>imag(data)</code>	Selects the imaginary parts of each signal in <code>iddata</code> object.
<code>isreal(data)</code>	1 when <code>data</code> (time-domain or frequency-domain) contains only real input and output signals, and returns 0 when <code>data</code> (time-domain or frequency-domain) contains complex signals.
<code>real(data)</code>	Real part of complex signals in <code>iddata</code> object.
<code>realdata(data)</code>	Returns a value of 1 when <code>data</code> is a real-valued, time-domain signal, and returns 0 otherwise.

For example, suppose that you create a frequency-domain `iddata` object `Datf` by applying `fft` to a real-valued time-domain signal to take the Fourier transform of the signal. The following is true for `Datf`:

```
isreal(Datf) = 0  
realdata(Datf) = 1
```

Creating and Manipulating Models

Creating Model Objects (p. 3-2)

Overview of model objects and summary of commands for constructing and estimating models, accessing model properties, and concatenating and merging models.

Subreferencing Models (p. 3-20)

Creating models with subsets of inputs and outputs from multivariable models at the command line.

Concatenating Model Objects (p. 3-25)

Horizontal and vertical concatenation of model objects at the command line.

Merging Model Objects (p. 3-29)

How to merge models to obtain a single model with parameters that are statistically weighed means of the parameters of the individual models.

Creating Model Objects

In this section...

“Definition of a Model Object” on page 3-2

“Basic Object-Oriented Concepts” on page 3-3

“Command for Creating Model Objects” on page 3-4

“Commands for Model Estimation” on page 3-5

“Example – Estimating Models Using `armax`” on page 3-7

“Model Properties” on page 3-13

“See Also” on page 3-19

Definition of a Model Object

When you estimate a model in System Identification Toolbox, all information about this model is stored in a model object. *Model objects* are entities that store information about a model, including the mathematical form of a model, names of input and output channels, units, names and values of estimated parameters, parameter uncertainties, algorithm specifications, and estimation information.

System Identification Toolbox provides different model objects to represent different model types. When you work at the command line or writing M-file scripts, you create and operate on model objects directly. For a tutorial on estimating models at the command line, see “Tutorial: Estimating Linear Models Using the Command Line”.

If you work in the System Identification Tool GUI, you create and operate on model objects graphically using menu options and fields. For a tutorial on estimating models using the GUI, see “Tutorial: Estimating Linear Models Using the GUI”.

Whenever you estimate a linear model, the resulting model object includes two components: the dynamic model G for measured data and the noise model H for unmeasured disturbances. You can extract the dynamic and noise models from the estimated model using subreferencing, as described in “Subreferencing Models” on page 3-20.

Basic Object-Oriented Concepts

There are four basic concepts you need to know about when working with objects in System Identification Toolbox. These concepts include classes, methods, constructors, and properties.

Model objects are based on model classes. Each *class* is a blueprint that defines the following information about your model:

- How the object stores data
- Which operations you can perform on the object

System Identification Toolbox includes nine different classes for representing models. For example, `idpoly` represents linear black-box polynomial models, and `idss` represents linear state-space models. For a complete list of available model objects, see “Command for Creating Model Objects” on page 3-4.

The way a model object stores information is defined by the *properties* of the corresponding class. For example, the `idpoly` model object has a property called `InputName` for storing one or more input channel names. The set of available properties differs for different model objects.

The allowed operations on an object, called *methods*, are also defined by the corresponding class. In System Identification Toolbox, some methods have the same name but apply to multiple model objects. For example, the method `bode` create a bode plot for all linear model types. However, other methods are unique to a specific model object. For example, the estimation method `n4sid` is unique to the state-space model object `idss`.

Every class has a method for creating objects based on this class, called a *constructor*. Using a constructor creates a specific instance of the corresponding class and is called *instantiating the object*. In System Identification Toolbox, the constructor name is the same as the class name. For example, `idpoly` is both the name of the class representing linear black-box polynomial models and the constructor for instantiating the model object with specific property values.

You use model constructors to create a model object by specifying all required model properties explicitly. You use the constructed model for simulation or as an initial guess for iterative estimation. In most cases, you use the

estimation commands instead of constructors. These estimation commands both construct the model object and estimate the model parameters. For an example, see “Example – Estimating Models Using `armax`” on page 3-7.

Command for Creating Model Objects

The following table summarizes the model classes available in System Identification Toolbox for representing various types of models. This table also specifies whether a specific model type supports single or multiple outputs. The name of the object is the same as the name of the constructor for that object.

For information about how to both construct and estimate models with a single command, see “Commands for Model Estimation” on page 3-5. After estimating models in System Identification Toolbox, you can recognize these model objects in the MATLAB Workspace browser by their class names.

Summary of Model Classes

Model Class	Model Type	Single or Multiple Outputs?
<code>idarx</code>	Represents parametric multiple-output ARX models. Also represents nonparametric transient-response models.	Single- or multiple-output models.
<code>idfrd</code>	Represents nonparametric frequency-response model.	Single- or multiple-output models.
<code>idproc</code>	Represents continuous-time, low-order process models.	Single-output models only.
<code>idpoly</code>	Represents linear, black-box polynomial models: <ul style="list-style-type: none"> • ARX • ARMAX • Output-Error • Box-Jenkins 	Single-output models only.

Summary of Model Classes (Continued)

Model Class	Model Type	Single or Multiple Outputs?
idss	Represents linear, state-space models.	Single- or multiple-output models.
idgrey	Represents linear state-space model (grey-box models) in terms of your own variables and parameters. You write an M-file that translates user parameters to state-space matrices.	Single- and multiple-output models.
idnlgrey	Represents nonlinear, grey-box models. You write an M-file or MEX-file to represent the set of first-order differential or difference equations.	Supports single- and multiple-output models.
idnlarx	Represents nonlinear ARX models, which define the predicted output as a nonlinear command of past inputs and outputs.	Single- or multiple-output models. Does not support time series.
idnlhw	Represents Hammerstein-Wiener models, which include a linear dynamic system with nonlinear static transformations of inputs and outputs.	Single- or multiple-output models. Does not support time series.

Commands for Model Estimation

The quickest way to both construct a model object and estimate the model in System Identification Toolbox is to use estimation commands. This approach differs from the standard object-oriented approach, where you first use a constructor to instantiate the object, and then use an estimation method to estimate the model.

Note The standard object-oriented approach does apply to grey-box models, where you must define the model structure before estimating the model parameters.

The estimator `pem` corresponds to the iterative prediction-error method.

For ARMAX, Box-Jenkins, and Output-Error Models—which can only be estimated using the iterative prediction-error method—use the `armax`, `bj`, and `oe` estimation commands, respectively. These commands are versions of `pem` with simplified syntax for these specific model structures.

The following table summarizes System Identification Toolbox estimation commands. For detailed information about using each estimation command, see the corresponding reference page.

Construction and Estimation Commands

Model Type	Model Estimation Commands
Continuous-time low-order process models	<code>pem</code>
Linear black-box polynomial models: <ul style="list-style-type: none"> • ARX • ARMAX • Box-Jenkins (BJ) • Output-Error (OE) 	<code>armax</code> (ARMAX only) <code>arx</code> (ARX only) <code>bj</code> (BJ only) <code>iv4</code> (ARX only) <code>oe</code> (OE only) <code>pem</code> (for all models)
Linear state-space models	<code>n4sid</code> <code>pem</code>
Linear time-series models	<code>ar</code> <code>arx</code> (for multiple outputs) <code>ivar</code>

Construction and Estimation Commands (Continued)

Model Type	Model Estimation Commands
Nonlinear ARX	nlarx
Hammerstein-Wiener	nlhw

Example – Estimating Models Using armax

In System Identification Toolbox, you can use estimation commands to both construct a model object and estimate the model parameters. In this example, you estimate a linear, polynomial model with an ARMAX structure for a three-input and single-output (MISO) system using the iterative estimation method `armax`. For a summary of all available estimation commands in System Identification Toolbox, see “Commands for Model Estimation” on page 3-5.

- 1 Load a sample data set `z8` with three inputs and one output, measured at 1-second intervals and containing 500 data samples:

```
load iddata8
```

- 2 Use `armax` to both construct the `idpoly` model object, and estimate the parameters:

$$A(q)y(t) = \sum_{i=1}^{nu} B_i(q)u_i(t - nk_i) + C(q)e(t)$$

Typically you try different model orders and compare results, ultimately choosing the simplest model that best describes the system dynamics. The following command specifies the estimation data set, `z8`, and the orders of the A , B , and C polynomials as `na`, `nb`, and `nc`, respectively. `nk` of `[0 0 0]` specifies that there is no input delay for all three input channels.

```
m_armax=armax(z8,'na',4,...
               'nb',[3 2 3],...
               'nc',4,...
               'nk',[0 0 0],...
               'focus', 'simulation',...
               'tolerance',1e-5,...
               'maxiter',50);
```

`covariance`, `focus`, `tolerance`, and `maxiter` are optional arguments specify additional information about the computation. `focus` specifies whether the model is optimized for simulation or prediction applications, `tolerance` and `maxiter` specify when to stop estimation. For more information about these properties, see the `algorithm` properties reference page.

`armax` is a version of `pem` with simplified syntax for the ARMAX model structure. The `armax` method both constructs the `idpoly` model object and estimates its parameters.

Tip Instead of specifying model orders and delays as individual property-value pairs, you can use the equivalent shorthand notation that includes all of the order integers in a single vector, as follows:

```
m_armax=armax(z8,[4 3 2 3 4 0 0 0],...
               'focus', 'simulation',...
               'tolerance',1e-5,...
               'maxiter',50);
```

- 3** To view information about the resulting model object, type the following at the MATLAB prompt:

```
m_armax
```

MATLAB provides the following information about this model object:

```
Discrete-time IDPOLY model: A(q)y(t) = B(q)u(t) + e(t)
```

```
A(q) = 1 - 1.255q^-1 + 0.2551q^-2 + 0.2948q^-3 - 0.0619q^-4
```

```
B1(q) = -0.09168 + 1.105q^-1 + 0.7399q^-2
```

```
B2(q) = 1.022 + 0.129q^-1
```

```
B3(q) = -0.07605 + 0.08681q^-1 + 0.5619q^-2
```

```
C(q) = 1-0.06117q^-1 - 0.1461q^-2 + 0.009862q^-3 - 0.04313q^-4
```

```
Estimated using ARMAX from data set z8
```

```
Loss command 2.23844 and FPE 2.35202
```

```
Sampling interval: 1
```

`m_armax` is an `idpoly` model object. The coefficients represent estimated parameters of this polynomial model.

Tip You can use `present(m_armax)` to show additional information about the model, including parameter uncertainties.

- 4** To view all property values for this model, type the following command:

```
get(m_armax)
```

MATLAB lists the following object properties and values:

```
ans =  
      a: [1 -1.2549 0.2551 0.2948 -0.0619]  
      b: [3x3 double]  
      c: [1 -0.0612 -0.1461 0.0099 -0.0431]  
      d: 1  
      f: [3x1 double]  
     da: []  
     db: [3x0 double]  
     dc: []  
     dd: []  
     df: [3x0 double]  
     na: 4  
     nb: [3 2 3]  
     nc: 4  
     nd: 0  
     nf: [0 0 0]  
     nk: [0 0 0]  
InitialState: 'Auto'  
      Name: ''  
      Ts: 1  
      InputName: {3x1 cell}  
      InputUnit: {3x1 cell}  
      OutputName: {'y1'}  
      OutputUnit: {''}  
      TimeUnit: ''  
ParameterVector: [16x1 double]  
      PName: {}  
CovarianceMatrix: [16x16 double]  
      NoiseVariance: 0.9932  
      InputDelay: [3x1 double]  
      Algorithm: [1x1 struct]  
EstimationInfo: [1x1 struct]  
      Notes: {}  
      UserData: []
```

- 5** The `Algorithm` and `EstimationInfo` model properties are structures. To view the properties and values inside these structure, use dot notation. For example:

```
m_armax.Algorithm
```

This action displays the complete list of `Algorithm` properties and values that specify the iterative computational algorithm:

```
ans =  
    Approach: 'Pem'  
    Focus: 'Simulation'  
    MaxIter: 50  
    Tolerance: 1.0000e-005  
    LimitError: 1.6000  
    MaxSize: 'Auto'  
    SearchDirection: 'Auto'  
    FixedParameter: []  
    Trace: 'Off'  
    N4Weight: 'Auto'  
    N4Horizon: 'Auto'  
    Advanced: [1x1 struct]
```

Similarly, to view the properties and values of the EstimationInfo structure, type the following command:

```
m_armax.EstimationInfo
```

This action displays the complete list of read-only EstimationInfo properties and values that describe the estimation data set, quantitative measures of model quality (loss command and FPE), the number of iterations actually used, and the behavior of the iterative model estimation.

```
ans =  
      Status: 'Estimated model (PEM)'  
      Method: 'ARMAX'  
      LossFcn: 0.9602  
           FPE: 1.0263  
      DataName: 'z8'  
      DataLength: 500  
           DataTs: 1  
      DataDomain: 'Time'  
      DataInterSample: {3x1 cell}  
           WhyStop: 'Near (local) minimum, (norm(g)<tol).'  
           UpdateNorm: 8.0572e-006  
      LastImprovement: '7.4611e-006%'  
           Iterations: 4  
      InitialState: 'Zero'  
           Warning: 'None'
```

- 6** If you want to repeat the model estimation using different model orders, but keep the algorithm properties the same, you can store the model properties used for `m_armax` in a variable, as follows:

```
myAlg=m_armax.Algorithm
```

This action stores the specified focus, tolerance, and maxiter, and the default algorithm.

- 7** To reuse the algorithm properties in estimating the ARMAX model with different orders, use the following command:

```
m_armax2=armax(z8,[4 3 2 3 3 1 1 1],...  
              'algorithm',myAlg);
```


Model Properties

- “Categories of Model Object Properties” on page 3-13
- “Specifying Model Properties in the Estimator” on page 3-15
- “Accessing Model Properties” on page 3-16
- “Help on Properties at the Command Line” on page 3-18

Categories of Model Object Properties

The way a model object stores information is defined by the fields, or *properties*, of the corresponding class.

Each model object has properties for storing information that is relevant only to that specific model type. However, the `idarx`, `idgrey`, `idpoly`, `idproc`, and `idss` model objects are based on the `idmodel` superclass and share all `idmodel` properties.

Similarly, the nonlinear models `idnlarx`, `idnlhw`, and `idnlgrey` are based on the `idnlmodel` superclass and share all `idnlmodel` properties.

In general, all model objects have properties that belong to the following general categories:

- Names of input and output channels, such as `InputName` and `OutputName`
- Sampling interval of the model, such as `Ts`
- Units for time or frequency
- Properties that store estimation results and model uncertainty
- User comments, such as `Notes` and `Userdata`

All model objects (except as notes) have the following properties:

- **Algorithm**

This structure includes fields that specify how the nonlinear optimization search method works. `Algorithm` includes another structure, called `Advanced`, which provides additional flexibility for setting the optimization algorithm. Different fields apply for different estimation techniques.

For linear parametric models, `Algorithm` specifies the frequency weighing of the estimation using the `Focus` property.

Note `Algorithm` does not apply to `idfrd` models.

- **EstimationInfo**

This structure includes read-only fields that describe the estimation data set, quantitative model quality measures, the number of iterations actually used, how the initial states were handled, and whether any warnings were encountered during the estimation.

For information about how to list available object properties, see “Help on Properties at the Command Line” on page 3-18.

Specifying Model Properties in the Estimator

If you are estimating a new model, you can specify model properties directly in the estimator syntax. For a complete list of model estimation commands, see “Commands for Model Estimation” on page 3-5.

The following commands load the sample data, `z8`, and estimate an ARMAX model. The arguments of the `armax` estimator specify model properties as property-value pairs.

```
load iddata8
m_armax=armax(z8, 'na',4,...
               'nb',[3 2 3],...
               'nc',4,...
               'nk',[0 0 0],...
               'focus', 'simulation',...
               'covariance', 'none',...
               'tolerance',1e-5,...
               'maxiter',50);
```

`focus`, `covariance`, `tolerance`, and `maxiter` are fields in the Algorithm model property.

When using the commands that let you both construct and estimate a model, as described in “Commands for Model Estimation” on page 3-5, you can specify all top-level model properties in the estimator syntax. Top-level properties are those listed when you type `get(object_name)`. You can also specify the top-level fields of the Algorithm structure directly in the estimator using property-value pairs—such as `focus` in the previous example—without having to define the structure fields first.

For linear models, you can use a shortcut to specify the second-level Algorithm properties, such as `Advanced`. With this syntax, you can reference the structure fields by name without specifying the structure to which these fields belong.

However, when estimating nonlinear black-box models, you must set the specific fields of the Advanced Algorithm structure and the nonlinearity estimators before estimation. For example, suppose you want to set the value of the `wavenet` object property `Options`, which is a structure. The following

commands set the Options values before estimation and include the modified wavenet object in the estimator:

```
% Define wavenet object with default properties
W = wavenet;
% Specify variable to represent Options field
O = W.Options;
% Modify values of specific Options fields
O.MaxLevels = 5 ;
O.DilationStep = 2;
% Estimate model using new Options settings
M = nlarx(data,[2 2 1],wavenet('options',O))
```

where O specifies the values of the Options structure fields and M is the estimated model. For more information about these and other commands, see the corresponding reference page.

Accessing Model Properties

To view all the properties and values of any model object, use the get command. For example, type the following at the MATLAB prompt to load sample data, compute an ARX model, and list the model properties:

```
load iddata8
m_arx=arx(z8,[4 3 2 3 0 0 0]);
get(m_arx)
```

To access a specific property, use dot notation. For example, to view the A matrix containing the estimated parameters in the previous model, type the following command:

```
m_arx.a
```

MATLAB returns the following result:

```
ans =
    1.0000   -0.8441   -0.1539    0.2278    0.1239
```

Similarly, to access the uncertainties in these parameter estimates, type the following command:

```
m_arx.da
```

MATLAB returns the following result:

```
ans =  
    0    0.0357    0.0502    0.0438    0.0294
```

Property names are not case sensitive. You do not need to type the entire property name if the portion of the name you do enter uniquely identifies the property.

To change property values for an existing model object, use the `set` command or dot notation. For example, to change the input delays for all three input channels to `[1 1 1]`, type the following at the MATLAB prompt:

```
set(m_arx, 'nk', [1 1 1])
```

or equivalently

```
m_arx.nk = [1 1 1]
```

Some model properties, such as `Algorithm`, are structures. To access the fields in this structure, use the following syntax:

```
model.algorithm.PropertyName
```

where *PropertyName* represents any of the `Algorithm` fields. For example, to change the maximum number of iterations using the `MaxIter` property, type the following command:

```
m_arx.algorithm.MaxIter=50
```

To verify the new property value, type the following:

```
m_arx.algorithm.MaxIter
```

Note *PropertyName* refers to fields in a structure and is case sensitive. You must type the entire property name. Use the **Tab** key when typing property names to get completion suggestions.

Help on Properties at the Command Line

If you need quick assistance on model properties while working at the command line, you can use the `idprops` command to list the properties and values for each object.

Some model objects are based on the superclasses `idmodel` and `idnlmodel` and inherit the properties of these superclasses. For such model object, you must look up for property for both the model object and its superclass.

The following table summarizes the commands for getting help on object properties.

Help Commands for Model Properties

Model Class	Help Commands
<code>idarx</code>	<code>idprops idarx</code> Also inherits properties from <code>idmodel</code> .
<code>idfrd</code>	<code>idprops idfrd</code>
<code>idnlmodel</code>	<code>idprops idnlmodel</code>
<code>idmodel</code>	<code>idprops idmodel</code> <code>idprops idmodel Algorithm</code> <code>idprops idmodel EstimationInfo</code> Also see the <code>Algorithm</code> and <code>EstimationInfo</code> reference page.
<code>idproc</code>	<code>idprops idproc</code> Also inherits properties from <code>idmodel</code> .
<code>idpoly</code>	<code>idprops idpoly</code> Also inherits properties from <code>idmodel</code> .
<code>idss</code>	<code>idprops idss</code> Also inherits properties from <code>idmodel</code> .
<code>idgrey</code>	<code>idprops idgrey</code> Also inherits properties from <code>idmodel</code> .
<code>idnlgrey</code>	<code>idprops idnlgrey</code> <code>idprops idnlgrey Algorithm</code> <code>idprops idnlgrey EstimationInfo</code> Also inherits properties from <code>idnlmodel</code> .

Help Commands for Model Properties (Continued)

Model Class	Help Commands
idnlarx	idprops idnlarx idprops idnlarx Algorithm idprops idnlarx EstimationInfo Also inherits properties from idnlmodel.
idnlhw	idprops idnlhw idprops idnlhw Algorithm idprops idnlhw EstimationInfo Also inherits properties from idnlmodel.

See Also

Validate each model directly after estimation to help you fine-tune your modeling strategy. When you do not achieve a satisfactory model, you can try a different model structure and order, or try another identification algorithm. For more information about validating and troubleshooting models, see Chapter 10, “Validating and Analyzing Models”.

After you have selected a good model to represent your system, see Chapter 12, “Simulating and Predicting Model Output”.

The object-oriented concepts also apply to the data objects used to represent data in System Identification Toolbox. For detailed information about working with data objects, see Chapter 2, “Preparing Data”.

Subreferencing Models

In this section...
“What Is Subreferencing?” on page 3-20
“Limitation on Supported Models” on page 3-20
“Subreferencing Specific Measured Channels” on page 3-20
“Subreferencing Measured and Noise Models” on page 3-21
“Treating Noise Channels as Measured Inputs” on page 3-23

What Is Subreferencing?

You can use subreferencing to create models with subsets of inputs and outputs from existing multivariable models. Subreferencing is also useful when you want to generate model plots for only certain channels, such as when you are exploring multiple-output models for input channels that have minimal effect on the output.

System Identification Toolbox supports subreferencing operations for `idarx`, `idgrey`, `idpoly`, `idproc`, `idss`, and `idfrd` model objects.

In addition to subreferencing the model for specific combinations of measured inputs and output, you can subreference dynamic and noise models individually.

Limitation on Supported Models

Subreferencing nonlinear models is not supported.

Subreferencing Specific Measured Channels

Use the following general syntax to subreference specific input and output channels in models:

```
model(outputs,inputs)
```

In this syntax, `outputs` and `inputs` specify channel indexes or channel names.

To select all output or all input channels, use a colon (:). To select no channels, specify an empty matrix ([]). If you need to reference several channel names, use a cell array of strings.

For example, to create a new model `m2` from `m` from inputs 1 ('power') and 4 ('speed') to output number 3 ('position'), use either of the following equivalent commands:

```
m2 = m('position', {'power', 'speed'})
```

or

```
m2 = m(3, [1 4])
```

For a single-output model, you can use the following syntax to subreference specific input channels without ambiguity:

```
m3 = m(inputs)
```

Similarly, for a single-input model, you can use the following syntax to subreference specific output channels:

```
m4 = m(outputs)
```

Subreferencing Measured and Noise Models

For linear models, the general symbolic model description is given by:

$$y = Gu + He$$

G is an operator that takes the measured inputs u to the outputs and captures the system dynamics.

H is an operator that describes the properties of the additive output disturbance and takes the hypothetical (unmeasured) noise source inputs to the outputs. H represents the noise model. When you estimate a noise model, System Identification Toolbox includes one noise channel at the input e for each output in your system.

Therefore, linear, parametric models represent input-output relationships for two kinds of input channels: measured inputs and (unmeasured) noise inputs. For example, consider the ARX model given by one of the following equations:

$$A(q)y(t) = B(q)u(t - nk) + e(t)$$

or

$$y(t) = \frac{B(q)}{A(q)}u(t) + \frac{1}{A(q)}e(t)$$

In this case, the dynamic model is the relationship between the measured input u and output y , $G = B(q)/A(q)$. The noise model is the contribution of the input noise e to the output y , given by $H = 1/A(q)$.

Suppose that the model `m` contains both a dynamic model G and a noise model H . To create a new model by subreferencing G due to measured inputs, use the following syntax:

```
m_G = m('measured')
```

Tip Alternatively, you can use the following shorthand syntax: `m_G = m('m')`

To create a new model by subreferencing H due to unmeasured inputs, use the following syntax:

```
m_H = m('noise')
```

Tip Alternatively, you can use the following shorthand syntax: `m_H = m('n')`

This operation creates a time-series model from m by ignoring the measured input.

The covariance matrix of e is given by the `idmodel` property `NoiseVariance`, which is the matrix Λ :

$$\Lambda = LL^T$$

The covariance matrix of e is related to v , as follows:

$$e = Lv$$

where v is white noise with an identity covariance matrix representing independent noise sources with unit variances.

Treating Noise Channels as Measured Inputs

To study noise contributions in more detail, it might be useful to convert the noise channels to measured channels using `noisecnv`:

```
m_GH = noisecnv(m)
```

This operation creates a model `m_GH` that represents both measured inputs u and noise inputs e , treating both sources as measured signals. `m_GH` is a model from u and e to y , describing the transfer functions G and H .

Converting noise channels to measured inputs loses information about the variance of the innovations e . For example, step response due to the noise channels does not take into consideration the magnitude of the noise contributions. To include this variance information, normalize e such that v becomes white noise with an identity covariance matrix, where

$$e = Lv$$

To normalize e , use the following command:

```
m_GH = noisecnv(m, 'Norm')
```

This command creates a model where u and v are treated as measured signals, as follows:

$$y(t) = Gu(t) + HLv = \begin{bmatrix} G & HL \end{bmatrix} \begin{bmatrix} u \\ v \end{bmatrix}$$

For example, the scaling by L causes the step responses from v to y to reflect the size of the disturbance influence.

The converted noise sources are named in a way that relates the noise channel to the corresponding output. Unnormalized noise sources e are assigned names such as 'e@y1', 'e@y2', ..., 'e@yn', where 'e@yn' refers to the noise input associated with the output yn. Similarly, normalized noise sources v , are named 'v@y1', 'v@y2', ..., 'v@yn'.

Note When you plot models in the GUI that include noise sources, you can select to view the response of the noise model corresponding to specific outputs. For more information, see “Selecting Measured and Noise Channels in Plots” on page 1-23.

Concatenating Model Objects

In this section...

“About Concatenating Models” on page 3-25

“Limitation on Supported Models” on page 3-25

“Horizontal Concatenation of Model Objects” on page 3-26

“Vertical Concatenation of Model Objects” on page 3-26

“Concatenating Noise Spectral Data of idfrd Objects” on page 3-27

“See Also” on page 3-28

About Concatenating Models

You can perform horizontal and vertical concatenation of model objects to grow the number of inputs or outputs in the model.

When you concatenate parametric models, such as `idarx`, `idgrey`, `idpoly`, `idproc`, and `idss` model objects, the resulting model combines the parameters of the individual models.

You can also concatenate nonparametric models, which contain the estimated impulse-response (`idarx` object) and frequency-response (`idfrd` object) of a system.

In case of `idfrd` models, concatenation combines information in the `ResponseData` properties of the individual model objects. `ResponseData` is an `ny-by-nu-by-nf` array that stores the response of the system, where `ny` is the number of output channels, `nu` is the number of input channels, and `nf` is the number of frequency values. The `(j,i,:)` vector of the resulting response data represents the frequency response from the `i`th input to the `j`th output at all frequencies.

For more information, see the corresponding reference page.

Limitation on Supported Models

Concatenation is supported for linear models only.

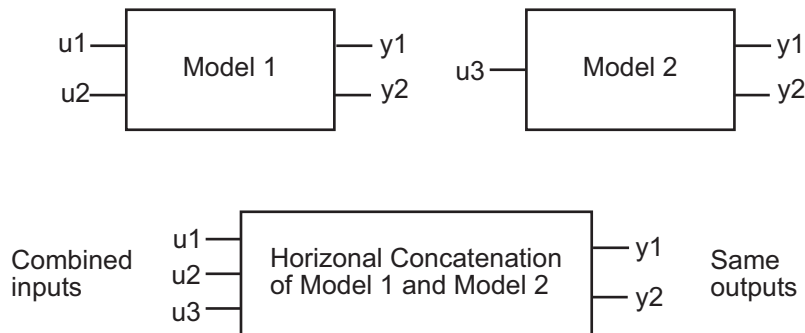
Horizontal Concatenation of Model Objects

Horizontal concatenation of model objects requires that they have the same outputs. If the output channel names are different and their dimensions are the same, the concatenation operation uses the names of output channels in the first model object you listed. Input channels must have unique names.

The following syntax creates a new model object m that contains the horizontal concatenation of m_1, m_2, \dots, m_N :

$$m = [m_1, m_2, \dots, m_N]$$

m takes all of the inputs of m_1, m_2, \dots, m_N to the same outputs as in the original models. The following diagram is a graphical representation of horizontal concatenation of the models.



Note Horizontal concatenation of `idarx` objects creates an `idss` object.

Vertical Concatenation of Model Objects

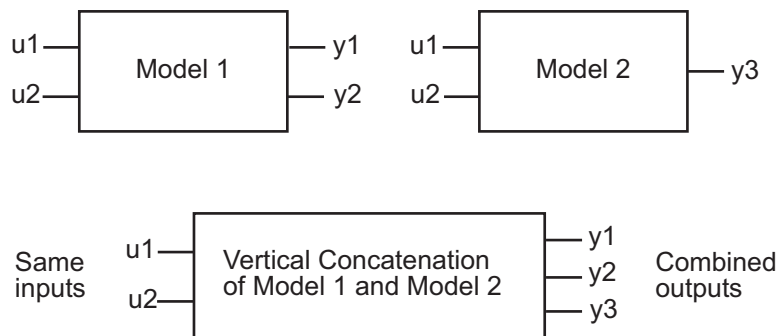
Vertical concatenation combines output channels of specified models. Vertical concatenation of model objects requires that they have the same inputs and frequency vectors. If the input channel names are different and their dimensions are the same, the concatenation operation uses the names of input channels in the first model object you listed. Output channels must have unique names.

Note You cannot concatenate the single-output `idproc` and `idpoly` model objects.

The following syntax creates a new model object `m` that contains the vertical concatenation of `m1`, `m2`, \dots , `mN`:

$$m = [m1;m2;\dots ;mN]$$

`m` takes the same inputs in the original models to all of the output of `m1`, `m2`, \dots , `mN`. The following diagram is a graphical representation of vertical concatenation of frequency-response data.



Concatenating Noise Spectral Data of `idfrd` Objects

When the `idfrd` objects contain the frequency-response data you measured or constructed manually, the concatenation operation combines only the `ResponseData` properties. Because noise spectral data does not exist (unless you also entered it manually), `SpectralData` is empty in both the individual `idfrd` objects and the concatenated `idfrd` object.

However, when the `idfrd` objects are spectral models that you estimated, the `SpectralData` property is not empty and contains the power spectra and cross spectra of the output noise in the system. For each output channel, this toolbox estimates one noise channel to explain the difference between the output of the model and the measured output.

When the `SpectralData` property of individual `idfrd` objects is not empty, horizontal and vertical concatenation handle `SpectralData`, as follows.

In case of horizontal concatenation, there is no meaningful way to combine the `SpectralData` of individual `idfrd` objects, and the resulting `SpectralData` property is empty. An empty property results because each `idfrd` object has its own set of noise channels, where the number of noise channels equals the number of outputs. When the resulting `idfrd` object contains the same output channels as each of the individual `idfrd` objects, it cannot accommodate the noise data from all the `idfrd` objects.

In case of vertical concatenation, this toolbox concatenates individual noise models diagonally. The following shows that `m.SpectrumData` is a block diagonal matrix of the power spectra and cross spectra of the output noise in the system:

$$m.s = \begin{pmatrix} m1.s & & \mathbf{0} \\ & \ddots & \\ \mathbf{0} & & mN.s \end{pmatrix}$$

`s` in `m.s` is the abbreviation for the `SpectrumData` property name.

See Also

If you have Control System Toolbox, see “Combining Model Objects” on page 13-6 about additional functionality for combining models.

Merging Model Objects

You can merge models of the same structure to obtain a single model with parameters that are statistically weighed means of the parameters of the individual models. When computing the merged model, the covariance matrices of the individual models determine the weights of the parameters.

You can perform the merge operation for the `idarx`, `idgrey`, `idpoly`, `idproc`, and `idss` model objects.

Note Each merge operation merges the same type of model object.

Merging models is an alternative to merging data sets into a single multiexperiment data set, and then estimating a model for the merged data. Whereas merging data sets assumes that the signal-to-noise ratios are about the same in the two experiments, merging models allows greater variations in model uncertainty, which might result from greater disturbances in an experiment.

When the experimental conditions are about the same, merge the data instead of models. This approach is more efficient and typically involves better-conditioned calculations. For more information about merging data sets into a multiexperiment data set, see “Creating Multiexperiment Data at the Command Line” on page 2-38.

For more information about merging models, see the `merge` reference page.

Estimating Nonparametric Models

Estimating Time-Domain
Correlation Models (p. 4-2)

Estimating impulse- and
step-response models using
correlation analysis.

Estimating Frequency-Response
Models (p. 4-10)

Estimating nonparametric
frequency-response models.

Estimating Time-Domain Correlation Models

In this section...

“What Is Time-Domain Correlation Analysis?” on page 4-2

“Data Supported by Correlation Analysis” on page 4-3

“How to Estimate Correlation Models Using the GUI” on page 4-3

“How to Estimate Correlation Models at the Command Line” on page 4-4

“How to Compute Response Values” on page 4-6

“How to Identify Delay Using Transient-Response Plots” on page 4-6

“Algorithm for Correlation Analysis” on page 4-8

What Is Time-Domain Correlation Analysis?

Time-domain correlation analysis is a nonparametric estimate of transient response of dynamic systems, which computes a finite impulse response (FIR) model from the data. Correlation analysis assumes a linear system and does not require a specific model structure.

There are two types of transient response for a dynamic model:

- Impulse response

Impulse response is the output signal that results when the input is an impulse and has the following definition for a discrete model:

$$u(t) = 0 \quad t > 0$$

$$u(t) = 1 \quad t = 0$$

- Step response

Step response is the output signal that results from a step input, defined as follows:

$$u(t) = 0 \quad t < 0$$

$$u(t) = 1 \quad t \geq 0$$

The response to an input $u(t)$ is equal to the convolution of the impulse response, as follows:

$$y(t) = \int_0^t h(t-z) \cdot u(z) dz$$

Data Supported by Correlation Analysis

You can estimate correlation analysis models from data with the following characteristics:

- Real or complex time-domain `iddata` object. To learn about estimating time-series models, see Chapter 6, “Estimating Time-Series Models”.
- Frequency-domain `iddata` or `idfrd` object with the sampling interval $T \neq 0$.
- Single- or multiple-output data.

How to Estimate Correlation Models Using the GUI

The following procedure assumes that you already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

To estimate impulse- and step-response models in the System Identification Tool GUI using time-domain correlation analysis:

- 1** In the System Identification Tool GUI, select **Estimate > Correlation models** to open the Correlation Model dialog box.
- 2** In the **Time span (s)** field, specify a scalar value as the time interval over which the impulse or step response is calculated. For a scalar time span T , the resulting response is plotted from $-T/4$ to T .

Tip You can also enter a 2-D vector in the format `[min_value max_value]`.

- 3** In the **Order of whitening filter** field, specify the filter order.

The prewhitening filter is determined by modeling the input as an Auto-Regressive (AR) process of order N . The algorithm applies a filter of the form $A(q)u(t)=u_F(t)$. That is, the input $u(t)$ is subjected to an FIR filter A to produce the filtered signal $u_F(t)$. *Prewhitening* the input by applying a whitening filter before estimation might improve the quality of the estimated impulse response g .

The order of the prewhitening filter, N , is the order of the A filter. N equals the number of lags. The default value of N is 10, which you can also specify as `[]`.

- 4** In the **Model Name** field, enter the name of the correlation analysis model. The name of the model should be unique in the Model Board.
- 5** Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
- 6** In the Correlation Model dialog box, click **Close**.
- 7** To view the transient response plot, select the **Transient resp** check box in the System Identification Tool GUI. For more information about working with this plot and selecting to view impulse- versus step-response, see “Using Impulse- and Step-Response Plots to Validate Models” on page 10-21.

You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

How to Estimate Correlation Models at the Command Line

You can use `impulse` and `step` commands to estimate the impulse and step response directly from time- or frequency-domain data using correlation analysis. Both `impulse` and `step` produce the same FIR model, but generate different plots.

Note `cra` is an alternative method for computing impulse response from time-domain data only.

The following tables summarize the commands for computing impulse- and step-response models. The resulting models are stored as `idarx` model objects and contain impulse-response coefficients in the model parameter `B`. For detailed information about these commands, see the corresponding reference page.

Commands for Impulse and Step Response

Command	Description	Example
<code>impulse</code>	Estimates a high-order, noncausal FIR model using correlation analysis.	<p>To estimate the model <code>m</code> and plot the impulse response, use the following syntax:</p> <pre>m=impulse(data,Time,'pw',N)</pre> <p>where <code>data</code> is a single- or multiple-output time-domain <code>iddata</code> object, and <code>Time</code> is a scalar value representing the time interval over which the impulse or step response is calculated. For a scalar time span T, the resulting response is plotted from $-T/4$ to T. <code>'pw'</code> and <code>N</code> is an option property-value pair that specifies the order <code>N</code> of the prewhitening filter <code>'pw'</code>.</p>
<code>step</code>	Estimates a high-order, noncausal FIR model correlation analysis.	<p>To estimate the model <code>m</code> and plot the step response, use the following syntax:</p> <pre>step(data,Time)</pre> <p>where <code>data</code> is a single- or multiple-output time-domain <code>iddata</code> object, and <code>Time</code> is the time span.</p>

To validate the model, see Chapter 10, “Validating and Analyzing Models”. For more information about continuing to work with models in the MATLAB Workspace browser, see Chapter 13, “Using Models with Other Products”.

How to Compute Response Values

You can use `impulse` and `step` commands with output arguments to get the numerical impulse- and step-response vectors as a command of time, respectively.

To get the numerical response values:

- 1 Compute the FIR model by applying either `impulse` or `step` commands on the data, as described in “How to Estimate Correlation Models at the Command Line” on page 4-4.
- 2 Apply the following syntax on the resulting model:


```
% To compute impulse-response data
[y,t,yzd] = impulse(model)
% To compute step-response data
[y,t,yzd] = step(model)
```

where `y` is the response data, `t` is the time vector, and `yzd` is the standard deviations of the response.

How to Identify Delay Using Transient-Response Plots

You can use transient-response plots to estimate the input delay, or *dead time*, of linear systems. Input delay represents the time it takes for the output to respond to the input.

In the System Identification Tool GUI. To view the transient response plot, select the **Transient resp** check box in the System Identification Tool GUI. For example, the following step response plot shows a time delay of about 0.25 s before the system responds to the input.



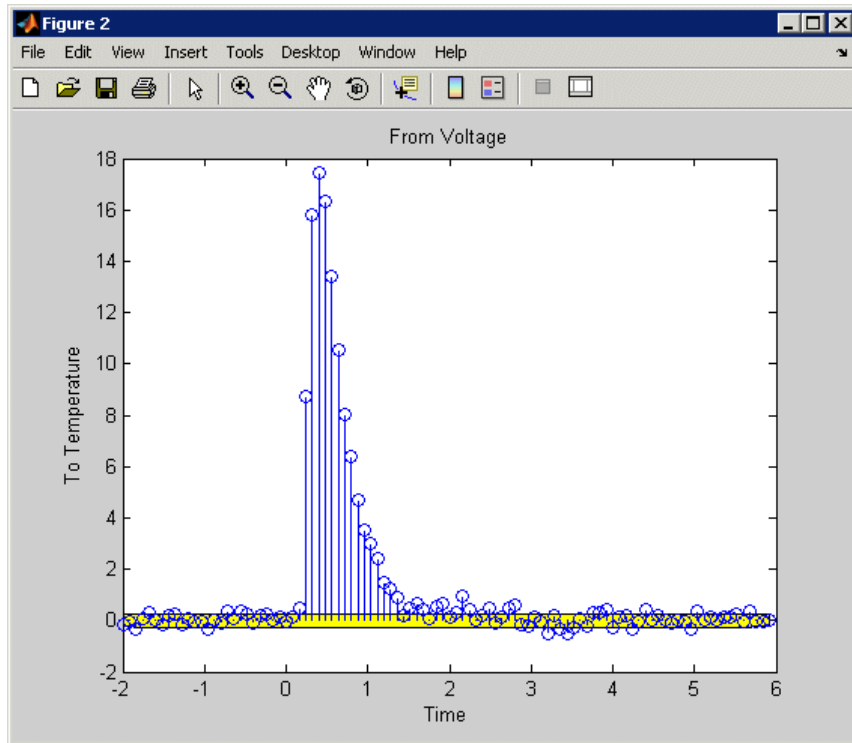
Step Response Plot

At the command line. You can use the `impulse` command to plot the impulse response. The time delay is equal to the first positive peak in the transient response magnitude that is greater than the confidence region for positive time values.

For example, the following commands create an impulse-response plot with a 1-standard-deviation confidence region:

```
% Load sample data
load dry2
% Split data into estimation and
% validation data sets
ze = dry2(1:500);
zr = dry2(501:1000);
impulse(ze,'sd',1,'fill')
```

The resulting figure shows that the first positive peak of the response magnitude, which is greater than the confidence region for positive time values, occurs at 0.24 s.



Algorithm for Correlation Analysis

To better understand the algorithm underlying correlation analysis, consider the following description of a dynamic system:

$$y(t) = G(q)u(t) + v(t)$$

where $u(t)$ and $y(t)$ are the input and output signals, respectively. $v(t)$ is the additive noise term. $G(q)$ is the transfer command of the system. The $G(q)u(t)$ notation represents the following operation:

$$G(q)u(t) = \sum_{k=1}^{\infty} g(k)u(t-k)$$

q is the *shift operator*, defined by the following equation:

$$G(q) = \sum_{k=1}^{\infty} g(k)q^{-k} \quad q^{-1}u(t) = u(t-1)$$

For impulse response, the algorithm estimates impulse response coefficients g for both the single- and multiple-output data. The impulse response is estimated as a high-order, noncausal FIR model:

$$y(t) = g(-m)u(t+m) + \dots + g(-1)u(t+1) + g(0)u(t) \\ + g(1)u(t-1) + \dots + g(n)u(t-n)$$

The estimation algorithm prefilters the data such that the input is as white as possible. It then computes the correlations from the prefiltered data to obtain the FIR coefficients.

g is also estimated for negative lags, which takes into account any noncausal effects from input to output. Noncausal effects can result from feedback. The coefficients are computed using the least-squares method.

For a multiple-input or multiple-output system, the impulse response g_k is an ny -by- nu matrix, where ny is the number of outputs and nu is the number of inputs. The i - j th element of the impulse response matrix describes the behavior of the i th output after an impulse in the j th input.

Estimating Frequency-Response Models

In this section...

“What Is a Frequency-Response Model?” on page 4-10

“Data Supported by Frequency-Response Models” on page 4-11

“How to Estimate Frequency-Response Models in the GUI” on page 4-11

“How to Estimate Frequency-Response Models at the Command Line” on page 4-13

“Options for Computing Spectral Models” on page 4-13

“Options for Frequency Resolution” on page 4-14

“Spectral Analysis Algorithm” on page 4-16

“Understanding Spectrum Normalization” on page 4-19

What Is a Frequency-Response Model?

You can estimate *frequency-response models* and visualize the responses on a Bode plot, which shows the amplitude change and the phase shift as a function of the sinusoid frequency.

For a discrete-time system sampled with a time interval T , the frequency-response model $G(z)$ relates the Z-transforms of the input $U(z)$ and output $Y(z)$:

$$Y(z) = G(z)U(z)$$

The frequency-response command describes the steady-state response of a system to sinusoidal inputs. For a linear system, a sinusoidal input of a specific frequency results in an output that is also a sinusoid with the same frequency, but with a different amplitude and phase. The frequency response describes the amplitude change and phase shift as a function of frequency.

In other words, the frequency-response command, $G(e^{i\omega T})$, is the Laplace transform of the impulse response that is evaluated on the imaginary axis. The frequency-response command is the transfer command $G(z)$ evaluated on the unit circle.

Data Supported by Frequency-Response Models

You can estimate spectral analysis models from data with the following characteristics:

- Complex or real data.
- Time- or frequency-domain `iddata` or `idfrd` data object. To learn more about estimating time-series models, see Chapter 6, “Estimating Time-Series Models”.
- Single- or multiple-output data.

How to Estimate Frequency-Response Models in the GUI

You must have already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

To estimate frequency-response models in the System Identification Tool GUI:

- 1** In the System Identification Tool GUI, select **Estimate > Spectral models** to open the Spectral Model dialog box.
- 2** In the **Method** list, select the spectral analysis method you want to use. For information about each method, see “Options for Computing Spectral Models” on page 4-13.
- 3** Specify the frequencies at which to compute the spectral model in *one* of the following ways:
 - In the **Frequencies** field, enter either a vector of values, a MATLAB expression that evaluates to a vector, or a variable name of a vector in the MATLAB Workspace browser. For example, `logspace(-1, 2, 500)`.
 - Use the combination of **Frequency Spacing** and **Frequencies** to construct the frequency vector of values:
 - In the **Frequency Spacing** list, select Linear or Logarithmic frequency spacing.

Note For `etfe`, only the `Linear` option is available.

- In the **Frequencies** field, enter the number of frequency points.

For time-domain data, the frequency ranges from 0 to the Nyquist frequency. For frequency-domain data, the frequency ranges from the smallest to the largest frequency in the data set.

- 4** In the **Frequency Resolution** field, enter the frequency resolution, as described in “Options for Frequency Resolution” on page 4-14. To use the default value, enter `default` or, equivalently, the empty matrix `[]`.
- 5** In the **Model Name** field, enter the name of the correlation analysis model. The model name should be unique in the Model Board.
- 6** Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
- 7** In the Spectral Model dialog box, click **Close**.
- 8** To view the frequency-response plot, select the **Frequency resp** check box in the System Identification Tool GUI. For more information about working with this plot, see “Using Frequency-Response Plots to Validate Models” on page 10-29.
- 9** To view the estimated disturbance spectrum, select the **Noise spectrum** check box in the System Identification Tool GUI. For more information about working with this plot, see “Creating Noise-Spectrum Plots” on page 10-37.
- 10** After estimating the model, see Chapter 10, “Validating and Analyzing Models”, to validate the model.

To export the model to the MATLAB Workspace browser, drag it to the **To Workspace** rectangle in the System Identification Tool GUI. You can retrieve the responses from the resulting `idfrd` model object using the `bode` or `nyquist` command.

How to Estimate Frequency-Response Models at the Command Line

You can use the `etfe`, `spa`, and `spafdr` commands to estimate spectral models. The following table provides a brief description of each command and usage examples.

The resulting models are stored as `idfrd` model objects. For detailed information about the commands and their arguments, see the corresponding reference page.

Commands for Frequency Response

Command	Description	Usage
<code>etfe</code>	Estimates an empirical transfer command using Fourier analysis.	To estimate a model <code>m</code> , use the following syntax: <code>m=etfe(data)</code>
<code>spa</code>	Estimates a frequency response with a fixed frequency resolution using spectral analysis.	To estimate a model <code>m</code> , use the following syntax: <code>m=spa(data)</code>
<code>spafdr</code>	Estimates a frequency response with a variable frequency resolution using spectral analysis.	To estimate a model <code>m</code> , use the following syntax: <code>m=spafdr(data,R,w)</code> where <code>R</code> is the resolution vector and <code>w</code> is the frequency vector.

After estimating the model, see Chapter 10, “Validating and Analyzing Models” to validate the model.

Options for Computing Spectral Models

This section describes how to select the method for computing spectral models in the estimation procedures “How to Estimate Frequency-Response Models in the GUI” on page 4-11 and “How to Estimate Frequency-Response Models at the Command Line” on page 4-13.

System Identification Toolbox provides the following three spectral-analysis methods:

- **etfe (Empirical Transfer Function Estimate)**

For input-output data. This method computes the ratio of the Fourier transform of the output to the Fourier transform of the input.

For time-series data. This method computes a periodogram as the normalized absolute squares of the Fourier transform of the time series.

ETFE works well for highly resonant systems or narrowband systems. The drawback of this method is that it requires linearly spaced frequency values, does not estimate the disturbance spectrum, and does not provide confidence intervals. ETFE also works well for periodic inputs and computes exact estimates at multiples of the fundamental frequency of the input and their ratio.

- **spa (SPectral Analysis)**

This method is the Blackman-Tukey spectral analysis method, where windowed versions of the covariance functions are Fourier transformed. For more information about this algorithm, see “Spectral Analysis Algorithm” on page 4-16.

- **spafdr (SPectral Analysis with Frequency Dependent Resolution)**

This method is a variant of the Blackman-Tukey spectral analysis method with frequency-dependent resolution. First, the algorithm computes Fourier transforms of the inputs and outputs. Next, the products of the transformed inputs and outputs with the conjugate input transform are smoothed over local frequency regions. The widths of the local frequency regions can vary as a command of frequency. The ratio of these averages computes the frequency-response estimate.

Options for Frequency Resolution

- “What Is Frequency Resolution?” on page 4-15
- “Frequency Resolution for etfe and spa” on page 4-15
- “Frequency Resolution for spafdr” on page 4-15
- “etfe Frequency Resolution for Periodic Input” on page 4-16

This section supports the estimation procedures “How to Estimate Frequency-Response Models in the GUI” on page 4-11 and “How to Estimate Frequency-Response Models at the Command Line” on page 4-13.

What Is Frequency Resolution?

Frequency resolution is the size of the smallest frequency for which details in the frequency response and the spectrum can be resolved by the estimate. A resolution of 0.1 rad/s means that the frequency response variations at frequency intervals at or below 0.1 rad/s are not resolved.

Note Finer resolution results in greater uncertainty in the model estimate.

Specifying the frequency resolution for `etfe` and `spa` is different than for `spafdr`.

Frequency Resolution for `etfe` and `spa`

For `etfe` and `spa`, the frequency resolution is approximately equal to the following value:

$$\frac{2\pi}{M} \left(\frac{\text{radians}}{\text{sampling interval}} \right)$$

M is a scalar integer that sets the size of the lag window.

A large value of M gives good resolution, but results in more uncertain estimates.

The default value of M for `spa` is good for systems without sharp resonances.

For `etfe`, the default value of M gives the maximum resolution.

Frequency Resolution for `spafdr`

In case of `etfe` and `spa`, the frequency response is defined over a uniform frequency range, 0 - $F_s/2$ radians per second, where F_s is the sampling frequency—equal to twice the Nyquist frequency. In contrast, `spafdr` lets

you increase the resolution in a specific frequency range, such as near a resonance frequency. Conversely, you can make the frequency grid coarser in the region where the noise dominates—at higher frequencies, for example. Such customizing of the frequency grid assists in the estimation process by achieving high fidelity in the frequency range of interest.

For `spafdr`, the frequency resolution around the frequency k is the value $R(k)$. You can enter $R(k)$ in any *one* of the following ways:

- Scalar value of the constant frequency resolution value in radians per second.

Note The scalar R is inversely related to the M value used for `etfe` and `spa`.

- Vector of frequency values the same size as the frequency vector.
- Expression using MATLAB workspace variables and evaluates to a resolution vector that is the same size as the frequency vector.

The default value of the resolution for `spafdr` is twice the difference between neighboring frequencies in the frequency vector.

etfe Frequency Resolution for Periodic Input

If the input data is marked as periodic and contains an integer number of periods (`data.Period` is an integer), `etfe` computes the frequency response at

frequencies $\frac{2\pi k}{T} \left(\frac{k}{\text{Period}} \right)$ where $k = 1, 2, \dots, \text{Period}$.

For periodic data, the frequency resolution is ignored.

Spectral Analysis Algorithm

You can estimate the frequency-response command of dynamic systems using spectral analysis.

To better understand the spectral analysis algorithm, consider the following description of a linear, dynamic system:

$$y(t) = G(q)u(t) + v(t)$$

where $u(t)$ and $y(t)$ are the input and output signals, respectively. $G(q)$ is called the transfer command of the system—it takes the input to the output and captures the system dynamics. The $G(q)u(t)$ notation represents the following operation:

$$G(q)u(t) = \sum_{k=1}^{\infty} g(k)u(t-k)$$

q is the *shift operator*, defined by the following equation:

$$G(q) = \sum_{k=1}^{\infty} g(k)q^{-k} \quad q^{-1}u(t) = u(t-1)$$

$G(q)$ that is evaluated on the unit circle, $G(q=e^{i\omega})$, is the *frequency-response command*.

Together, $G(q=e^{i\omega})$ and the output noise spectrum $\hat{\Phi}_v(\omega)$ comprise the frequency-domain description of the system.

According to the Blackman-Tukey approach, the estimated frequency-response command is given by the following equation:

$$\hat{G}_N(e^{i\omega}) = \frac{\hat{\Phi}_{yu}(\omega)}{\hat{\Phi}_u(\omega)}$$

In this case, $\hat{}$ represents approximate quantities. For a derivation of this equation, see the chapter on nonparametric time- and frequency-domain methods in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

The output noise spectrum is given by the following equation:

$$\hat{\Phi}_v(\omega) = \hat{\Phi}_y(\omega) - \frac{|\hat{\Phi}_{yu}(\omega)|^2}{\hat{\Phi}_u(\omega)}$$

This equation for the noise spectrum is derived by assuming the linear relationship $y(t) = G(q)u(t) + v(t)$, that $u(t)$ is independent of $v(t)$, and the following relationships between the spectra:

$$\Phi_y(\omega) = \left|G(e^{i\omega})\right|^2 \Phi_u(\omega) + \Phi_v(\omega)$$

$$\Phi_{yu}(\omega) = G(e^{i\omega})\Phi_u(\omega)$$

where the noise spectrum is given by the following equation:

$$\Phi_v(\omega) \equiv \sum_{\tau=-\infty}^{\infty} R_v(\tau)e^{-i\omega\tau}$$

$\hat{\Phi}_{yu}(\omega)$ is the output-input cross-spectrum and $\hat{\Phi}_u(\omega)$ is the input spectrum.

When you use System Identification Toolbox to estimate the frequency-response command, the spectral estimation algorithms (such as `spa`) perform the following steps:

- 1 Compute the covariances and cross-covariance from $u(t)$ and $y(t)$, as follows:

$$\hat{R}_y(\tau) = \frac{1}{N} \sum_{t=1}^N y(t+\tau)y(t)$$

$$\hat{R}_u(\tau) = \frac{1}{N} \sum_{t=1}^N u(t+\tau)u(t)$$

$$\hat{R}_{yu}(\tau) = \frac{1}{N} \sum_{t=1}^N y(t+\tau)u(t)$$

- 2** Compute the Fourier transforms of the covariances and the cross-covariance, as follows:

$$\hat{\Phi}_y(\omega) = \sum_{\tau=-M}^M \hat{R}_y(\tau) W_M(\tau) e^{-i\omega\tau}$$

$$\hat{\Phi}_u(\omega) = \sum_{\tau=-M}^M \hat{R}_u(\tau) W_M(\tau) e^{-i\omega\tau}$$

$$\hat{\Phi}_{yu}(\omega) = \sum_{\tau=-M}^M \hat{R}_{yu}(\tau) W_M(\tau) e^{-i\omega\tau}$$

where $W_M(\tau)$ is called the *lag window* with the width M .

- 3** Compute the frequency-response command $\hat{G}_N(e^{i\omega})$ and the output noise spectrum $\hat{\Phi}_v(\omega)$

Alternatively, the disturbance $v(t)$ can be described as filtered white noise:

$$v(t) = H(q)e(t)$$

where $e(t)$ is the white noise with variance λ and the noise power spectrum is given by the following equation:

$$\Phi_v(\omega) = \lambda \left| H(e^{i\omega}) \right|^2$$

Understanding Spectrum Normalization

The *spectrum* of a signal is the square of the Fourier transform of the signal. The spectral estimate using the commands `spa`, `spafdr`, and `etfe` is normalized by the sampling interval T :

$$\Phi_y(\omega) = T \sum_{k=-M}^M R_y(kT) e^{-i\omega T} W_M(k)$$

where $W_M(k)$ is the lag window, and M is the width of the lag window. The output covariance $R_y(kT)$ is given by the following discrete representation:

$$\hat{R}_y(kT) = \frac{1}{N} \sum_{l=1}^N y(lT - kT)y(lT)$$

Because there is no scaling in a discrete Fourier transform of a vector, the purpose of T is to relate the discrete transform of a vector to the physically meaningful transform of the measured signal. This normalization sets the units of $\Phi_y(\omega)$ as power per radians per unit time, and makes the frequency units radians per unit time.

The scaling factor of T is necessary to preserve the energy density of the spectrum after interpolation or decimation.

By Parseval's theorem, the average energy of the signal must equal the average energy in the estimated spectrum, as follows:

$$Ey^2(t) = \frac{1}{2\pi} \int_{-\pi/T}^{\pi/T} \Phi_y(\omega) d\omega$$

$$S1 \equiv Ey^2(t)$$

$$S2 \equiv \frac{1}{2\pi} \int_{-\pi/T}^{\pi/T} \Phi_y(\omega) d\omega$$

To compare the left side of the equation (S1) to the right side (S2), enter the following commands in MATLAB:

```
load iddata1
% Create time-series iddata object
y = z1(:,1,[]);
% Define sample interval from the data
T = y.Ts;
% Estimate frequency response
sp = spa(y);
% Remove spurious dimensions
phiy = squeeze(sp.spec);
% Compute average energy from the estimated
```

```
% energy spectrum, where S1 is scaled by T
S1 = sum(phiy)/length(phiy)/T
% Compute average energy of the signal
S2 = sum(y.y.^2)/size(y,1)
```

In this code, phiy contains $\Phi_y(\omega)$ between $\omega=0$ and $\omega=\pi/T$ with the frequency step given as follows:

$$\left(\frac{\pi}{T \cdot \text{length}(\text{phiy})} \right)$$

MATLAB responds with the following values for S1 and S2:

```
S1 =
    19.2076
S2 =
    19.4646
```

Thus, the average energy of the signal approximately equals the average energy in the estimated spectrum.

Estimating Linear Parametric Models

Estimating Process Models (p. 5-2)

Estimating the parameters of linear, low-order, continuous-time transfer functions.

Estimating Polynomial Models (p. 5-21)

Estimating the parameters of linear black-box polynomial models.

Estimating State-Space Models (p. 5-48)

Estimating the parameters of linear state-space models.

Refining Linear Parametric Models (p. 5-78)

Procedures for refining model parameters after estimating a model or constructing the model with initial parameter guesses.

Estimating Process Models

In this section...

- “What Is a Process Model?” on page 5-2
- “Data Supported by a Process Model” on page 5-3
- “How to Estimate Process Models Using the GUI” on page 5-3
- “Estimating Process Models at the Command Line” on page 5-9
- “Options for Specifying the Process-Model Structure” on page 5-15
- “Options for Multiple-Input Models” on page 5-16
- “Options for the Disturbance Model Structure” on page 5-17
- “Options for Frequency-Weighing Focus” on page 5-18
- “Options for Initial States” on page 5-19

What Is a Process Model?

The structure of a *continuous-time process model* is a simple transfer function that describes linear system dynamics in terms of one or more of the following elements:

- Static gain K_p .
- One or more time constants T_{pk} . For complex poles, the time constant is called T_ω —equal to the inverse of the natural frequency—and the damping coefficient is ζ (zeta).
- Process zero T_z .
- Possible time delay T_d before the system output responds to the input (*dead time*).
- Possible enforced integration.

Process models are popular for describing system dynamics in many industries and apply to various production environments. The primary advantages of these models are that they provide delay estimation, and the model coefficients have a physical interpretation.

System Identification Toolbox lets you create different process-model structures for both time-domain and frequency-domain data by varying the number of poles, adding an integrator, or adding or removing a time delay or a zero. In this toolbox, you can specify a first-, second-, or third-order model, and the poles can be real or complex (underdamped modes).

Note Continuous-time process models let you estimate the input delay.

For example, the following model structure is a first-order continuous-time process model, where K is the static gain, T_{p1} is a time constant, and T_d is the input-to-output delay:

$$G(s) = \frac{K}{1 + sT_{p1}} e^{-sT_d}$$

To learn more about estimating continuous-time process models in the GUI, see “Tutorial: Estimating Process Models Using the GUI” in *Getting Started with System Identification Toolbox*.

Data Supported by a Process Model

You can estimate low-order (up to third order), continuous-time transfer functions from data with the following characteristics:

- Time- or frequency-domain `iddata` or `idfrd` data object
- Real data, or complex data in the time domain only
- Single-output data

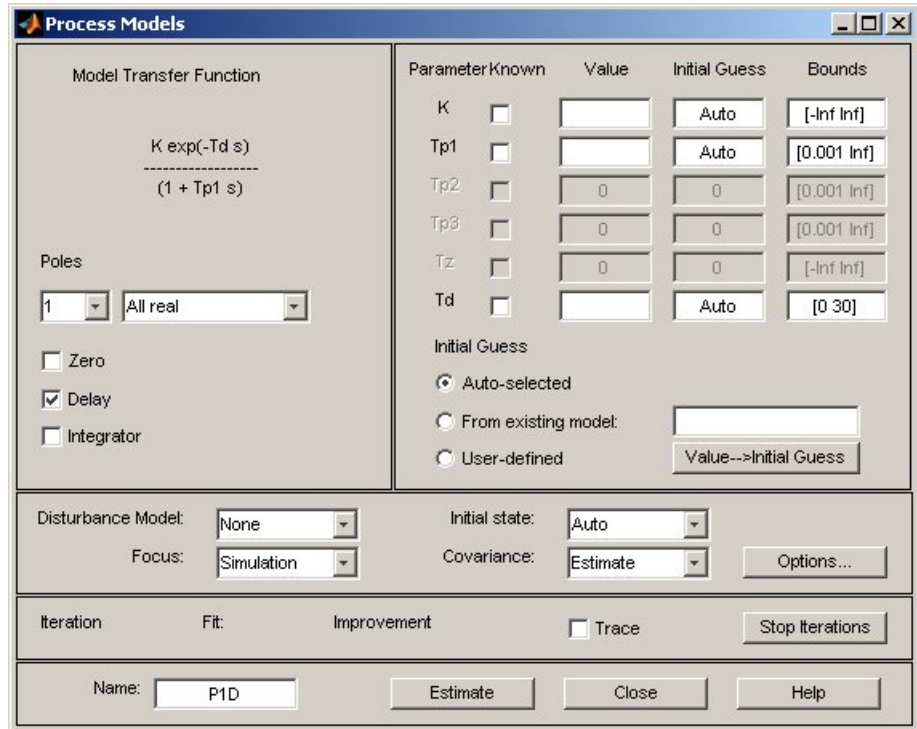
You must import your data into the MATLAB Workspace browser, as described in Chapter 2, “Preparing Data”.

How to Estimate Process Models Using the GUI

The following procedure assumes that you have already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

To estimate a low-order transfer function (process model) using the System Identification Tool GUI:

- 1 In the System Identification Tool GUI, select **Estimate > Process models** to open the Process Models dialog box.



- 2 If your model contains multiple inputs, select the input channel in the **Input** list. This list only appears when you have multiple inputs. For more information, see “Options for Multiple-Input Models” on page 5-16.
- 3 In the **Model Transfer Function** area, specify the model structure using the following options:

- Under **Poles**, select the number of poles, and then select All real or Underdamped.

Note You need at least two poles to allow underdamped modes (complex-conjugate pair).

- Select the **Zero** check box to include a zero, which is a numerator term other than a constant, or clear the check box to exclude the zero.
- Select the **Delay** check box to include a delay, or clear the check box to exclude the delay.
- Select the **Integrator** check box to include an integrator (self-regulating process), or clear the check box to exclude the integrator.

The **Parameter** area shows as many active parameters as you included in the model structure.

Note By default, the model **Name** is set to the acronym that reflects the model structure, as described in “Options for Specifying the Process-Model Structure” on page 5-15.

- 4** In the **Initial Guess** area, select Auto-selected to calculate the initial parameter values for the estimation. The **Initial Guess** column in the

Parameter table displays Auto. If you do not have a good guess for the parameter values, Auto works better than entering an ad hoc value.

Parameter	Known	Value	Initial Guess	Bounds
K	<input type="checkbox"/>		Auto	[-Inf Inf]
Tw	<input type="checkbox"/>		Auto	[0.001 Inf]
Zeta	<input type="checkbox"/>		Auto	[0.001 Inf]
Tp3	<input type="checkbox"/>	0	0	[0.001 Inf]
Tz	<input type="checkbox"/>	0	0	[-Inf Inf]
Td	<input type="checkbox"/>		Auto	[0 30]

Initial Guess

Auto-selected

From existing model:

User-defined

- 5 (Optional) If you approximately know a parameter value, enter this value in the **Initial Guess** column of the Parameter table. The estimation algorithm uses this value as a starting point. If you know a parameter value exactly, enter this value in the **Initial Guess** column, and also select the corresponding **Known** check box in the table to fix its value.

If you know the range of possible values for a parameter, enter these values into the corresponding **Bounds** field to help the estimation algorithm.

For example, the following figure shows that the delay value T_d is fixed at 2 s and is not estimated.

Parameter	Known	Value	Initial Guess	Bounds
K	<input type="checkbox"/>		Auto	[-Inf Inf]
Tw	<input type="checkbox"/>		Auto	[0.001 Inf]
Zeta	<input type="checkbox"/>		Auto	[0.001 Inf]
Tp3	<input type="checkbox"/>	0	0	[0.001 Inf]
Tz	<input type="checkbox"/>	0	0	[-Inf Inf]
Td	<input checked="" type="checkbox"/>	2	2	[0 30]

Initial Guess

Auto-selected
 From existing model:
 User-defined

Value-->Initial Guess

- 6 In the **Disturbance Model** list, select one of the available options. For more information about each option, see “Options for the Disturbance Model Structure” on page 5-17.
- 7 In the **Focus** list, select how to weigh the relative importance of the fit at different frequencies. For more information about each option, see “Options for Frequency-Weighing Focus” on page 5-18.
- 8 In the **Initial state** list, specify how you want the algorithm to treat initial states. For more information about the available options, see “Options for Initial States” on page 5-19.

Tip If you get a bad fit, you might try setting a specific method for handling initial states, rather than choosing it automatically.

- 9 In the **Covariance** list, select Estimate if you want the algorithm to compute parameter uncertainties. Effects of such uncertainties are displayed on plots as model confidence regions.

To omit estimating uncertainty, select **None**. Skipping uncertainty computation might reduce computation time for complex models and large data sets.

- 10** In the **Model Name** field, edit the name of the model or keep the default. The name of the model should be unique in the Model Board.
- 11** To view the estimation progress in the MATLAB Command Window, select the **Trace** check box. During estimation, the following information is displayed for each iteration:
 - **Loss command** — Equals the determinant of the estimated covariance matrix of the input noise.
 - **Parameter values** — Values of the model structure coefficients you specified.
 - **Search direction** — Change in parameter values from the previous iteration.
 - **Fit improvements** — Shows the actual versus expected improvements in the fit.
- 12** Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
- 13** To stop the search and save the results after the current iteration has been completed, click **Stop Iterations**. To continue iterations from the current model, click the **Continue iter** button to assign current parameter values as initial guesses for the next search.
- 14** To plot the model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. For more information about validating models, see “Overview of Model Validation and Plots” on page 10-3.
- 15** To refine the current estimate, click the **Value —> Initial Guess** button to assign current parameter values as initial guesses for the next search, edit the **Model Name** field, and click **Estimate**.

If your model is not sufficiently accurate, try another model structure.

You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

Estimating Process Models at the Command Line

- “Using pem to Estimate Process Models” on page 5-9
- “Example – Estimating Process Models with Free Parameters at the Command Line” on page 5-10
- “Example – Estimating Process Models with Fixed Parameters at the Command Line” on page 5-12

Using pem to Estimate Process Models

You can estimate process models using the iterative estimation method pem that minimizes the prediction errors to obtain maximum likelihood estimates. The resulting models are stored as idproc model objects.

You can use the following general syntax to both configure and estimate process models:

```
m = pem(data,mod_struct,'Property1',Value1,...,  
        'PropertyN',ValueN)
```

data is the estimation data and mod_struct is a string that represents the process model structure, as described in “Options for Specifying the Process-Model Structure” on page 5-15.

Tip You do not need to construct the model object using idproc before estimation unless you want to specify initial parameter guesses or fixed parameter values, as described in “Example – Estimating Process Models with Fixed Parameters at the Command Line” on page 5-12.

The property-value pairs specify any model properties that configure the estimation algorithm and the initial conditions. For more information about accessing and setting model properties, see “Model Properties” on page 3-13.

Note You can specify all property-value pairs in pem as a simple, comma-separated list without worrying about the hierarchy of these properties in the idproc model object.

For more information about validating a process model, see “Overview of Model Validation and Plots” on page 10-3.

You can use pem to refine parameter estimates of an existing process model, as described in “Refining Linear Parametric Models” on page 5-78.

For detailed information about pem and idproc, see the corresponding reference page.

Example – Estimating Process Models with Free Parameters at the Command Line

This example demonstrates how to estimate the parameters of a first-order process model:

$$G(s) = \frac{K}{1 + sT_{p1}} e^{-sT_d}$$

This process has two inputs and the response from each input is estimated by a first-order process model. All parameters are free to vary.

Use the following commands to estimate a model m from sample data:

```
% Load sample data
load co2data
% Sampling interval is 0.5 min (known)
Ts = 0.5;
% Split data set into estimation data ze
% and validation data zv
ze = iddata(Output_exp1,Input_exp1,Ts,...
            'TimeUnit','min');
zv = iddata(Output_exp2,Input_exp2,Ts,...
            'TimeUnit','min');
% Estimate model with one pole and a delay
m = pem(ze,'P1D')
```

MATLAB responds with the following output:

```
Process model with 2 inputs:
y = G_1(s)u_1 + G_2(s)u_2
where
      K
G_1(s) = ----- * exp(-Td*s)
          1+Tp1*s

with   K = -3.2168
      Tp1 = 23.033
      Td = 10.101

      K
G_2(s) = ----- * exp(-Td*s)
          1+Tp1*s

with   K = 9.9877
      Tp1 = 2.0314
      Td = 4.8368
```

Use dot notation to get the value of any model parameter. For example, to get the Value field in the K structure, type the following command:

```
m.K.value
```

Example – Estimating Process Models with Fixed Parameters at the Command Line

When you know the values of certain parameters in the model and want to estimate only the values you do not know, you must specify the fixed parameters after creating the `idproc` model object.

Use the following commands to prepare the data and construct a process model with one pole and a delay:

```
% Load sample data
load co2data
% Sampling interval is 0.5 min (known)
Ts = 0.5;
% Split data set into estimation data ze
% and validation data zv
ze = iddata(Output_exp1,Input_exp1,Ts,...
            'TimeUnit','min');
zv = iddata(Output_exp2,Input_exp2,Ts,...
            'TimeUnit','min');
mod=idproc('P1D')
```

MATLAB responds with the following output:

```
Process model with transfer command
      K
      |
G(s) = ----- * exp(-Td*s)
      1+Tp1*s

with  K = NaN
      Tp1 = NaN
      Td = NaN
```

This model was not estimated from data.

The model parameters `K`, `Tp1`, and `Td` are assigned `NaN` values, which means that the parameters have not yet been estimated from the data.

All process-model parameters are structures with the following fields:

- `status` field specifies whether to estimate the parameter, or keep the initial value fixed (do not estimate), or set the value to zero. This field can have the values 'estimate', 'fixed', or 'zero'. For more information, see “Options for Initial States” on page 5-19.
- `min` specifies the minimum bound on the parameter.
- `max` specifies the maximum bound on the parameter.
- `value` specifies the numerical value of the parameter, if known.

To set the value of `K` to 12 and keep it fixed, use the following commands:

```
mod.K.value=12;  
mod.K.status='fixed';
```

Note `mod` is defined for one input. This model is automatically adjusted to have a duplicate for each input.

To estimate `Tp1` and `Td` only, use the following command:

```
mod_proc=pem(ze,mod)
```

MATLAB returns the following result:

Process model with 2 inputs:

$$y = G_1(s)u_1 + G_2(s)u_2$$

where

$$G_1(s) = \frac{K}{1+Tp1*s} * \exp(-Td*s)$$

with $K = 12$
 $Tp1 = 7.0998e+007$
 $Td = 15$

$$G_2(s) = \frac{K}{1+Tp1*s} * \exp(-Td*s)$$

with $K = 12$
 $Tp1 = 3.6962$
 $Td = 3.817$

In this case, the value of K is fixed at 12, but Tp1 and Td are estimated.

If you prefer to specify parameter constraints directly in the estimator syntax, the following table provides examples of pem commands.

Action	Examples
Fix the value of K to 12.	<code>m=pem(ze,'p1d','k','fix','k',12)</code>
Initialize K for the iterative search without fixing this value.	<code>m=pem(ze,'p1d','k',12)</code>
Constrain the value of K between 3 and 4.	<code>m=pem(ze,'p1d','k',... {'min',3},'k',{'max',4})</code>

Options for Specifying the Process-Model Structure

This section describes how to specify the model structure in the estimation procedures “How to Estimate Process Models Using the GUI” on page 5-3 and “Estimating Process Models at the Command Line” on page 5-9.

In the System Identification Tool GUI. Specify the model structure by selecting the number of real or complex poles, and whether to include a zero, delay, and integrator. The resulting transfer command is displayed in the Process Models dialog box.

At the command line. Specify the model structure using an acronym that includes the following letters and numbers:

- (Required) P for a process model
- (Required) 0, 1, 2 or 3 for the number of poles
- (Optional) D to include a time-delay term e^{-sT_d}
- (Optional) Z to include a process zero (numerator term)
- (Optional) U to indicate possible complex-valued (underdamped) poles
- (Optional) I to indicate enforced integration

Typically, you specify the model-structure acronym as a string argument in the estimation command pem:

- pem(data, 'P1D') to estimate the following structure:

$$G(s) = \frac{K}{1 + sT_{p1}} e^{-sT_d}$$

- pem(data, 'P2ZU') to estimate the following structure:

$$G(s) = \frac{K_p (1 + sT_z)}{1 + 2s\zeta T_w + s^2 T_w^2}$$

- pem(data, 'POID') to estimate the following structure:

$$G(s) = \frac{K_p}{s} e^{-sT_d}$$

- pem(data, 'P3Z') to estimate the following structure:

$$G(s) = \frac{K_p (1 + sT_z)}{(1 + sT_{p1})(1 + sT_{p2})(1 + sT_{p3})}$$

For more information about estimating models, see “Estimating Process Models at the Command Line” on page 5-9.

Options for Multiple-Input Models

If your model contains multiple inputs, you can specify whether to estimate the same transfer command for all inputs, or a different transfer command for each input. The information in this section supports the estimation procedures “How to Estimate Process Models Using the GUI” on page 5-3 and “Estimating Process Models at the Command Line” on page 5-9.

In the System Identification Tool GUI. To fit a data set with multiple inputs in the Process Models dialog box, configure the process model settings for one input at a time. When you finish configuring the model and the

estimation settings for one input, select a different input in the **Input Number** list.

If you want the same transfer command to apply to all inputs, select the **Same structure for all channels** check box. To apply a different structure to each channel, leave this check box clear, and create a different transfer command for each input.

At the command line. Specify the model structure as a cell array of acronym strings in the estimation command pem. For example, use this command to specify the first-order transfer command for the first input, and a second-order model with a zero and an integrator for the second input:

```
m = idproc({'P1','P2ZI'})
m = pem(data,m)
```

To apply the same structure to all inputs, define a single structure in idproc.

Options for the Disturbance Model Structure

This section describes how to specify a noise model in the estimation procedures “How to Estimate Process Models Using the GUI” on page 5-3 and “Estimating Process Models at the Command Line” on page 5-9.

In addition to the transfer command G , a linear system can include an additive noise term He , as follows:

$$y = Gu + He$$

where e is white noise.

You can estimate only the dynamic model G , or estimate both the dynamic model and the disturbance model H . For process models, H is a rational transfer command C/D , where the C and D polynomials for a first- or second-order ARMA model.

In the GUI. To specify whether to include or exclude a noise model in the Process Models dialog box, select one of the following options from the **Disturbance Model** list:

- None — The algorithm does not estimate a noise model ($C=D=1$). This option also sets **Focus** to Simulation.
- Order 1 — Estimates a noise model as a continuous-time, first-order ARMA model.
- Order 2 — Estimates a noise model as a continuous-time, second-order ARMA model.

At the command line. Specify the disturbance model as an argument in the estimation command `pem`. For example, use this command to estimate a first-order transfer command and a first-order noise model:

```
pem(data, 'P1D', 'DisturbanceModel', 'ARMA1')
```

Tip You can type 'dis' instead of 'DisturbanceModel'.

For a complete list of values for the `DisturbanceModel` model property, see the `idproc` reference page.

Options for Frequency-Weighing Focus

You can specify how the estimation algorithm weighs the fit at various frequencies. This information supports the estimation procedures “How to Estimate Process Models Using the GUI” on page 5-3 and “Estimating Process Models at the Command Line” on page 5-9.

In the System Identification Tool GUI. Set **Focus** to one of the following options:

- Prediction — Uses the inverse of the noise model H to weigh the relative importance of how closely to fit the data in various frequency ranges. Corresponds to minimizing one-step-ahead prediction, which typically favors the fit over a short time interval. Optimized for output prediction applications.
- Simulation — Uses the input spectrum to weigh the relative importance of the fit in a specific frequency range. Does not use the noise model to weigh the relative importance of how closely to fit the data in various frequency ranges. Optimized for output simulation applications.

- **Stability** — Behaves the same way as the **Prediction** option, but also forces the model to be stable. For more information about model stability, see “Unstable Models” on page 10-70.
- **Filter** — Specify a custom filter to open the Estimation Focus dialog box, where you can enter a filter, as described in “Simple Passband Filter” on page 2-91 or “Defining a Custom Filter” on page 2-92. This prefiltering applies only for estimating the dynamics from input to output. The disturbance model is determined from the estimation data.

At the command line. Specify the focus as an argument in the estimation command `pem` using the same options as in the GUI. For example, use this command to optimize the fit for simulation and estimate a disturbance model:

```
pem(data, 'P1D', 'dist', 'arma2', 'Focus', 'Simulation')
```

Options for Initial States

Because the process models are dynamic, you need initial states that capture past input properties. Thus, you must specify how the iterative algorithm treats initial states. This information supports the estimation procedures “How to Estimate Process Models Using the GUI” on page 5-3 and “Estimating Process Models at the Command Line” on page 5-9.

In the System Identification Tool GUI. Set **Initial state** to one of the following options:

- **Zero** — Sets all initial states to zero.
- **Estimate** — Treats the initial states as an unknown vector of parameters and estimates these states from the data.
- **Backcast** — Estimates initial states using a backward filtering method (least-squares fit).
- **U-level est** — Estimates both the initial states and the `InputLevel` model property that represents the input offset level. For multiple inputs, the input level for each input is estimated individually. Use if you included an integrator in the transfer command.
- **Auto** — Automatically chooses one of the preceding options based on the estimation data. If initial states have negligible effect on the prediction errors, the initial states are set to zero to optimize algorithm performance.

At the command line. Specify the initial states as an argument in the estimation command `pem` using the same options as in the GUI. For example, use this command to estimate a first-order transfer command and set the initial states to zero:

```
m=pem(data,'P1D','InitialState','zero')
```

For a complete list of values for the `InitialState` model property, see the `idproc` reference page.

Estimating Polynomial Models

In this section...

“What Are Black-Box Polynomial Models?” on page 5-21

“Data Supported by Polynomial Models” on page 5-28

“Preliminary Step – Estimating Model Orders and Input Delays ” on page 5-29

“How to Estimate Polynomial Models in the GUI” on page 5-37

“How to Estimate Polynomial Models at the Command Line” on page 5-40

“Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44

“Option for Frequency-Weighing Focus” on page 5-45

“Options for Initial States” on page 5-46

“Algorithms for Estimating Polynomial Models” on page 5-46

What Are Black-Box Polynomial Models?

- “Polynomial Model Structure” on page 5-22
- “Understanding the Time-Shift Operator q ” on page 5-23
- “Definition of a Discrete-Time Polynomial Model” on page 5-23
- “Definition of a Continuous-Time Polynomial Model” on page 5-26
- “Definition of Multiple-Output ARX Models” on page 5-26

Polynomial Model Structure

System Identification Toolbox supports linear polynomial model structures that have the following general representation:

$$A(q)y(t) = \sum_{i=1}^{nu} \frac{B_i(q)}{F_i(q)} u_i(t - nk_i) + \frac{C(q)}{D(q)} e(t)$$

The polynomials A , B_i , C , D , and F_i contain the time-shift operator q . u_i is the i th input, nu is the total number of inputs, and nk_i is the i th input delay that characterizes the delay response time. The variance of the white noise $e(t)$ is assumed to be λ . For more information about the time-shift operator, see “Understanding the Time-Shift Operator q ” on page 5-23.

Note This form is completely equivalent to the Z-transform form: q corresponds to z .

To estimate polynomial models, you must specify the *model order* as a set of integers that represent the number of coefficients for each polynomial you include in your selected structure— na for A , nb for B , nc for C , nd for D , and nf for F . You must also specify the number of samples nk corresponding to the input delay—*dead time*—given by the number of samples before the output responds to the input.

The number of coefficients in denominator polynomials is equal to the number of poles, and the number of coefficients in the numerator polynomials is equal to the number of zeros plus 1. When the dynamics from $u(t)$ to $y(t)$ contain a delay of nk samples, then the first nk coefficients of B are zero.

For more information about the family of transfer-command models, see the corresponding section in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

Understanding the Time-Shift Operator q

The general polynomial equation is written in terms of the time-shift operator q . To understand this time-shift operator, consider the following discrete-time difference equation:

$$y(t) + a_1 y(t - T) + a_2 y(t - 2T) = b_1 u(t - T) + b_2 u(t - 2T)$$

where $y(t)$ is the output, $u(t)$ is the input, and T is the sampling interval. q^{-1} is a time-shift operator that compactly represents such difference equations using $qu(t) = u(t - T)$:

$$y(t) + a_1 q^{-1} y(t) + a_2 q^{-2} y(t) = b_1 q^{-1} u(t) + b_2 q^{-2} u(t)$$

or

$$A(q)y(t) = B(q)u(t)$$

In this case, $A(q) = 1 + a_1 q^{-1} + a_2 q^{-2}$ and $B(q) = b_1 q^{-1} + b_2 q^{-2}$.

Note This q description is completely equivalent to the Z -transform form: q corresponds to z .

Definition of a Discrete-Time Polynomial Model

These model structures are subsets of the following general polynomial equation:

$$A(q)y(t) = \sum_{i=1}^{nu} \frac{B_i(q)}{F_i(q)} u_i(t - nk_i) + \frac{C(q)}{D(q)} e(t)$$

The model structures differ by how many of these polynomials are included in the structure. Thus, different model structures provide varying levels of flexibility for modeling the dynamics and noise characteristics. For more

information about the time-shift operator, see “Understanding the Time-Shift Operator q ” on page 5-23.

The following table summarizes common linear polynomial model structures supported by System Identification Toolbox. If you have a specific structure in mind for your application, you can decide whether the dynamics and the noise have common or different poles. $A(q)$ corresponds to poles that are common for the dynamic model and the noise model. Using common poles for dynamics and noise is useful when the disturbances enter the system at the input. F_i determines the poles unique to the system dynamics, and D determines the poles unique to the disturbances.

Model Structure	Discrete-Time Form	Noise Model
ARX	$A(q)y(t) = \sum_{i=1}^{nu} B_i(q)u_i(t - nk_i) + e(t)$	The noise model is $\frac{1}{A}$ and the noise is coupled to the dynamics model. ARX does not let you model noise and dynamics independently. Use ARX to have a simple model at good signal-to-noise ratios.
ARMAX	$A(q)y(t) = \sum_{i=1}^{nu} B_i(q)u_i(t - nk_i) + C(q)e(t)$	Extends the ARX structure by providing more flexibility for modeling noise using the C parameters (a M oving A verage of white noise). Use ARMAX when the dominating disturbances enter at the input. Such disturbances are called <i>load disturbances</i> .

Model Structure	Discrete-Time Form	Noise Model
Box-Jenkins (BJ)	$y(t) = \sum_{i=1}^{nu} \frac{B_i(q)}{F_i(q)} u_i(t - nk_i) + \frac{C(q)}{D(q)} e(t)$	Provides completely independent parameterization for the dynamics and the noise using rational polynomial functions. Use BJ models when the noise does not enter at the input, but is primary a measurement disturbance. This structure provides additional flexibility for modeling noise.
Output-Error (OE)	$y(t) = \sum_{i=1}^{nu} \frac{B_i(q)}{F_i(q)} u_i(t - nk_i) + e(t)$	Use when you want to parameterize dynamics, but do not want to estimate a noise model. <hr/> <p>Note In this case, the noise models is $H = 1$ in the general equation and the white noise source $e(t)$ affects only the output.</p> <hr/>

The System Identification Tool GUI supports only the polynomial models listed in the table. However, you can use pem to estimate all five polynomial or any subset of polynomials in the general equation. For more information about working with pem, see “Using pem to Estimate Polynomial Models” on page 5-42.

Definition of a Continuous-Time Polynomial Model

In continuous time, the general frequency-domain equation is written in terms of the Laplace transform variable s , which corresponds to a differentiation operation:

$$A(s)Y(s) = \frac{B(s)}{F(s)}U(s) + \frac{C(s)}{D(s)}E(s)$$

In the continuous-time case, the underlying time-domain model is a differential equation and the model order integers represent the number of estimated numerator and denominator coefficients. For example, $n_a=3$ and $n_b=2$ correspond to the following model:

$$\begin{aligned}A(s) &= s^4 + a_1s^3 + a_2s^2 + a_3 \\ B(s) &= b_1s + b_2\end{aligned}$$

The simplest way to estimate continuous-time polynomial models of arbitrary structure is to first estimate a discrete-time model of arbitrary order and then use `d2c` to convert this model to continuous time. For more information, see “Transforming Between Discrete-Time and Continuous-Time Representations” on page 11-2.

You can also use System Identification Toolbox to estimate continuous-time polynomial models directly using continuous-time frequency-domain data. In this case, you must set the `Ts` data property to 0 to indicate that you have continuous-time frequency-domain data.

Definition of Multiple-Output ARX Models

You can use a multiple-output ARX model to model a multiple-output dynamic system. The ARX model structure is given by the following equation:

$$A(q)y(t) = B(q)u(t - nk) + e(t)$$

For a system with nu inputs and ny outputs, $A(q)$ is an ny -by- ny matrix. $A(q)$ can be represented as a polynomial in the shift operator q^{-1} :

$$A(q) = I_{ny} + A_1 q^{-1} + \dots + A_{na} q^{-na}$$

For more information about the time-shift operator, see “Understanding the Time-Shift Operator q ” on page 5-23.

$A(q)$ can also be represented as a matrix:

$$A(q) = \begin{pmatrix} a_{11}(q) & a_{12}(q) & \dots & a_{1ny}(q) \\ a_{21}(q) & a_{22}(q) & \dots & a_{2ny}(q) \\ \dots & \dots & \dots & \dots \\ a_{ny1}(q) & a_{ny2}(q) & \dots & a_{nyny}(q) \end{pmatrix}$$

where the matrix element a_{kj} is a polynomial in the shift operator q^{-1} :

$$a_{kj}(q) = \delta_{kj} + a_{kj}^1 q^{-1} + \dots + a_{kj}^{na_{kj}} q^{-na_{kj}}$$

δ_{kj} represents the Kronecker delta, which equals 1 for $k=j$ and equals 0 for $k \neq j$. This polynomial describes how the old values of the j th output are affected by the k th output. The i th row of $A(q)$ represents the contribution of the past output values for predict the current value of the i th output.

$B(q)$ is an ny -by- ny matrix. $B(q)$ can be represented as a polynomial in the shift operator q^{-1} :

$$B(q) = B_0 + B_1 q^{-1} + \dots + B_{nb} q^{-nb}$$

$B(q)$ can also be represented as a matrix:

$$B(q) = \begin{pmatrix} b_{11}(q) & b_{12}(q) & \dots & b_{1nu}(q) \\ b_{21}(q) & b_{22}(q) & \dots & b_{2nu}(q) \\ \dots & \dots & \dots & \dots \\ b_{ny1}(q) & b_{ny2}(q) & \dots & b_{nynu}(q) \end{pmatrix}$$

where the matrix element b_{kj} is a polynomial in the shift operator q^{-1} :

$$b_{kj}(q) = a_{kj}^1 q^{-nb_{kj}} + \dots + a_{kj}^{nk_{kj}} q^{-nb_{kj} - nb_{kj} + 1}$$

nk_{kj} is the delay from the j th input to the k th output. $B(q)$ represents the contributions of inputs to predicting all output values.

Data Supported by Polynomial Models

- “Types of Supported Data” on page 5-28
- “Designating Data for Estimating Continuous-Time Models” on page 5-29
- “Designating Data for Estimating Discrete-Time Models” on page 5-29

Types of Supported Data

You can estimate linear, black-box polynomial models from data with the following characteristics:

- Time- or frequency-domain data (iddata or idfrd data objects).

Note For frequency-domain data, you can only estimate ARX and OE models.

To estimate black-box polynomial models for time-series data, see Chapter 6, “Estimating Time-Series Models”.

- Real data or complex data in any domain.

- Single-output and multiple-output.

You must import your data into the MATLAB Workspace browser, as described in Chapter 2, “Preparing Data”.

Designating Data for Estimating Continuous-Time Models

To get a linear, continuous-time model of arbitrary structure for time-domain data, you can estimate a discrete-time model, and then use `d2c` to transform it to a continuous-time model.

For continuous-time frequency-domain data, you can estimate directly only the ARX and Output-Error (OE) continuous-time models. Other structures include noise models, which is not supported for frequency-domain data.

Tip To denote continuous-time frequency-domain data, set the data sampling interval to 0. You can set the sampling interval when you import data into the GUI or set the `Ts` property of the data object at the command line.

Designating Data for Estimating Discrete-Time Models

You can estimate arbitrary-order, linear state-space models for both time- or frequency-domain data.

Your data must have the data property `Ts` set to the experimental data sampling interval.

Tip You can set the sampling interval when you import data into the GUI or set the `Ts` property of the data object at the command line.

Preliminary Step – Estimating Model Orders and Input Delays

- “Why Estimate Model Orders and Delays?” on page 5-30
- “Estimating Orders and Delays in the GUI” on page 5-30

- “Estimating Model Orders at the Command Line” on page 5-33
- “Estimating Delays at the Command Line” on page 5-35
- “Selecting Model Orders from the Best ARX Structure” on page 5-35

Why Estimate Model Orders and Delays?

To estimate polynomial models, you must provide input delays and model orders. If you already have insight into the physics of your system, you can specify the number of poles and zeros.

In most cases, you do not know the model orders in advance. To get initial model orders and delays for your system, you can estimate several ARX models with a range of orders and delays and compare the performance of these models. You choose the model orders that correspond to the best model performance and use these orders as an initial guess for further modeling.

Because this estimation procedure uses the ARX model structure, which includes the A and B polynomials, you only get estimates for the na , nb , and nk parameters. However, you can use these results as initial guesses for the corresponding polynomial orders and input delays in other model structures, such as ARMAX, OE, and BJ.

If the estimated nk is too small, the leading nb coefficients are much smaller than their standard deviations. Conversely, if the estimated nk is too large, there is a significant correlation between the residuals and the input for lags that correspond to the missing B terms. For information about residual analysis plots, see “Using Residual Analysis Plots to Validate Models” on page 10-15.

Estimating Orders and Delays in the GUI

The following procedure assumes that you have already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

To estimate model orders and input delays in the System Identification Tool GUI:

- 1 In the System Identification Tool GUI, select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box.


The ARX model is already selected by default in the **Structure** list.

Note For time-series models, select the AR model structure.

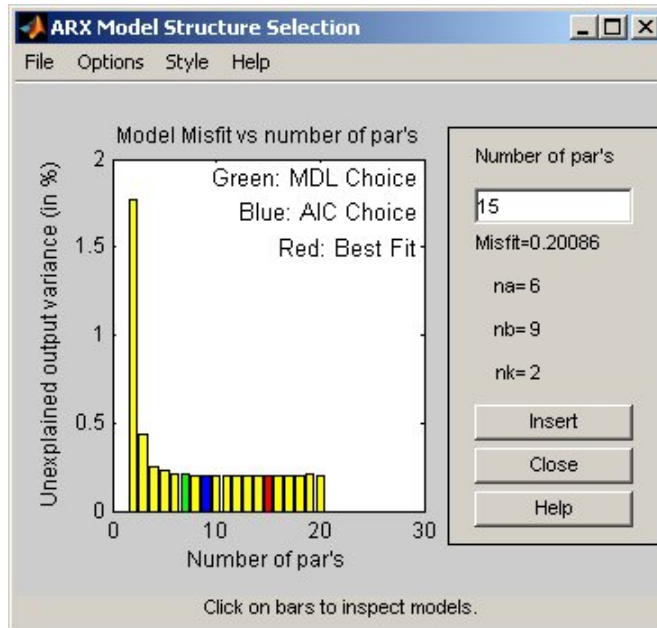
- 2 Edit the **Orders** field to specify a range of poles, zeros, and delays. For example, enter the following values for na , nb , and nk :

[1:10 1:10 1:10]

Tip As a shortcut for entering 1:10 for each required model order, click **Order Selection**.



- 3 Click **Estimate** to open the ARX Model Structure Selection window, which displays the model performance for each combination of model parameters. The following figure shows an example plot.



- 4 Select a rectangle that represents the optimum parameter combination and click **Insert** to estimate a model with these parameters. For information about using this plot, see “Selecting Model Orders from the Best ARX Structure” on page 5-35.

This action adds a new model to the Model Board in the System Identification Tool GUI. The default name of the parametric model contains the model type and the number of poles, zeros, and delays. For example, arx692 is an ARX model with $n_a=6$, $n_b=9$, and a delay of two samples.

- 5 Click **Close** to close the ARX Model Structure Selection window.

After estimating model orders and delays, use these values as initial guesses for estimating other model structures, as described in “How to Estimate Polynomial Models in the GUI” on page 5-37.

Estimating Model Orders at the Command Line

You can estimate model orders using the `struc`, `arxstruc`, and `selstruc` commands in combination.

If you are working with a multiple-output system, you must use `struc`, `arxstruc`, and `selstruc` commands for each output. In this case, you must subreference the correct output channel in your estimation and validation data sets.

For each estimation, you use two independent data sets—an estimation data set and a validation data set. These independent data set can be from different experiments, or you can select these data sets from a single experiment. For more information about subreferencing data, see “Subreferencing `iddata` Objects” on page 2-40 and “Subreferencing `idfrd` Objects” on page 2-55.

For an example of estimating model orders for a multiple-input system, see “Estimating Delays in the System” in *Getting Started with System Identification Toolbox*.

struc. The `struc` command creates a matrix of possible model-order combinations for a specified range of n_a , n_b , and n_k values.

For example, the following command defines the range of model orders and delays `na=2:5`, `nb=1:5`, and `nk=1:5`:

```
NN = struc(2:5,1:5,1:5)
```

Note `struc` applies only to single-input/single-output models. If you have multiple inputs and want to use `struc`, apply this command to one input-output pair at a time.

arxstruc. The `arxstruc` command takes the output from `struc`, estimates an ARX model for each model order, and compares the model output to the measured output. `arxstruc` returns the *loss* for each model, which is the normalized sum of squared prediction errors.

For example, the following command uses the range of specified orders `NN` to compute the loss command for single-input/single-output estimation data `data_e` and validation data `data_v`:

```
V = arxstruc(data_e,data_v,NN)
```

Each row in `NN` corresponds to one set of orders:

```
[na nb nk]
```

selstruc. The `selstruc` command takes the output from `arxstruc` and opens the ARX Model Structure Selection window to guide your choice of the model order with the best performance.

For example, to open the ARX Model Structure Selection window and interactively choose the optimum parameter combination, use the following command:

```
selstruc(V)
```

For more information about working with the ARX Model Structure Selection window, see “Selecting Model Orders from the Best ARX Structure” on page 5-35.

To find the structure that minimizes Akaike’s Information Criterion, use the following command:

```
nn = selstruc(V, 'AIC')
```

where `nn` contains the corresponding `na`, `nb`, and `nk` orders.

Similarly, to find the structure that minimizes the Rissanen’s Minimum Description Length (MDL), use the following command:

```
nn = selstruc(V, 'MDL')
```

To select the structure with the smallest loss command, use the following command:

```
nn = selstruc(V,0)
```

After estimating model orders and delays, use these values as initial guesses for estimating other model structures, as described in “Using pem to Estimate Polynomial Models” on page 5-42.

Estimating Delays at the Command Line

The `delayest` command estimates the time delay in a dynamic system by estimating a low-order, discrete-time ARX model and treating the delay as an unknown parameter.

By default, `delayest` assumes that $n_a=n_b=2$ and that there is a good signal-to-noise ratio, and uses this information to estimate n_k .

To estimate the delay for a data set `data`, type the following at the MATLAB prompt:

```
delayest(data)
```

If your data has a single input, MATLAB responds with a scalar value for the input delay—equal to the number of data samples. If your data has multiple inputs, MATLAB returns a vector, where each value is the delay for the corresponding input signal.

To compute the actual delay time, you must multiply the input delay by the sampling interval of the data.

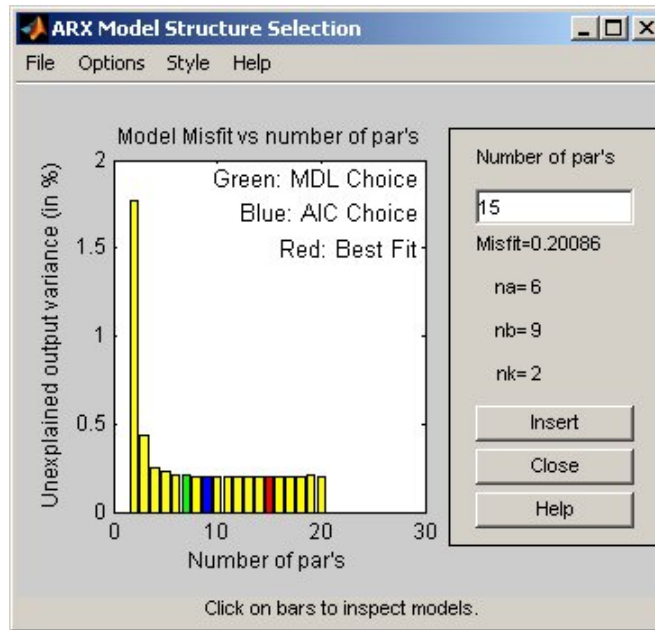
You can also use the ARX Model Structure Selection window to estimate input delays and model order together, as described in “Estimating Model Orders at the Command Line” on page 5-33.

Selecting Model Orders from the Best ARX Structure

You generate the ARX Model Structure Selection window for your data to select the best-fit model.

For a procedure on generating this plot in the System Identification Tool GUI, see “Estimating Orders and Delays in the GUI” on page 5-30. To open this plot at the command line, see “Estimating Model Orders at the Command Line” on page 5-33.

The following figure shows a sample plot in the ARX Model Structure Selection window.



The horizontal axis in the ARX Model Structure Selection window is the total number of ARX parameters:

$$\text{Number of parameters} = n_a + n_b$$

The vertical axis, called **Unexplained output variance (in %)**, is the ARX model prediction error for a specific number of parameters. The *prediction error* is the sum of the squares of the differences between the validation data output and the model output. In other words, **Unexplained output variance (in %)** is the portion of the output not explained by the model.

Three rectangles are highlighted on the plot—green, blue, and red. Each color indicates a type of best-fit criterion, as follows:

- Red minimizes the sum of the squares of the difference between the validation data output and the model output. This option is considered the overall best fit.
- Green minimizes Rissanen MDL criterion.
- Blue minimizes Akaike AIC criterion.

In the ARX Model Structure Selection window, click any bar to view the orders that give the best fit. The area on the right is dynamically updated to show the orders and delays that give the best fit.

For more information about the AIC criterion, see “Using Akaike’s Criteria to Validate Models” on page 10-58.

How to Estimate Polynomial Models in the GUI

- “Before You Begin” on page 5-37
- “Estimating Polynomial Models in the GUI” on page 5-37

Before You Begin

Before you estimate polynomial models, you must have already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

This procedure also requires that you select a model structure and specify model orders and delays. For more information about how to estimate model orders and delays, see “Estimating Orders and Delays in the GUI” on page 5-30.

If you are estimating a multiple-output ARX model, you must specify order matrices in the MATLAB Workspace browser before estimation, as described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44.

Estimating Polynomial Models in the GUI

To estimate a polynomial model in the System Identification Tool GUI.

- 1 In the System Identification Tool GUI, select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box.
- 2 In the **Structure** list, select the polynomial model structure you want to estimate from the following options:
 - ARX: [na nb nk]
 - ARMAX: [na nb nc nk]
 - OE: [nb nf nk]
 - BJ: [nb nc nd nf nk]

This action updates the options in the Linear Parametric Models dialog box to correspond with this model structure. For information about each model structure, see “What Are Black-Box Polynomial Models?” on page 5-21.

Note For time-series data, only AR and ARMA models are available. For more information about estimating time-series models, see Chapter 6, “Estimating Time-Series Models”.

- 3 In the **Orders** field, specify the model orders and delays, as follows:
 - **For single-output polynomial models.** Enter the model orders and delays according to the sequence displayed in the **Structure** field. For multiple-input models, specify nb and nk as row vectors with as many elements as there are inputs. If you are estimating BJ and OE models, you must also specify nf as a vector.

For example, for a three-input system, nb can be [1 2 4], where each element corresponds to an input.
 - **For multiple-output ARX models.** Enter the model orders, as described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44.

Tip To enter model orders and delays using the Order Editor dialog box, click **Order Editor**.

- 4** (ARX models only) Select the estimation **Method** as **ARX** or **IV** (instrumental variable method). For information about the algorithms, see “Algorithms for Estimating Polynomial Models” on page 5-46.
- 5** In the **Name** field, edit the name of the model or keep the default. The name of the model should be unique in the Model Board.
- 6** In the **Focus** list, select how to weigh the relative importance of the fit at different frequencies. For more information about each option, see “Option for Frequency-Weighing Focus” on page 5-45.
- 7** In the **Initial state** list, specify how you want the algorithm to treat initial states. For more information about the available options, see “Options for Initial States” on page 5-19.

Tip If you get an inaccurate fit, try setting a specific method for handling initial states rather than choosing it automatically.

- 8** In the **Covariance** list, select **Estimate** if you want the algorithm to compute parameter uncertainties. Effects of such uncertainties are displayed on plots as model confidence regions.

To omit estimating uncertainty, select **None**. Skipping uncertainty computation for large, multiple-output ARX models might reduce computation time.

- 9** (ARMAX, OE, and BJ models only) To view the estimation progress in the MATLAB Command Window, select the **Trace** check box. During estimation, the following information is displayed for each iteration:
 - **Loss command** — Equals the determinant of the estimated covariance matrix of the input noise.
 - **Parameter values** — Values of the model structure coefficients you specified.
 - **Search direction** — Change in parameter values from the previous iteration.
 - **Fit improvements** — Shows the actual versus expected improvements in the fit.

- 10** Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
- 11** (Prediction-error method only) To stop the search and save the results after the current iteration has been completed, click **Stop Iterations**. To continue iterations from the current model, click the **Continue iter** button to assign current parameter values as initial guesses for the next search.
- 12** To plot the model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. For more information about validating models, see “Overview of Model Validation and Plots” on page 10-3.

If you get an inaccurate fit, try estimating a new model with different orders or structure. You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

How to Estimate Polynomial Models at the Command Line

- “Using `arx` and `iv4` to Estimate ARX Models” on page 5-40
- “Using `pem` to Estimate Polynomial Models” on page 5-42

Using `arx` and `iv4` to Estimate ARX Models

You can estimate single-output and multiple-output ARX models using the `arx` and `iv4` commands. For information about the algorithms, see “Algorithms for Estimating Polynomial Models” on page 5-46.

If you are estimating a multiple-output ARX model, you must specify order matrices in the MATLAB Workspace browser before estimation, as described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44.

For single-output data, the `arx` and `iv4` commands produce an `idpoly` model object, and for multiple-output data these commands produce an `idarx` model object.

You can use the following general syntax to both configure and estimate ARX models:

```
% Using ARX method
m = arx(data,[na nb nk], 'Property1',Value1,...,
         'PropertyN',ValueN)

% Using IV method
m = iv4(data,[na nb nk], 'Property1',Value1,...,
         'PropertyN',ValueN)
```

`data` is the estimation data and `[na nb nk]` specifies the model orders, as discussed in “What Are Black-Box Polynomial Models?” on page 5-21.

The property-value pairs specify any model properties that configure the estimation algorithm and the initial conditions. For more information about accessing and setting model properties, see “Model Properties” on page 3-13.

Note You can specify all property-value pairs as a comma-separated list.

To get discrete-time models, use the time-domain data (`iddata` object). To get a single-output continuous-time model, apply `d2c` to a discrete-time model or use continuous-time frequency-domain data—either `idfrd` object, or frequency-domain `iddata` with `Ts=0`.

Note System Identification Toolbox does not support multiple-output continuous-time `idarx` models.

For more information about validating your model, see “Overview of Model Validation and Plots” on page 10-3.

You can use `pem` to refine parameter estimates of an existing polynomial model, as described in “Refining Linear Parametric Models” on page 5-78.

For detailed information about these commands, see the corresponding reference page.

Using pem to Estimate Polynomial Models

You can estimate any single-output polynomial model using the iterative prediction-error estimation method `pem`. For Gaussian disturbances, this method gives the maximum likelihood estimate, that minimizes the prediction errors to obtain maximum-likelihood values. The resulting models are stored as `idpoly` model objects.

Use the following general syntax to both configure and estimate polynomial models:

```
m = pem(data, 'na', na,  
          'nb', nb,  
          'nc', nc,  
          'nd', nd,  
          'nf', nf,  
          'nk', nk,  
          'Property1', Value1, ...,  
          'PropertyN', ValueN)
```

where `data` is the estimation data. `na`, `nb`, `nc`, `nd`, `nf` are integers that specify the model orders, and `nk` specifies the input delays for each input. If you skip any property-value pair, the corresponding parameter value is set to zero—except `nk`, which has the default value 1. For more information about model orders, see “What Are Black-Box Polynomial Models?” on page 5-21.

Tip You do not need to construct the model object using `idoly` before estimation.

If you want to estimate the coefficients of all five polynomials, A , B , C , D , and F , you must specify an integer order for each polynomial. However, if you want to specify an ARMAX model for example, which includes only the A , B , and C polynomials, you must set `nd` and `nf` to 0.

Note To get faster estimation of ARX models, use `arx` or `iv4` instead of `pem`.

In addition to the polynomial models listed in “What Are Black-Box Polynomial Models?” on page 5-21, you can use `pem` to model the ARARX structure—called the *generalized least-squares model*—by setting `nc=nf=0`. You can also model the ARARMAX structure—called the *extended matrix model*—by setting `nf=0`.

The property-value pairs specify any model properties that configure the estimation algorithm and the initial conditions. You can enter all property-value pairs in `pem` as a comma-separated list without worrying about the hierarchy of these properties in the `idpoly` model object. For more information about accessing and setting model properties, see “Model Properties” on page 3-13.

For multiple inputs, `nb`, `nf`, and `nk` are row vectors of the same lengths as the number of input channels:

```
nb = [nb1 ... nbnu];  
nf = [nf1 ... nfnu];  
nk = [nk1 ... nknu];
```

For ARMAX, Box-Jenkins, and Output-Error models—which can only be estimated using the iterative prediction-error method—use the `armax`, `bj`, and `oe` estimation commands, respectively. These commands are versions of `pem` with simplified syntax for these specific model structures, as follows:

```
m = armax(Data,[na nb nc nk])  
m = oe(Data,[nb nf nk])  
m = bj(Data,[nb nc nd nf nk])
```

Tip If your data is sampled fast, it might help to apply a lowpass filter to the data before estimating the model. For example, to model only data in the frequency range 0-10 rad/s, use the `Focus` property, as follows:

```
m = oe(Data,[nb nf nk], 'Focus', [0 10])
```

For more information about validating your model, see “Overview of Model Validation and Plots” on page 10-3.

You can use `pem` to refine parameter estimates of an existing polynomial model, as described in “Refining Linear Parametric Models” on page 5-78.

For detailed information about `pem` and `idpoly`, see the corresponding reference page.

Options for Multiple-Input and Multiple-Output ARX Orders

To estimate a multiple-input and multiple-output (MIMO) ARX model, you must first specify the model order matrices, as follows:

- **NA** — An n_y -by- n_y matrix whose i - j th entry is the order of the polynomial that relates the j th output to the i th output.
- **NB** — An n_y -by- n_u matrix whose i - j th entry is the order of the polynomial that relates the j th input to the i th output.
- **NK** — An n_y -by- n_u matrix whose i - j th entry is the delay from the j th input to the i th output.
- For n_y outputs and n_u inputs, the A coefficients are n_y -by- n_y matrices and the B coefficients are n_y -by- n_u matrices. For more information about MIMO ARX structure, see “Definition of Multiple-Output ARX Models” on page 5-26.

Note For multiple-output time-series models, only AR models are supported. AR models require only the NA matrix.

In the System Identification Tool GUI. You can enter the matrices directly in the **Orders** field.

At the command line. Define variables that store the model order matrices and specify these variables in the `mdoel-estimation` command. You can use the following syntax to estimate a model with these orders:

```
arx(data, 'na', NA, 'nb', NB, 'nk', NK)
```

Tip To simplify entering large matrices orders in the System Identification Tool GUI, define the variable `NN=[NA NB NK]` at the command line. You can specify this variable in the **Orders** field.

Option for Frequency-Weighing Focus

You can specify how the estimation algorithm weighs the fit at various frequencies. This information supports the estimation procedures “How to Estimate Polynomial Models in the GUI” on page 5-37 and “Using pem to Estimate Polynomial Models” on page 5-42.

In the System Identification Tool GUI. Set **Focus** to one of the following options:

- **Prediction** — Uses the inverse of the noise model H to weigh the relative importance of how closely to fit the data in various frequency ranges. Corresponds to minimizing one-step-ahead prediction, which typically favors the fit over a short time interval. Optimized for output prediction applications.
- **Simulation** — Uses the input spectrum to weigh the relative importance of the fit in a specific frequency range. Does not use the noise model to weigh the relative importance of how closely to fit the data in various frequency ranges. Optimized for output simulation applications.
- **Stability** — Estimates the best stable model. For more information about model stability, see “Unstable Models” on page 10-70.
- **Filter** — Specify a custom filter to open the Estimation Focus dialog box, where you can enter a filter, as described in “Simple Passband Filter” on page 2-91 or “Defining a Custom Filter” on page 2-92. This prefiltering applies only for estimating the dynamics from input to output. The disturbance model is determined from the unfiltered estimation data.

At the command line. Specify the focus as an argument in the model-estimation command using the same options as in the GUI. For example, use this command to estimate an ARX model and emphasize the frequency content related to the input spectrum only:

```
m=arx(data,[2 2 3],'Focus','Simulation')
```

This Focus setting might produce more accurate simulation results.

Options for Initial States

When you use the iterative estimation algorithm PEM to estimate ARMAX, Box-Jenkins (BJ), Output-Error (OE), you must specify how the algorithm treats initial states.

This information supports the estimation procedures “How to Estimate Polynomial Models in the GUI” on page 5-37 and “Using pem to Estimate Polynomial Models” on page 5-42.

In the System Identification Tool GUI. For ARMAX, OE, and BJ models, set **Initial state** to one of the following options:

- **Auto** — Automatically chooses Zero, Estimate, or Backcast based on the estimation data. If initial states have negligible effect on the prediction errors, the initial states are set to zero to optimize algorithm performance.
- **Zero** — Sets all initial states to zero.
- **Estimate** — Treats the initial states as an unknown vector of parameters and estimates these states from the data.
- **Backcast** — Estimates initial states using a smoothing filter.

At the command line. Specify the initial states as an argument in the model-estimation command. For example, use this command to estimate an ARMAX model and set the initial states to zero:

```
m=armax(data,[2 2 2 3],'InitialState','zero')
```

For a complete list of values for the InitialState model property, see the `idpoly` reference page.

Algorithms for Estimating Polynomial Models

For linear ARX and AR models, you can choose between the ARX and IV algorithms. *ARX* implements the least-squares estimation method that uses QR-factorization for overdetermined linear equations. *IV* is the *instrumental variable method*. For more information about IV, see the section

on variance-optimal instruments in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

The ARX and IV algorithms treat noise differently. ARX assumes white noise. However, the instrumental variable algorithm, IV, is not sensitive to noise color. Thus, use IV when the noise in your system is not completely white and it is incorrect to assume white noise. If the models you obtained using ARX are inaccurate, try using IV.

Note AR models apply to time-series data, which has no input. For more information, see Chapter 6, “Estimating Time-Series Models”. For more information about working with AR and ARX models, see “Estimating Polynomial Models” on page 5-21.

Estimating State-Space Models

In this section...
“What Are State-Space Models?” on page 5-48
“Data Supported by State-Space Models” on page 5-52
“Supported State-Space Parameterizations” on page 5-53
“Preliminary Step – Estimating State-Space Model Orders” on page 5-54
“How to Estimate State-Space Models in the GUI” on page 5-59
“How to Estimate State-Space Models at the Command Line” on page 5-62
“How to Estimate Free-Parameterization State-Space Models” on page 5-65
“How to Estimate State-Space Models with Canonical Parameterization” on page 5-66
“How to Estimate State-Space Models with Structured Parameterization” on page 5-68
“How to Estimate the State-Space Equivalent of ARMAX and OE Models” on page 5-75
“Options for Frequency-Weighing Focus” on page 5-75
“Options for Initial States” on page 5-76
“Algorithms for Estimating State-Space Models” on page 5-76

What Are State-Space Models?

- “Definition of State-Space Models” on page 5-49
- “Continuous-Time Representation” on page 5-49
- “Discrete-Time Representation” on page 5-50
- “Relationship Between Continuous-Time and Discrete-Time State Matrices” on page 5-50
- “State-Space Representation of Transfer Functions” on page 5-51

Definition of State-Space Models

State-space models are models that use state variables to describe a system by a set of first-order differential or difference equations, rather than by one or more n th-order differential or difference equations. State variables $x(t)$ can be reconstructed from the measured input-output data, but are not themselves measured during an experiment.

The state-space model structure is a good choice for quick estimation because it requires only two parameters:

- n — Model order or the number of poles (size of the A matrix).
- nk — One or more input delays.

The *model order* for state-space models is an integer equal to the dimension of $x(t)$ and relates to the number of delayed inputs and outputs used in the corresponding linear difference equation.

Continuous-Time Representation

In continuous-time, the state-space description has the following form:

$$\begin{aligned}\dot{x}(t) &= Fx(t) + Gu(t) + \tilde{K}w(t) \\ y(t) &= Hx(t) + Du(t) + w(t) \\ x(0) &= x_0\end{aligned}$$

It is often easier to define a parameterized state-space model in continuous time because physical laws are most often described in terms of differential equations. In this case, the matrices F , G , H , and D contain elements with physical significance—for example, material constants. x_0 specifies the initial states.

Note $K=0$ gives the state-space representation of an Output-Error model. For more information about Output-Error models, see “What Are Black-Box Polynomial Models?” on page 5-21.

Discrete-Time Representation

Discrete-time state-space models provide the same type of linear difference relationship between the inputs and the outputs as the linear ARX model, but are rearranged such that there is only one delay in the expressions. The discrete-time state-space model structure is often written in the *innovations form* that describes noise:

$$\begin{aligned}x(kT + T) &= Ax(kT) + Bu(kT) + Ke(kT) \\y(kT) &= Cx(kT) + Du(kT) + e(kT) \\x(0) &= x_0\end{aligned}$$

where T is the sampling interval, $u(kT)$ is the input at time instant kT , and $y(kT)$ is the output at time instant kT .

Note $K=0$ gives the state-space representation of an Output-Error model. For more information about Output-Error models, see “What Are Black-Box Polynomial Models?” on page 5-21.

Relationship Between Continuous-Time and Discrete-Time State Matrices

The relationships between the discrete state-space matrices A , B , C , D , and K and the continuous-time state-space matrices F , G , H , D , and \tilde{K} are given for piece-wise-constant input, as follows:

$$\begin{aligned}A &= e^{FT} \\B &= \int_0^T e^{F\tau} G d\tau \\C &= H\end{aligned}$$

These relationships assume that the input is piece-wise-constant over time intervals $kT \leq t < (k+1)T$.

The exact relationship between K and \tilde{K} is complicated. However, for short sampling intervals T , the following approximation works well:

$$K = \int_0^T e^{F\tau} \tilde{K} d\tau$$

State-Space Representation of Transfer Functions

For linear models, the general symbolic model description is given by:

$$y = Gu + He$$

G is a transfer command that takes the input u to the output y . H is a transfer command that describes the properties of the additive output noise model.

The discrete-time state-space representation is given by the following equation:

$$\begin{aligned} x(kT + T) &= Ax(kT) + Bu(kT) + Ke(kT) \\ y(kT) &= Cx(kT) + Du(kT) + e(kT) \\ x(0) &= x_0 \end{aligned}$$

where T is the sampling interval, $u(kT)$ is the input at time instant kT , and $y(kT)$ is the output at time instant kT .

The relationships between the transfer functions and the discrete-time state-space matrices are given by the following equations:

$$\begin{aligned} G(q) &= C(qI_{nx} - A)^{-1}B + D \\ H(q) &= C(qI_{nx} - A)^{-1}K + I_{ny} \end{aligned}$$

where I_{nx} is the nx -by- nx identity matrix, I_{ny} is the ny -by- ny identity matrix, and ny is the dimension of y and e .

Data Supported by State-Space Models

- “Types of Supported Data” on page 5-52
- “Estimating Continuous-Time Models” on page 5-52
- “Designating Data for Estimating Discrete-Time Models” on page 5-53

Types of Supported Data

You can estimate linear state-space models from data with the following characteristics:

- Real data or complex data in any domain
- Single-output and multiple-output
- Time- or frequency-domain data

To estimate state-space models for time-series data, see Chapter 6, “Estimating Time-Series Models”.

You must first import your data into the MATLAB Workspace browser, as described in Chapter 2, “Preparing Data”.

Estimating Continuous-Time Models

Use either of the following ways to estimate continuous-time, state-space models:

- To get a linear, continuous-time model of arbitrary structure for time-domain data, you can estimate a discrete-time model, and then use `d2c` to transform it to a continuous-time model.
- Use continuous-time frequency-domain data.

To denote continuous-time frequency-domain data, set the data sampling interval to 0. You can set the sampling interval when you import data into the GUI or set the `Ts` property of the data object at the command line.

Tip Continuous state-space models are available for canonical and structured parameterizations and grey-box models. In this case, no disturbance model can be estimated.

Designating Data for Estimating Discrete-Time Models

You can estimate arbitrary-order, linear state-space models for both time- or frequency-domain data.

You must specify your data to have the sampling interval equal to the experimental data sampling interval.

You can set the sampling interval when you import data into the GUI or set the `Ts` property of the data object at the command line.

Supported State-Space Parameterizations

System Identification Toolbox supports the following parameterizations that indicate which parameters are estimated and which remain fixed at specific values:

- Free parameterization results in the estimation of all system matrix elements A , B , C , D , and K .
- Canonical forms of A , B , C , D , and K matrices.

Canonical parameterization represents a state-space system in its minimal form, using the minimum number of free parameters to capture the dynamics. Thus, free parameters appear in only a few of the rows and columns in system matrices A , B , C , and D , and the remaining matrix elements are fixed to zeros and ones.

- Structured parameterization lets you specify the values of specific parameters and exclude these parameters from estimation.
- Completely arbitrary mapping of parameters to state-space matrices. For more information, see “Estimating Linear Grey-Box Models” on page 9-5.

You can only estimate free state-space models in discrete time. Continuous state-space models are available for canonical and structured parameterizations and grey-box models.

Note To estimate canonical and structured state-space models in the System Identification Tool GUI, define the corresponding model structures at the command line and import them into the System Identification Tool GUI.

Preliminary Step – Estimating State-Space Model Orders

- “Why Estimate Model Orders?” on page 5-54
- “Estimating Model Order in the GUI” on page 5-54
- “Estimating the Model Order at the Command Line” on page 5-57
- “Using the Model Order Selection window” on page 5-58

Why Estimate Model Orders?

To estimate a state-space model using System Identification Toolbox, you must provide a model order and one or more input delays.

To get an initial model order for your system, you can estimate a group of state-space models with a range of orders for a specific delay and compare the performance of these models. You choose the model order that include states with the highest contribution to the input-output behavior of the model and use this order as an initial guess for further modeling.

The model order is always a single integer—regardless of the number of inputs and outputs. However, the number of input delays must correspond to the number of input channels.

Estimating Model Order in the GUI

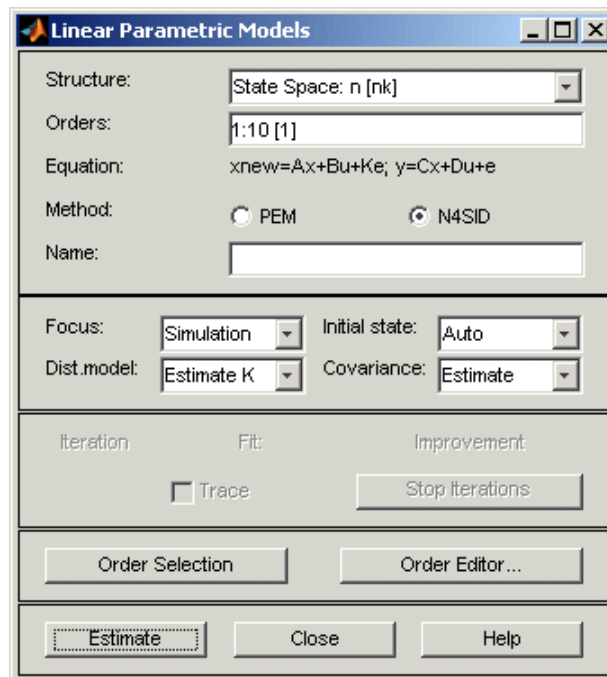
You must have already imported your data into the GUI, as described in “Importing Data into the GUI” on page 2-14.

To estimate model orders for a specific input delay:

- 1** In the System Identification Tool GUI, select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box.
- 2** In the **Structure** list, select State Space: $n [nk]$.
- 3** Edit the **Orders** field to specify a range of orders for a specific delay. For example, enter the following values for n and nk :

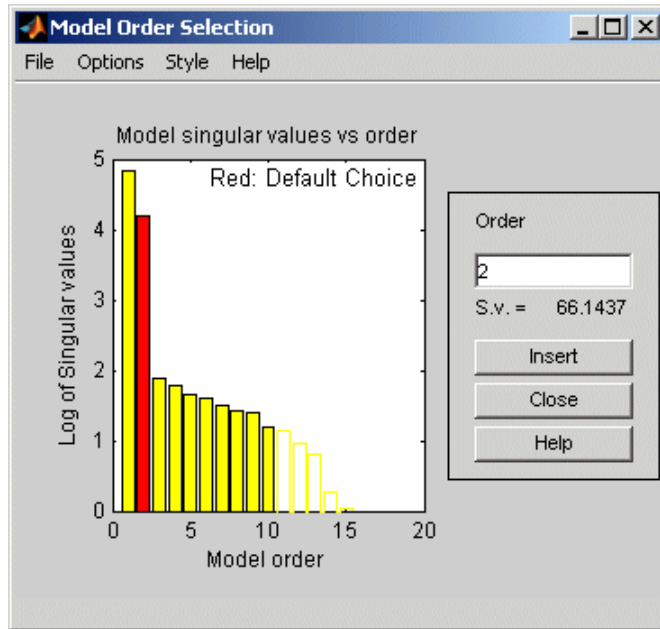
1:10 [1]

Tip As a shortcut for entering 1:10 [1], click **Order Selection**.



- 4** Verify that the **Method** is set to **N4SID**.

- 5 Click **Estimate** to open the Model Order Selection window, which displays the relative measure of how much each state contributes to the input-output behavior of the model (*log of singular values of the covariance matrix*). The following figure shows an example plot.



- 6 Select the rectangle that represents the cutoff for the states on the left that provide a significant contribution to the input-output behavior, and click **Insert** to estimate a model with this order. Red indicates the recommended choice. States 1 and 2 provide the most significant contribution. The contributions to the right of state 2 drop significantly. For information about using the Model Order Selection window, see “Using the Model Order Selection window” on page 5-58.

This action adds a new model to the Model Board in the System Identification Tool GUI. The default name of the parametric model combines the string n4s and the selected model order.

- 7 Click **Close** to close the Model Order Selection window.

After estimating model orders, use this value as an initial guess for estimating other state-space models, as described in “How to Estimate State-Space Models in the GUI” on page 5-59.

Estimating the Model Order at the Command Line

You can estimate the state-space model order using the `n4sid` command.

Use following syntax to specify the range of model orders to try for a specific input delay.

```
m = n4sid(data,n1:n2,'nk',nk);
```

where `data` is the estimation data set, `n1` and `n2` specify the range of orders, and `nk` specifies the input delay. For multiple-input systems, `nk` is a vector of input delays.

This command opens the Model Order Selection window. For information about using this plot, see “Using the Model Order Selection window” on page 5-58.

Alternatively, you can use the `pem` command to open the Model Order Selection window, as follows:

```
m = pem(Data,'nx',nn)
```

where `nn = [n1,n2,...,nN]` specifies the vector or range of orders you want to try.

To omit opening the Model Order Selection window and automatically select the best order, use the following syntax:

```
m = pem(Data,'best')
```

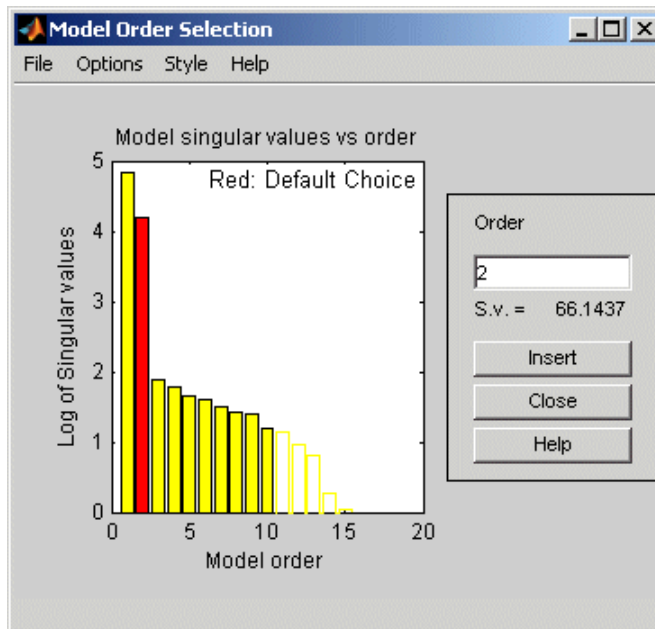
For a tutorial on estimating model orders for a multiple-input system, see “Estimating a State-Space Model” in *Getting Started with System Identification Toolbox*.

Using the Model Order Selection window

You can generate the Model Order Selection window for your data to select the number of states that provide the highest relative contribution to the input-output behavior of the model (*log of singular values of the covariance matrix*).

For a procedure on generating this plot in the System Identification Tool GUI, see “Estimating Model Order in the GUI” on page 5-54. To open this plot at the command line, see “Estimating the Model Order at the Command Line” on page 5-57.

The following figure shows a sample Model Order Selection window.



The horizontal axis corresponds to the model order n . The vertical axis, called **Log of Singular values**, shows the singular values of a covariance matrix constructed from the observed data.

You use this plot to decide which states provide a significant relative contribution to the input-output behavior, and which states provide the

smallest contribution. Based on this plot, select the rectangle that represents the cutoff for the states on the left that provide a significant contribution to the input-output behavior. The recommended choice is red.

For example, in the previous figure, states 1 and 2 provide the most significant contribution. However, the contributions of the states to the right of state 2 drop significantly. This sharp decrease in the log of the singular values after $n=2$ indicates that using two states is sufficient to get an accurate model.

How to Estimate State-Space Models in the GUI

- “Supported State-Space Models in the GUI” on page 5-59
- “Before You Begin” on page 5-59
- “Estimating State-Space Models in the GUI” on page 5-59

Supported State-Space Models in the GUI

Only free parameterization is directly supported in the System Identification Tool GUI. You can also estimate canonical and structured parameterizations at the command line and import them into the System Identification Tool GUI for parameter estimation. For more information about state-space parameterization, see “Supported State-Space Parameterizations” on page 5-53.

Before You Begin

Before you estimate state-space models, you must have already imported your data into the GUI and performed any necessary preprocessing operations. For more information, see Chapter 2, “Preparing Data”.

The following procedure also requires that you specify model order and input delays. For more information about how to estimate model orders, see “Estimating Model Order in the GUI” on page 5-54.

Estimating State-Space Models in the GUI

To estimate a state-space model with free parameterization in the System Identification Tool GUI:

1 In the System Identification Tool GUI, select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box.

2 In the **Structure** list, select State Space: $n \ [nk]$.

This action updates the options in the Linear Parametric Models dialog box to correspond with this model structure. For information about each model structure, see “What Are State-Space Models?” on page 5-48.

3 In the **Orders** field, specify the model order and delay, as follows:

- **For single-input models.** Enter the model order integer and the input delay in terms of the number of samples. Omitting nk uses the default value $nk=1$.

For example, enter 4 [2] for a fourth-order model and $nk=2$.

- **For multiple-input models.** Enter the model order integer and the input delay vector—which is a 1-by- nu vector whose i th entry is the delay for the i th input.

For example, for a two-input system, enter 4 [1 1] for a fourth-order model and a delay of 1 for each input.

- **For multiple-output models.** Enter the model order integer the same way as for single-input models.

Tip To enter model order and any delays using the Order Editor dialog box, click **Order Editor**.

4 Select the estimation **Method** as **N4SID** or **PEM**. For more information about these methods, “Algorithms for Estimating State-Space Models” on page 5-76.

5 In the **Name** field, edit the name of the model or keep the default. The name of the model should be unique in the Model Board.

6 In the **Focus** list, select how to weigh the relative importance of the fit at different frequencies. For more information about each option, see “Options for Frequency-Weighing Focus” on page 5-75.

- 7 (PEM only) In the **Initial state** list, specify how you want the algorithm to treat initial states. For more information about the available options, see “Options for Initial States” on page 5-76.

Tip If you get an inaccurate fit, try setting a specific method for handling initial states rather than choosing it automatically.

- 8 In the **Covariance** list, select **Estimate** if you want the algorithm to compute parameter uncertainties. Effects of such uncertainties are displayed on plots as model confidence regions.

To omit estimating uncertainty, select **None**. Skipping uncertainty computation reduces computation time for complex models and large data sets.

- 9 (PEM only) To view the estimation progress in the MATLAB Command Window, select the **Trace** check box. During estimation, the following information is displayed for each iteration:
 - Loss command — Equals the determinant of the estimated covariance matrix of the input noise.
 - Parameter values — Values of the model structure coefficients you specified.
 - Search direction — Change in parameter values from the previous iteration.
 - Fit improvements — Shows the actual versus expected improvements in the fit.
- 10 Click **Estimate** to add this model to the System Identification Tool GUI.
- 11 (PEM only) To stop the search and save the results after the current iteration has been completed, click **Stop Iterations**. To continue iterations from the current model, click the **Continue iter** button to assign current parameter values as initial guesses for the next search.
- 12 To plot the model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. For information about

validating your model, see “Overview of Model Validation and Plots” on page 10-3.

Tip You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

How to Estimate State-Space Models at the Command Line

- “Supported State-Space Models” on page 5-62
- “Estimating State-Space Models Using pem and n4sid” on page 5-62
- “Common Properties to Specify Model Estimation” on page 5-63
- “Choosing to Estimate D, K, and X0 Matrices” on page 5-64

Supported State-Space Models

You can only estimate discrete-time state-space models with free parameterization. Continuous state-space models are available for canonical and structured parameterizations.

Estimating State-Space Models Using pem and n4sid

You can estimate continuous-time and discrete-time polynomial model using the iterative estimation command pem that minimizes the prediction errors to obtain maximum-likelihood values. You can also use the noniterative subspace command n4sid.

You must have already estimated the model order, as described in “Preliminary Step – Estimating State-Space Model Orders” on page 5-54. You use this model order as input to the estimation functions.

Use the following general syntax to both configure and estimate state-space models:

```
m = pem(data,n,  
         'nk',nk,
```

```
'Property1',Value1,...,
'PropertyN',ValueN)
```

where `data` is the estimation data, `n` is the model order, and `nk` specifies the input delays for each input.

As an alternative to `pem`, you can use `n4sid`:

```
m = n4sid(data,n,
           'nk',nk,
           'Property1',Value1,...,
           'PropertyN',ValueN)
```

Note `pem` uses `n4sid` to initialize the state-space matrices.

For more information about the most common property-value pairs you can specify, see “Common Properties to Specify Model Estimation” on page 5-63.

For detailed information about the syntax, see the corresponding references page.

For more information about estimating model order, see “Estimating the Model Order at the Command Line” on page 5-57.

For information about validating your model, see “Overview of Model Validation and Plots” on page 10-3

Common Properties to Specify Model Estimation

The following properties are common to specify in the estimation syntax:

- `SSparameterization` — Specifies the state-space parameterization form. For more information about estimating a specific state-space parameterization, see the following topics:
 - “How to Estimate Free-Parameterization State-Space Models” on page 5-65
 - “How to Estimate State-Space Models with Canonical Parameterization” on page 5-66

- “How to Estimate State-Space Models with Structured Parameterization” on page 5-68
- `Focus` — Specifies the frequency-weighting of the noise model during estimation. See “Options for Frequency-Weighing Focus” on page 5-75.
- `DisturbanceModel` — Specifies to estimate or omit the noise model for time-domain data. See “K Matrix” on page 5-64.
- `InitialStates` — Specifies to set or estimate the initial states. See “Options for Initial States” on page 5-76

For more information about these properties, see the `idss` references page.

Choosing to Estimate D , K , and $X0$ Matrices

For state-space models with any parameterization, you can specify whether to estimate the K and $X0$ matrices, which represent the noise model and the initial states, respectively.

For state-space models with structured parameterization, you can also specify to estimate the D matrix. However, for free and canonical forms, the structure of the D matrix is set based on your choice of `nk`.

For more information about state-space structure, see “What Are State-Space Models?” on page 5-48.

D Matrix. By default, the D matrix is not estimated. Set the model property `nk` to estimate the D matrix, as follows:

- To estimate the k th column of D (corresponding to the k th input), set `nk` to 0. For `nu` inputs, `nk` is a 1-by-`nu` vector.
- To estimate the full D matrix, set all `nk` values to 0. For example, for two inputs:

```
m = pem(Data,n,'nk',[0 0])
```

To omit estimating the D matrix, set the `nk` value or values to 1, which is the default.

K Matrix. K represents the noise model.

For frequency-domain data, no noise model is estimated and K is set to 0. For time-domain data, K is estimated by default.

To modify whether K is estimated for time-domain data, you can specify the `DisturbanceModel` property in the estimator syntax.

Initially, you can omit estimating the noise parameters in K to focus on achieving a reasonable model for the system dynamics. After estimating the dynamic model, you can use `pem` to refine the model and set the K parameters to be estimated. For example:

```
m = pem(Data,md, 'DisturbanceModel', 'Estimate')
```

where `md` is the dynamic model without noise.

To set K to zero, set the value of the `DisturbanceModel` property to `'None'`. For example:

```
m = pem(Data,n, 'DisturbanceModel', 'None')
```

XO Matrices. XO stores the estimated or specified initial states of the model.

To specify how to handle the initial states, set the value of the `InitialStates` model property. For example, to set the initial states to zero, set the `InitialStates` property to `'zero'`, as follows:

```
m = pem(Data,n, 'InitialStates', 'zero')
```

When you estimate models using multiexperiment data and `InitialStates` is set to `'Estimate'`, XO stores the estimated initial states corresponding to the last experiment in the data set.

For a complete list of values for the `InitialStates` property, see “Options for Initial States” on page 5-76.

How to Estimate Free-Parameterization State-Space Models

The default parameterization of the state-space matrices A , B , C , D , and K is free; that is, any elements in the matrices are adjustable by the estimation

routines. Because the parameterization of A , B , and C is free, a basis for the state-space realization is automatically selected to give well-conditioned calculations.

You can only estimate discrete-time state-space models with any parameterization. Continuous state-space models are available for canonical and structured parameterizations only.

To estimate the disturbance model K , you must use time domain data.

Suppose that you have no knowledge about the internal structure of the discrete-time state-space model. To quickly get started, use the following syntax:

```
m = pem(data)
```

where `data` is your estimation data. This command estimates a state-space model for an automatically selected order between 1 and 10.

To find a black-box model of a specific order n , use the following syntax:

```
m = pem(Data,n)
```

The iterative algorithm `pem` is initialized by the subspace method `n4sid`. You can use `n4sid` directly, as an alternative to `pem`:

```
m = n4sid(Data,n)
```

How to Estimate State-Space Models with Canonical Parameterization

- “What Is Canonical Parameterization?” on page 5-66
- “Estimating Canonical State-Space Models” on page 5-67

What Is Canonical Parameterization?

Canonical parameterization represents a state-space system in its minimal form, using the minimum number of free parameters to capture the dynamics. Thus, free parameters appear in only a few of the rows and columns in system

matrices A , B , C , and D , and the remaining matrix elements are fixed to zeros and ones.

Of the two popular canonical forms, which include *controllable canonical form* and *observable canonical form*, System Identification Toolbox supports only controllable forms. Controllable canonical structures include free parameters in output rows of the A matrix, free B and K matrices, and the fixed C matrix. The representation within controllable canonical forms is not unique and the exact form depends on the actual choices of canonical indices. For more information about the distribution of free parameters in canonical forms, see the appendix on identifiability of black-box multivariable model structures in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999 (equation 4A.16).

Estimating Canonical State-Space Models

You can estimate state-space models with canonical parameterization at the command line.

To specify a canonical form for A , B , C , and D , set the `SSparameterization` model property directly in the estimator syntax, as follows:

```
m = pem(data,n,'SSparameterization','canonical')
```

If you have time-domain data, the preceding command estimates a discrete-time model.

Note When you estimate the D matrix in canonical form, you must set the `nk` property. See “Choosing to Estimate D , K , and X_0 Matrices” on page 5-64.

If you have continuous-time frequency-domain data, the preceding syntax estimates an n th order continuous-time state-space model with no direct contribution from the input to the output ($D=0$). To include a D matrix, set the `nk` property to 0 in the estimation, as follows:

```
m = pem(data,n,'SSparameterization','canonical',
         'nk',0)
```

You can specify additional property-value pairs similar to the free-parameterization case, as described in “How to Estimate Free-Parameterization State-Space Models” on page 5-65.

For information about validating your model, see “Overview of Model Validation and Plots” on page 10-3

How to Estimate State-Space Models with Structured Parameterization

- “What Is Structured Parameterization?” on page 5-68
- “Specifying the State-Space Structure” on page 5-69
- “Are Grey-Box Models Similar to State-Space Models with Structured Parameterization?” on page 5-71
- “Example – Estimating Structured Discrete-Time State-Space Models” on page 5-71
- “Example – Estimating Structured Continuous-Time State-Space Models” on page 5-72

What Is Structured Parameterization?

Structured parameterization lets you exclude specific parameters from estimation by setting these parameters to specific values. This approach is useful when you can derive state-space matrices from physical principles and provide initial parameter values based on physical insight. You can use this approach to discover what happens if you fix specific parameter values or if you free certain parameters.

In the case of structured parameterization, there are two stages to the estimation procedure:

- 1** Using the `idss` command to specify the structure of the state-space matrices and the initial values of the free parameters
- 2** Using the `pem` estimation command to estimate the free model parameters

This approach differs from estimating models with free and canonical parameterizations, where it is not necessary to specify initial parameter values before the estimation. For free parameterization, there is no structure to specify because it is assumed to be unknown. For canonical parameterization, the structure is fixed to a specific form.

For information about validating your model, see “Overview of Model Validation and Plots” on page 10-3

Specifying the State-Space Structure

To specify the state-space model structure, first define the A, B, C, D, K and X0 matrices in the MATLAB Workspace browser.

To define a discrete-time state-space structure, use the following syntax:

```
m = idss(A,B,C,D,K,X0,
         'Ts',T,
         'SSparameterization','structured')
```

where A, B, C, D, and K specify both the fixed parameter values and the initial values for the free parameters. T is the sampling interval. Setting SSparameterization to 'structured' flags that you want to estimate a partial structure for this state-space model.

Similarly, to define a continuous-time state-space structure, use the following syntax:

```
m = idss(A,B,C,D,K,X0,
         'Ts',0,
         'SSparameterization','structured')
```

In the continuous-time case, you must set the sampling interval property Ts to zero.

After you create the nominal model structure, you must specify which parameters to estimate and which to set to specific values. To accomplish this, you must edit the structures of the following model properties: As, Bs, Cs, Ds, Ks, and x0s. These *structure matrices* are properties of the nominal model you constructed and have the same sizes as A, B, C, D, K, and x0, respectively. Initially, the structure matrices contain NaN values.

Specify the structure matrix values, as follows:

- Set a NaN value to flag free parameters at the corresponding locations in A, B, C, D, K, and x0.
- Specify the values of fixed parameters at the corresponding locations in A, B, C, D, K, and x0.

For example, suppose that you constructed a nominal state-space model `m` with the following A matrix:

$$A = [2 \ 0; \ 0 \ 3]$$

Suppose you want to fix $A(1,2)=A(2,1)=0$. To specify the parameters you want to fix, enter their values at the corresponding locations in the structure matrix `As`:

$$m.As = [NaN \ 0; \ 0 \ NaN]$$

The estimation algorithm only estimates the parameters in A that have a NaN value in `As`.

Finally, use `pem` to estimate the model, as described in “How to Estimate State-Space Models at the Command Line” on page 5-62.

Use physical insight, whenever possible, to initialize the parameters for the iterative search algorithm. Because it is possible that the numerical minimization gets stuck in a local minimum, try several different initialization values for the parameters. For random initialization, use the `init` command. When the model structure contains parameters with different orders of magnitude, try to scale the variables so that the parameters are all roughly the same magnitude.

The iterative search computes gradients of the prediction errors with respect to the parameters using numerical differentiation. The step size is specified by the `nuderst` M-file. The default step size is equal to 10^{-4} times the absolute value of a parameter or equal to 10^{-7} , whichever is larger. To specify a different step size, edit the `nuderst` M-file.

Are Grey-Box Models Similar to State-Space Models with Structured Parameterization?

Structured parameterization state-space models are similar to grey-box modeling. However, the state-space models are simpler to estimate than grey-box models. To learn more about grey-box models, see Chapter 9, “Estimating Grey-Box Models”.

Example – Estimating Structured Discrete-Time State-Space Models

In this example, you estimate the unknown parameters $(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5)$ in the following discrete-time model:

$$\begin{aligned}x(t+1) &= \begin{bmatrix} 1 & \theta_1 \\ 0 & 1 \end{bmatrix} x(t) + \begin{bmatrix} \theta_2 \\ \theta_3 \end{bmatrix} u(t) + \begin{bmatrix} \theta_4 \\ \theta_5 \end{bmatrix} e(t) \\ y(t) &= [1 \quad 0] x(t) + e(t) \\ x(0) &= \begin{bmatrix} 0 \\ 0 \end{bmatrix}\end{aligned}$$

Suppose that the nominal values of the unknown parameters $(\theta_1, \theta_2, \theta_3, \theta_4, \theta_5)$ are -1, 2, 3, 4, and 5, respectively.

The discrete-time state-space model structure is defined by the following equation:

$$\begin{aligned}x(kT+T) &= Ax(kT) + Bu(kT) + Ke(kT) \\ y(kT) &= Cx(kT) + Du(kT) + e(kT) \\ x(0) &= x_0\end{aligned}$$

To construct and estimate the parameters of this discrete-time state-space model:

- 1 Construct the parameter matrices and initialize the parameter values using the nominal parameter values:

$$\begin{aligned}A &= [1, -1; 0, 1]; \\ B &= [2; 3];\end{aligned}$$

```
C = [1,0];  
D = 0;  
K = [4;5];
```

2 Construct the state-space model object:

```
m = idss(A,B,C,D,K);
```

3 Specify the parameter values in the structure matrices that you do not want to estimate:

```
m.As = [1, NaN; 0 ,1];  
m.Bs = [NaN;NaN];  
m.Cs = [1, 0];  
m.Ds = 0;  
m.Ks = [NaN;NaN];  
m.x0s = [0;0];
```

4 Estimate the model structure:

```
m = pem(data,m)
```

where data is name of the iddata object containing time-domain or frequency-domain data. The iterative search starts with the nominal values in the A, B, C, D, K, and x0 matrices.

Example – Estimating Structured Continuous-Time State-Space Models

In this example, you estimate the unknown parameters $(\theta_1, \theta_2, \theta_3)$ in the following continuous-time model:

$$\dot{x} = \begin{bmatrix} 0 & 1 \\ 0 & \theta_1 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ \theta_2 \end{bmatrix} u(t)$$
$$y(t) = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} x(t) + e(t)$$
$$x(0) = \begin{bmatrix} \theta_3 \\ 0 \end{bmatrix}$$

This equation represents an electrical motor, where $y_1(t) = x_1(t)$ is the angular position of the motor shaft, and $y_2(t) = x_2(t)$ is the angular velocity.

The parameter $-\theta_1$ is the inverse time constant of the motor, and $-\theta_2/\theta_1$ is the static gain from the input to the angular velocity.

The motor is at rest at $t=0$, but its angular position θ_3 is unknown. Suppose that the approximate nominal values of the unknown parameters are $\theta_1 = -1$ and $\theta_2 = 0.25$. The variance of the errors in the position measurement is 0.01, and the variance in the angular velocity measurements is 0.1. For more information about this example, see the section on state-space models in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

The continuous-time state-space model structure is defined by the following equation:

$$\begin{aligned}\dot{x}(t) &= Fx(t) + Gu(t) + \tilde{K}w(t) \\ y(t) &= Hx(t) + Du(t) + w(t) \\ x(0) &= x_0\end{aligned}$$

To construct and estimate the parameters of this continuous-time state-space model:

- 1 Construct the parameter matrices and initialize the parameter values using the nominal parameter values:

Note The following matrices correspond to continuous-time representation. However, to be consistent with the `idss` object property name, this example uses A, B, and C instead of F, G, and H.

$$\begin{aligned}A &= [0 \ 1; 0 \ -1]; \\ B &= [0; 0.25]; \\ C &= \text{eye}(2); \\ D &= [0; 0];\end{aligned}$$

```
K = zeros(2,2);  
x0 = [0;0];
```

2 Construct the continuous-time state-space model object:

```
m = idss(A,B,C,D,K,x0,'Ts',0);
```

3 Specify the parameter values in the structure matrices that you do not want to estimate:

```
m.As = [0 1;0 NaN];  
m.Bs = [0;NaN];  
m.Cs = m.c;  
m.Ds = m.d;  
m.Ks = m.k;  
m.x0s = [NaN;0]  
m.NoiseVariance = [0.01 0; 0 0.1];
```

4 Estimate the model structure:

```
m = pem(data,m)
```

where `data` is name of the `iddata` object containing time-domain or frequency-domain data. The iterative search for a minimum is initialized by the parameters in the nominal model `m`. The continuous-time model is sampled using the same sampling interval as the data.

5 To simulate this system using the sampling interval $T = 0.1$ for input `u` and the noise realization `e`, use the following commands:

```
e = randn(300,2);  
u = idinput(300);  
simdat = iddata([], [u e], 'Ts', 0.1);  
y = sim(m, simdat)
```

The continuous system is automatically sampled using $T_s=0.1$. The noise sequence is scaled according to the matrix `m.noisevar`.

If you discover that the motor was not initially at rest, you can estimate $x_2(0)$ by setting the second element of the `x0s` structure matrix to `NaN`, as follows:

```
m_new = pem(data,m,'x0s',[NaN;NaN])
```

How to Estimate the State-Space Equivalent of ARMAX and OE Models

You can estimate the equivalent of ARMAX and output-error (OE) multiple-output models using state-space model structures. For the ARMAX case, specify to estimate the K matrix for the state-space model. For the OE case, set K to zero.

For more information about ARMAX and OE models, see “Estimating Polynomial Models” on page 5-21.

Options for Frequency-Weighing Focus

You can specify how the estimation algorithm weighs the fit at various frequencies. This information supports the estimation procedures “How to Estimate State-Space Models in the GUI” on page 5-59 and “How to Estimate State-Space Models at the Command Line” on page 5-62.

In the System Identification Tool GUI. Set **Focus** to one of the following options:

- **Prediction** — Uses the inverse of the noise model H to weigh the relative importance of how closely to fit the data in various frequency ranges. Corresponds to minimizing one-step-ahead prediction, which typically favors the fit over a short time interval. Optimized for output prediction applications.
- **Simulation** — Uses the input spectrum to weigh the relative importance of the fit in a specific frequency range. Does not use the noise model to weigh the relative importance of how closely to fit the data in various frequency ranges. Optimized for output simulation applications.
- **Stability** — Estimates the best stable model. For more information about model stability, see “Unstable Models” on page 10-70.
- **Filter** — Specify a custom filter to open the Estimation Focus dialog box, where you can enter a filter, as described in “Simple Passband Filter” on page 2-91 or “Defining a Custom Filter” on page 2-92. This prefiltering applies only for estimating the dynamics from input to output. The disturbance model is determined from the estimation data.

At the command line. Specify the focus as an argument in the model-estimation command using the same options as in the GUI. For example, use this command to emphasize the fit between the 5 and 8 rad/s:

```
pem(data,4,'Focus',[5 8])
```

Options for Initial States

If you estimate state-space models using the iterative estimation algorithm `pem`, you must specify how the algorithm treats initial states. This information supports the estimation procedures “How to Estimate State-Space Models in the GUI” on page 5-59 and “How to Estimate State-Space Models at the Command Line” on page 5-62.

In the System Identification Tool GUI. Set **Initial state** to one of the following options:

- **Auto** — Automatically chooses **Zero**, **Estimate**, or **Backcast** based on the estimation data. If initial states have negligible effect on the prediction errors, the initial states are set to zero to optimize algorithm performance.
- **Zero** — Sets all initial states to zero.
- **Estimate** — Treats the initial states as an unknown vector of parameters and estimates these states from the data.
- **Backcast** — Estimates initial states using a backward filtering method (least-squares fit).

At the command line. Specify the initial states as an argument in the estimation command `pem`. For example, use this command to estimate a fourth-order state-space model and set the initial states to be estimated from the data:

```
m=pem(data,4,'InitialState','estimate')
```

For a complete list of values for the `InitialState` model property, see the `idss` reference page.

Algorithms for Estimating State-Space Models

For linear state-space models, you can use the subspace method, called *N4SID*. You can use the subspace method `N4SID` to get an initial model (see

the `n4sid` reference page), and then try to refine the initial estimate using the iterative prediction-error method PEM (see the `pem` reference page).

N4SID is faster than PEM, but is typically less accurate and robust, and requires additional arguments that might be difficult to specify.

You can use the iterative *prediction-error minimization (PEM)* (maximum likelihood) algorithm for all linear and nonlinear model types.

Refining Linear Parametric Models

In this section...
“When to Refine Models” on page 5-78
“What You Specify to Refine a Model” on page 5-78
“How to Refine Linear Parametric Models in the GUI” on page 5-79
“How to Refine Linear Parametric Models at the Command Line” on page 5-80

When to Refine Models

There are two situations where you can refine estimates of linear parametric models.

In the first situation, you have already estimated a parametric model and wish to refine the model. However, if your model captures the essential dynamics, it is usually not necessary to continue improving the fit—especially when the improvement is a fraction of a percent.

In the second situation, you might have constructed a model using one of the model constructors described in “Command for Creating Model Objects” on page 3-4. In this case, you built initial parameter guesses into the model structure and wish to refine these parameter values.

Note Because it is difficult to specify nonlinear model parameters in advance, you typically only estimate nonlinear models.

What You Specify to Refine a Model

When you refine a model, you must provide two inputs:

- Parametric model
- Data — You can either use the same data set for refining the model as the one you originally used to estimate the model, or you can use a different data set.

How to Refine Linear Parametric Models in the GUI

The following procedure assumes that the model you want to refine is already in the System Identification Tool GUI. You might have estimated this model in the current session or imported the model from the MATLAB Workspace browser. For information about importing models into the GUI, see “Importing Models into the System Identification Tool GUI” on page 1-11.

To refine your model:

- 1 In the System Identification Tool GUI, verify that you have the correct data set in the **Working Data** area for refining your model.

If you are using a different data set than the one you used to estimate the model, drag the correct data set into the **Working Data** area. For more information about specifying estimation data, see “Specifying Estimation and Validation Data” on page 1-11.

- 2 Select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box, if this dialog box is not already open.
- 3 In the Linear Parametric Models dialog box, select **By Initial Model** from the **Structure** list.
- 4 Enter the model name into the **Initial model** field, and press **Enter**.

The model name must be in the Model Board of the System Identification Tool GUI or a variable in the MATLAB Workspace browser.

Tip As a shortcut for specifying a model in the Model Board, you can drag the model icon from the System Identification Tool GUI into the **Initial model** field.

When you enter the model name, algorithm settings in the Linear Parametric Models dialog box override the initial model settings.

- 5 Modify the algorithm settings, displayed in the Linear Parametric Models dialog box, if necessary.
- 6 Click **Estimate** to refine the model.

- 7 Validate the new model, as described in Chapter 10, “Validating and Analyzing Models”.

Tip To continue refining the model using additional iterations, click **Continue iter**. This action continues parameter estimation using the most recent model.

How to Refine Linear Parametric Models at the Command Line

If you are working at the command line, you can use `pem` to refine parametric model estimates.

The general syntax for refining initial models is as follows:

```
m = pem(data,init_model)
```

`pem` uses the properties of the initial model unless you specify different properties. For more information about specifying model properties directly in the estimator, see “Specifying Model Properties in the Estimator” on page 3-15.

Example – Refining an Initial ARMAX Model at the Command Line

The following example shows to estimate an initial model and try to refine this model using pem:

```
load iddata8

% Split the data z8 into two parts.
% Create new data object with first hundred samples
z8a = z8(1:100);

% Create new data object with remaining samples
z8b = z8(101:end);

% Estimate ARMAX model with default Algorithm
% properties, na=4, nb=[3 2 3], nc=2, and nk=[0 0 0]
m1 = armax(z8a,[4 3 2 3 2 0 0 0]);

% Refine the initial model m1 using the data set z8b,
% and stricter algorithm settings with increased number
% of maximum iterations (MaxIter) and smaller tolerance
m2 = pem(z8b,m1,'tol',1e-5,'maxiter',50);
```

For more information about estimating polynomial models, see “Estimating Polynomial Models” on page 5-21.

Example – Refining an ARMAX Model with Initial Parameter Guesses at the Command Line

The following example shows how to refine models for which you have initial parameter guesses. This example estimates an ARMAX model for the data and requires you to initialize the A , B , and C polynomials.

In this case, you must first create a model object and set the initial parameter values in the model properties. Next, you provide this initial model as input to `pem`, which refines the initial parameter guesses using the data.

```
load iddata8
% Define model parameters
A = [1 -1.2 0.7];
B(1,:) = [0 1 0.5 0.1]; % first input
B(2,:) = [0 1.5 -0.5 0]; % second input
B(3,:) = [0 -0.1 0.5 -0.1]; % third input
C = [1 0 0 0 0];
Ts = 1;
% Leading zeros in B matrix indicate input delay (nk),
% which is 1 for each input channel. The trailing zeros
% in B(2,:) make the number of coefficients equal
% for all channels.

% Create model object
init_model = idpoly(A,B,C,'Ts',1);

% Use pem to refine initial model
model = pem(z8,init_model)

% Compare the two models
compare(z8,init_model,model)
```

For more information about estimating polynomial models, see “Estimating Polynomial Models” on page 5-21.

Estimating Time-Series Models

What Are Time-Series Models?
(p. 6-2)

Definition of time-series models.

Preparing Time-Series Data (p. 6-3)

Where you can learn more about importing and preparing time-series data for modeling.

Estimating Time-Series Power Spectra (p. 6-4)

How to estimate power spectra for time-series data in the GUI and at the command line.

Estimating AR and ARMA Models
(p. 6-7)

How to estimate polynomial AR and ARMA models for time-series data in the GUI and at the command line.

Estimating State-Space Time-Series Models (p. 6-12)

How to estimate state-space models for time-series data in the GUI and at the command line.

Example – Estimating Time Series at the Command Line (p. 6-14)

How to simulate a time-series model, compare spectral estimates, covariance estimates, and predicted output of the model.

Estimating Nonlinear Models for Time-Series Data (p. 6-15)

Where to learn more about estimating nonlinear ARX models for time-series data.

What Are Time-Series Models?

A *time series* is one or more measured output channels with no measured input.

System Identification Toolbox lets you estimate time-series spectra using both time- and frequency-domain data. Time-series spectra describe time-series variations using cyclic components at different frequencies.

You can also estimate parametric autoregressive (AR), autoregressive and moving average (ARMA), and state-space time-series models. For a definition of these models, see “Definition of AR and ARMA Models” on page 6-7.

Note ARMA and state-space models are supported for time-domain data only. Only single-output ARMA models are supported.

Preparing Time-Series Data

Before you can estimate models for time-series data, you must import your data into MATLAB. System Identification Toolbox supports both time-domain and frequency-domain data. For more information about how to import data into MATLAB and prepare it for modeling, see “Ways to Prepare Data for Modeling” on page 2-3.

Note For information about which variables you need to represent time-series data, see “Representing Time-Series Data” on page 2-8.

If your data is already in the MATLAB Workspace browser, you can import it directly into the System Identification Tool GUI. If you prefer to work at the command line, you must represent the data as a System Identification Toolbox data object instead.

In the System Identification Tool GUI. When you import scalar or multiple-output time series data into the GUI, leave the **Input** field empty. For more information about importing data, see “Importing Data into the System Identification Tool GUI” on page 2-14.

At the command line. To represent a time series vector or a matrix s as an `iddata` object, use the following syntax:

```
y = iddata(s,[],Ts);
```

s contains as many columns as there are measured outputs. For time-domain data, set T_s to the sampling interval. For continuous-time frequency domain data, set T_s to 0.

Estimating Time-Series Power Spectra

In this section...

“How to Estimate Time-Series Power Spectra Using the GUI” on page 6-4

“How to Estimate Time-Series Power Spectra at the Command Line” on page 6-5

How to Estimate Time-Series Power Spectra Using the GUI

You must have already imported your data into the GUI, as described in “Preparing Time-Series Data” on page 6-3.

To estimate time-series spectral models in the System Identification Tool GUI:

- 1** In the System Identification Tool GUI, select **Estimate > Spectral models** to open the Spectral Model dialog box.
- 2** In the **Method** list, select the spectral analysis method you want to use. For information about each method, see “Options for Computing Spectral Models” on page 4-13.
- 3** Specify the frequencies at which to compute the spectral model in either of the following ways:
 - In the **Frequencies** field, enter either a vector of values, a MATLAB expression that evaluates to a vector, or a variable name of a vector in the MATLAB Workspace browser. For example, `logspace(-1,2,500)`.
 - Use the combination of **Frequency Spacing** and **Frequencies** to construct the frequency vector of values:
 - In the **Frequency Spacing** list, select **Linear** or **Logarithmic** frequency spacing.

Note For `etfe`, only the **Linear** option is available.

- In the **Frequencies** field, enter the number of frequency points.

For time-domain data, the frequency ranges from 0 to the Nyquist frequency. For frequency-domain data, the frequency ranges from the smallest to the largest frequency in the data set.

- 4** In the **Frequency Resolution** field, enter the frequency resolution, as described in “Options for Frequency Resolution” on page 4-14. To use the default value, enter `default` or leave the field empty.
- 5** In the **Model Name** field, enter the name of the correlation analysis model. The model name should be unique in the Model Board.
- 6** Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
- 7** In the Spectral Model dialog box, click **Close**.
- 8** To view the estimated disturbance spectrum, select the **Noise spectrum** check box in the System Identification Tool GUI. For more information about working with this plot, see “Creating Noise-Spectrum Plots” on page 10-37.

To export the model to the MATLAB Workspace browser, drag it to the **To Workspace** rectangle in the System Identification Tool GUI. You can view the power spectrum and the confidence intervals of the resulting `idfrd` model object using the `bode` command.

How to Estimate Time-Series Power Spectra at the Command Line

You can use the `etfe`, `spa`, and `spafdr` commands to estimate power spectra of time series for both time-domain and frequency-domain data. The following table provides a brief description of each command.

You must have already prepared your data, as described in “Preparing Time-Series Data” on page 6-3.

The resulting models are stored as an `idfrd` model object, which contains `SpectrumData` and its variance. For multiple-output data, `SpectrumData` contains power spectra of each output and the cross-spectra between each output pair.

Estimating Frequency Response of Time Series

Command	Description
etfe	Estimates a periodogram using Fourier analysis.
spa	Estimates the power spectrum with its standard deviation using spectral analysis.
spafdr	Estimates the power spectrum with its standard deviation using a variable frequency resolution.

For example, suppose y is time-series data. The following commands estimate the power spectrum g and the periodogram p , and plot both models with three standard deviation confidence intervals:

```
g = spa(y)
p = etfe(y)
bode(g,p,'sd',3)
```

For detailed information about these commands, see the corresponding reference page.

Estimating AR and ARMA Models

In this section...

“Definition of AR and ARMA Models” on page 6-7

“Estimating Polynomial Models in the GUI” on page 6-7

“Estimating AR and ARMA Models at the Command Line” on page 6-10

Definition of AR and ARMA Models

For a single-output signal $y(t)$, the AR model is given by the following equation:

$$A(q)y(t) = e(t)$$

The AR model is a special case of the ARX model with no input.

The ARMA model for a single-output time-series is given by the following equation:

$$A(q)y(t) = C(q)e(t)$$

The ARMA structure reduces to the AR structure for $C(q)=1$. The ARMA model is a special case of the ARMAX model with no input.

For more information about polynomial models, see “What Are Black-Box Polynomial Models?” on page 5-21.

Estimating Polynomial Models in the GUI

Before you begin, you must have accomplished the following:

- Prepared the data, as described in “Preparing Time-Series Data” on page 6-3.
- Estimated model order, as described in “Preliminary Step – Estimating Model Orders and Input Delays ” on page 5-29.

- (Multiple-output AR models only) Specified the model-order matrix in the MATLAB Workspace browser before estimation, as described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44.

To estimate AR and ARMA models using the System Identification Tool GUI:

- 1** In the System Identification Tool GUI, select **Estimate > Linear parametric models** to open the Linear Parametric Models dialog box.
- 2** In the **Structure** list, select the polynomial model structure you want to estimate from the following options:
 - AR: [na]
 - ARMA: [na nc]

This action updates the options in the Linear Parametric Models dialog box to correspond with this model structure. For information about each model structure, see “Definition of AR and ARMA Models” on page 6-7.

Note OE and BJ structures are not available for time-series models.

- 3** In the **Orders** field, specify the model orders, as follows:
 - **For single-output models.** Enter the model orders according to the sequence displayed in the **Structure** field.
 - **For multiple-output ARX models.** (AR models only) Enter the model orders directly, as described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44. Alternatively, enter the name of the matrix NA in the MATLAB Workspace browser that stores model orders, which is Ny-by-Ny.

Tip To enter model orders and delays using the Order Editor dialog box, click **Order Editor**.

-
- 4** (AR models only) Select the estimation **Method** as **ARX** or **IV** (instrumental variable method). For more information about these methods, see “Algorithms for Estimating Polynomial Models” on page 5-46.

Note IV is not available for multiple-output data.

- 5** In the **Name** field, edit the name of the model or keep the default. The name of the model should be unique in the Model Board.
- 6** In the **Initial state** list, specify how you want the algorithm to treat initial states. For more information about the available options, see “Options for Initial States” on page 5-19.

Tip If you get an inaccurate fit, try setting a specific method for handling initial states rather than choosing it automatically.

- 7** In the **Covariance** list, select **Estimate** if you want the algorithm to compute parameter uncertainties. Effects of such uncertainties are displayed on plots as model confidence regions.

To omit estimating uncertainty, select **None**. Skipping uncertainty computation might reduce computation time for complex models and large data sets.

- 8** (ARMA only) To view the estimation progress at the command line, select the **Trace** check box. During estimation, the following information is displayed for each iteration:
- **Loss command** — Equals the determinant of the estimated covariance matrix of the input noise.
 - **Parameter values** — Values of the model structure coefficients you specified.
 - **Search direction** — Changes in parameter values from the previous iteration.

- Fit improvements — Shows the actual versus expected improvements in the fit.
- 9 Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.
 - 10 (Prediction-error method only) To stop the search and save the results after the current iteration has been completed, click **Stop Iterations**. To continue iterations from the current model, click the **Continue iter** button to assign current parameter values as initial guesses for the next search.
 - 11 To plot the model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. For more information about validating models, see “Overview of Model Validation and Plots” on page 10-3.

You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI. For more information about working with models, see Chapter 13, “Using Models with Other Products”.

Estimating AR and ARMA Models at the Command Line

You can estimate AR and ARMA models at the command line. For single-output time-series, the resulting models are `idpoly` model objects. For multiple-output time-series, the resulting models are `idarx` model objects. For more information about models objects, see “Creating Model Objects” on page 3-2.

The following table summarizes the command and specifies whether single-output or multiple-output models are supported.

Commands for Estimating Polynomial Time-Series Models

Method Name	Description	Supported Data
ar	Noniterative, least-squares method to estimate linear, discrete-time single-output AR models.	Time-domain, time-series iddata data object.
armax	Iterative prediction-error method to estimate linear, single-output ARMAX models.	Time-domain, time-series iddata data object.
arx	Noniterative, least-squares method for estimating single-output and multiple-output linear AR models.	Supports time- and frequency-domain time-series iddata data.
ivar	Noniterative, instrumental variable method for estimating single-output AR models.	Supports time-domain, time-series iddata data.

The following code shows usage examples for estimating AR models:

```
% For scalar signals
m = ar(y,na)
% For multiple-output vector signals
m = arx(y,na)
% Instrumental variable method
m = ivar(y,na)
% For ARMA, do not need to specify nb and nk
th = armax(y,[na nc])
```

The ar command provides additional options to let you choose the algorithm for computing the least-squares from a group of several popular techniques from the following methods:

- Burg (geometric lattice)
- Yule-Walker
- Covariance

For more information about validating models, see “Overview of Model Validation and Plots” on page 10-3.

Estimating State-Space Time-Series Models

In this section...
“Definition of State-Space Time-Series Model” on page 6-12
“Estimating State-Space Models at the Command Line” on page 6-12

Definition of State-Space Time-Series Model

The discrete-time state-space model for a time series is given by the following equations:

$$x(kT + T) = Ax(kT) + Ke(kT)$$

$$y(kT) = Cx(kT) + e(kT)$$

where T is the sampling interval and $y(kT)$ is the output at time instant kT .

The time-series structure corresponds to the general structure with empty B and D matrices.

For information about general discrete-time and continuous-time structures for state-space models, see “What Are State-Space Models?” on page 5-48.

Estimating State-Space Models at the Command Line

You can estimate single-output and multiple-output state-space models at the command line for time-domain and frequency-domain data (`iddata` object).

The following table provides a brief description of each command. The resulting models are `idss` model objects.

Commands for Estimating State-Space Time-Series Models

Method Name	Description
n4sid	<p data-bbox="520 361 1262 418">Noniterative, subspace method for estimating discrete-time linear state-space models.</p> <hr data-bbox="520 475 1323 479"/> <p data-bbox="520 487 1314 545">Note When you use pem to estimate a state-space model, n4sid creates the initial model.</p> <hr data-bbox="520 557 1323 560"/>
pem	<p data-bbox="520 598 1326 656">Estimates linear, discrete-time time-series models using iterative estimation method that minimizes the prediction error.</p>

Example – Estimating Time Series at the Command Line

Here is an example of how you simulate a time-series model, compare spectral estimates, covariance estimates, and predict output of the model:

```
ts0 = idpoly([1 -1.5 0.7],[]);
ir = sim(ts0,[1;zeros(24,1)]);
% Define the true covariance command
Ry0 = conv(ir,ir(25:-1:1));
e = idinput(200,'rgs');
% Define y vector
y = sim(ts0,e);
% iddata object with sampling time 1
y = iddata(y)
plot(y)
per = etfe(y);
speh = spa(y);
ffplot(per,speh,ts0)
% Estimate a second-order AR model
ts2 = ar(y,2);
ffplot(speh,ts2,ts0,'sd',3)
% Get covariance command estimates
Ryh = covf(y,25);
Ryh = [Ryh(end:-1:2),Ryh]';
ir2 = sim(ts2,[1;zeros(24,1)]);
Ry2 = conv(ir2,ir2(25:-1:1));
plot([-24:24]*ones(1,3),[Ryh,Ry2,Ry0])
% The prediction ability of the model
compare(y,ts2,5)
```


Estimating Nonlinear Models for Time-Series Data

When a linear model provides an insufficient description of the dynamics, you can try estimating a nonlinear models. To learn more about when to estimate nonlinear models, see “When to Estimate Linear Versus Nonlinear Models”.

Before you can estimate models for time-series data, you must have already prepared the data as described in “Preparing Time-Series Data” on page 6-3.

For black-box modeling, System Identification Toolbox supports nonlinear ARX models for time-series data. To learn how to estimate this type of model, see “Estimating Nonlinear ARX Models” on page 8-4.

If you understand the underlying physics of the system, you can specify an ordinary differential or difference equation and estimate the coefficients. To learn how to estimate this type of model, see “Estimating Nonlinear Grey-Box Models” on page 9-13.

For more information about validating models, see “Overview of Model Validation and Plots” on page 10-3.

Estimating Linear Models Using Recursive Techniques

What Is Recursive Estimation?
(p. 7-2)

Definition of recursive estimation
and its applications.

Commands for Recursive Estimation
(p. 7-3)

Summary of commands and syntax
for recursive estimation.

Algorithms for Recursive Estimation
(p. 7-6)

Description of supported algorithms
for recursive estimation.

Data Segmentation (p. 7-14)

Use of data segmentation to model
systems exhibiting abrupt changes.

What Is Recursive Estimation?

Many real-world applications, such as adaptive control, adaptive filtering, and adaptive prediction, require a model of the system to be available online while the system is in operation. Estimating models for batches of input-output data is useful for addressing the following types of questions regarding system operation:

- Which input should be applied at the next sampling instant?
- How should the parameters of a matched filter be tuned?
- What are the predictions of the next few outputs?
- Has a failure occurred? If so, what type of failure?

You might also use online models to investigate time variations in system and signal properties.

The methods for computing online models are called *recursive identification methods*. Recursive algorithms are also called *recursive parameter estimation*, *adaptive parameter estimation*, *sequential estimation*, and *online algorithms*.

For examples of recursive estimation and data segmentation using System Identification Toolbox, run the Recursive Estimation and Data Segmentation demo by typing the following command at the MATLAB prompt:

```
iddemo5
```

For detailed information about recursive parameter estimation algorithms, see the corresponding chapter in *System Identification: Theory for the User* by Lennart Ljung (Prentice Hall PTR, Upper Saddle River, NJ, 1999).

Commands for Recursive Estimation

You can use System Identification Toolbox commands to estimate linear polynomial models, such as ARX, ARMAX, Box-Jenkins, and Output-Error models. If you are working with time-series data that contains no inputs and a single output, you can estimate AR (Auto-Regressive) and ARMA (Auto-Regressive Moving Average) single-output models.

Before estimating models using recursive algorithms, you must import your data into the MATLAB Workspace browser and represent your data in either of the following formats:

- Matrix of the form $[y \ u]$. y represents the output data using one or more column vectors. Similarly, u represents the input data using one or more column vectors.
- `iddata` or `idfrd` object. For more information about creating these objects, see Chapter 2, “Preparing Data”.

The general syntax for recursive estimation commands is as follows:

```
[params, y_hat]=command(data, nn, adm, adg)
```

`params` matrix contains the values of the estimated parameters, where the k th row contains the parameters associated with time k , which are computed using the data values in the rows up to and including the row k .

`y_hat` contains the predicted output values such that the k th row of `y_hat` is computed based on the data values in the rows up to and including the row k .

Tip `y_hat` contains the adaptive predictions of the output and is useful for adaptive filtering applications, such as noise cancelation.

`nn` specified the model orders and delay according to the specific polynomial structure of the model. For example, `nn=[na nb nk]` for ARX models. For more information about specifying polynomial model orders and delays, see “Estimating Polynomial Models” on page 5-21.

adm and adg specify any of the four recursive algorithm, as described in “Algorithms for Recursive Estimation” on page 7-6.

The following table summarizes the recursive estimation commands supported by System Identification Toolbox. The command description indicates whether you can estimate single-input, single-output, multiple-input, and multiple-output, and time-series (no input) models. For details about each command, see the corresponding reference page.

Tip For ARX and AR models, use rarx. For single-input/single-output ARMAX or ARMA, Box-Jenkins, and Output-Error models, use rarmax, rbj, and roe, respectively.

Commands for Linear Recursive Estimation

Command	Description
rarmax	Estimate parameters of single-input/single-output ARMAX and ARMA models.
rarx	Estimate parameters of single- or multiple-input and single-output ARX and AR models. Does not support multiple-output system.
rbj	Estimate parameters of single-input/single-output Box-Jenkins models.
roe	Estimate parameters of single-input/single-output Output-Error models.

Commands for Linear Recursive Estimation (Continued)

Command	Description
rpem	<p data-bbox="795 371 1329 562">Estimate parameters of multiple-input and single-output ARMAX/ARMA, Box-Jenkins, or Output-Error models using the general recursive prediction-error algorithm for estimating the parameter gradient.</p> <hr data-bbox="795 614 1329 618"/> <p data-bbox="795 626 1307 690">Note Unlike pem, rpem does not support state-space models.</p> <hr data-bbox="795 699 1329 703"/>
rp1r	<p data-bbox="795 737 1310 894">Use as an alternative to rpem to estimate parameters of multiple-input and single-output systems when you want to use recursive pseudolinear regression method.</p>

Algorithms for Recursive Estimation

In this section...

“Types of Recursive Estimation Algorithms” on page 7-6

“General Form of Recursive Estimation Algorithm” on page 7-6

“Kalman Filter Algorithm” on page 7-8

“Forgetting Factor Algorithm” on page 7-10

“Unnormalized and Normalized Gradient Algorithms” on page 7-11

Types of Recursive Estimation Algorithms

System Identification Toolbox supports four recursive estimation algorithms.

- “General Form of Recursive Estimation Algorithm” on page 7-6
- “Kalman Filter Algorithm” on page 7-8
- “Forgetting Factor Algorithm” on page 7-10
- “Unnormalized and Normalized Gradient Algorithms” on page 7-11

You specify the type of recursive estimation algorithms as arguments `adm` and `adg` of the recursive estimation commands in “Commands for Recursive Estimation” on page 7-3.

For detailed information about these algorithms, see the corresponding chapter in *System Identification: Theory for the User* by Lennart Ljung (Prentice Hall PTR, Upper Saddle River, NJ, 1999).

General Form of Recursive Estimation Algorithm

The general recursive identification algorithm is given by the following equation:

$$\hat{\theta}(t) = \hat{\theta}(t-1) + K(t)(y(t) - \hat{y}(t))$$

$\hat{\theta}(t)$ is the parameter estimate at time t . $y(t)$ is the observed output at time t and $\hat{y}(t)$ is the prediction of $y(t)$ based on observations up to time $t-1$. The gain, $K(t)$, determines how much the current prediction error $y(t) - \hat{y}(t)$ affects the update of the parameter estimate. The estimation algorithms minimize the prediction-error term $y(t) - \hat{y}(t)$.

The gain has the following general form:

$$K(t) = Q(t)\psi(t)$$

The recursive algorithms supported by System Identification Toolbox differ based on different approaches for choosing the form of $Q(t)$ and computing $\psi(t)$, where $\psi(t)$ represents the gradient of the predicted model output $\hat{y}(t | \theta)$ with respect to the parameters θ .

The simplest way to visualize the role of the gradient $\psi(t)$ of the parameters, is to consider models with a linear-regression form:

$$y(t) = \psi^T(t)\theta_0(t) + e(t)$$

In this equation, $\psi(t)$ is the *regression vector* that is computed based on previous values of measured inputs and outputs. $\theta_0(t)$ represents the true parameters. $e(t)$ is the noise source (*innovations*), which is assumed to be white noise. The specific form of $\psi(t)$ depends on the structure of the polynomial model.

For linear regression equations, the predicted output is given by the following equation:

$$\hat{y}(t) = \psi^T(t)\hat{\theta}(t-1)$$

For models that do not have the linear regression form, it is not possible to compute exactly the predicted output and the gradient $\psi(t)$ for the current parameter estimate $\hat{\theta}(t-1)$. To learn how you can compute approximation for $\psi(t)$ and $\hat{\theta}(t-1)$ for general model structures, see the section on recursive prediction-error methods in *System Identification: Theory for the User* by Lennart Ljung (Prentice Hall PTR, Upper Saddle River, NJ, 1999).

Kalman Filter Algorithm

- “Mathematics of the Kalman Filter Algorithm” on page 7-8
- “Using the Kalman Filter Algorithm” on page 7-9

Mathematics of the Kalman Filter Algorithm

The following set of equations summarizes the *Kalman filter* adaptation algorithm:

$$\hat{\theta}(t) = \hat{\theta}(t-1) + K(t)(y(t) - \hat{y}(t))$$

$$\hat{y}(t) = \psi^T(t) \hat{\theta}(t-1)$$

$$K(t) = Q(t) \psi(t)$$

$$Q(t) = \frac{P(t-1)}{R_2 + \psi(t)^T P(t-1) \psi(t)}$$

$$P(t) = P(t-1) + R_1 - \frac{P(t-1) \psi(t) \psi(t)^T P(t-1)}{R_2 + \psi(t)^T P(t-1) \psi(t)}$$

This formulation assumes the linear-regression form of the model:

$$y(t) = \psi^T(t) \theta_0(t) + e(t)$$

The Kalman filter is used to obtain $Q(t)$.

This formulation also assumes that the true parameters $\theta_0(t)$ are described by a random walk:

$$\theta_0(t) = \theta_0(t-1) + w(t)$$

$w(t)$ is Gaussian white noise with the following covariance matrix, or *drift matrix* R_1 :

$$Ew(t)w(t)^T = R_1$$

R_2 is the variance of the innovations $e(t)$ in the following equation:

$$y(t) = \psi^T(t)\theta_0(t) + e(t)$$

The Kalman filter algorithm is entirely specified by the sequence of data $y(t)$, the gradient $\psi(t)$, R_1 , R_2 , and the initial conditions $\theta(t=0)$ (initial guess of the parameters) and $P(t=0)$ (covariance matrix that indicates parameters errors).

Note To simplify the inputs, you can scale R_1 , R_2 , and $P(t=0)$ of the original problem by the same value such that R_2 is equal to 1. This scaling does not affect the parameters estimates.

Using the Kalman Filter Algorithm

The general syntax for the command described in “Algorithms for Recursive Estimation” on page 7-6 is the following:

```
[params, y_hat] = command(data, nn, adm, adg)
```

To specify the Kalman filter algorithm, set adm to 'kf' and adg to the value of the drift matrix R_1 (described in “Mathematics of the Kalman Filter Algorithm” on page 7-8).

Forgetting Factor Algorithm

- “Mathematics of the Forgetting Factor Algorithm” on page 7-10
- “Using the Forgetting Factor Algorithm” on page 7-11

Mathematics of the Forgetting Factor Algorithm

The following set of equations summarizes the *forgetting factor* adaptation algorithm:

$$\hat{\theta}(t) = \hat{\theta}(t-1) + K(t)(y(t) - \hat{y}(t))$$

$$\hat{y}(t) = \psi^T(t) \hat{\theta}(t-1)$$

$$K(t) = Q(t) \psi(t)$$

$$Q(t) = P(t) = \frac{P(t-1)}{\lambda + \psi(t)^T P(t-1) \psi(t)}$$

$$P(t) = \frac{1}{\lambda} \left(P(t-1) - \frac{P(t-1) \psi(t) \psi(t)^T P(t-1)}{\lambda + \psi(t)^T P(t-1) \psi(t)} \right)$$

To obtain $Q(t)$, the following function is minimized at time t :

$$\sum_{k=1}^t \lambda^{t-k} e^2(k)$$

This approach discounts old measurements exponentially such that an observation that is τ samples old carries a weight that is equal to λ^τ times the weight of the most recent observation. $\tau = \frac{1}{1-\lambda}$ represents the *memory*

horizon of this algorithm. Measurements older than $\tau = 1/(1-\lambda)$ typically carry a weight that is less than about 0.3.

λ is called the forgetting factor and typically has a positive value between 0.97 and 0.995.

Note In the linear regression case, the forgetting factor algorithm is known as the *recursive least-squares* (RLS) algorithm. The forgetting factor algorithm for $\lambda = 1$ is equivalent to the Kalman filter algorithm with $R_1=0$ and $R_2=1$. For more information about the Kalman filter algorithm, see “Kalman Filter Algorithm” on page 7-8.

Using the Forgetting Factor Algorithm

The general syntax for the command described in “Algorithms for Recursive Estimation” on page 7-6 is the following:

```
[params,y_hat]=command(data,nn,adm,adg)
```

To specify the forgetting factor algorithm, set `adm` to 'ff' and `adg` to the value of the forgetting factor λ (described in “Mathematics of the Forgetting Factor Algorithm” on page 7-10).

Tip λ typically has a positive value from 0.97 to 0.995.

Unnormalized and Normalized Gradient Algorithms

- “Mathematics of the Unnormalized and Normalized Gradient Algorithm” on page 7-12
- “Using the Unnormalized and Normalized Gradient Algorithms” on page 7-12

Mathematics of the Unnormalized and Normalized Gradient Algorithm

In the linear regression case, the gradient methods are also known as the *least mean squares* (LMS) methods.

The following set of equations summarizes the *unnormalized gradient* and *normalized gradient* adaptation algorithm:

$$\hat{\theta}(t) = \hat{\theta}(t-1) + K(t)(y(t) - \hat{y}(t))$$

$$\hat{y}(t) = \psi^T(t)\hat{\theta}(t-1)$$

$$K(t) = Q(t)\psi(t)$$

In the unnormalized gradient approach, $Q(t)$ is the product of the gain γ and the identity matrix:

$$Q(t) = \gamma I$$

In the normalized gradient approach, $Q(t)$ is the product of the gain γ , and the identity matrix is normalized by the magnitude of the gradient $\psi(t)$:

$$Q(t) = \frac{\gamma}{|\psi(t)|^2} I$$

These choices of $Q(t)$ update the parameters in the negative gradient direction, where the gradient is computed with respect to the parameters.

Using the Unnormalized and Normalized Gradient Algorithms

The general syntax for the command described in “Algorithms for Recursive Estimation” on page 7-6 is the following:

```
[params, y_hat] = command(data, nn, adm, adg)
```

To specify the unnormalized gain algorithm, set `adm` to 'ug' and `adg` to the value of the gain γ (described in “Mathematics of the Unnormalized and Normalized Gradient Algorithm” on page 7-12).

To specify the normalized gain algorithm, set `adm` to 'ng' and `adg` to the value of the gain γ .

Data Segmentation

For systems that exhibit abrupt changes while the data is being collected, you might want to develop models for separate data segments such that the system does not change during a particular data segment. Such modeling requires identification of the time instants when the changes occur in the system, breaking up the data into segments according to these time instants, and identification of models for the different data segments.

The following cases are typical applications for *data segmentation*:

- Segmentation of speech signals, where each data segment corresponds to a phonem.
- Detection of trend breaks in time series.
- Failure detection, where the data segments correspond to operation with and without failure.
- Estimating different working modes of a system.

Use `segment` to build polynomial models, such as ARX, ARMAX, AR, and ARMA, so that the model parameters are piece-wise constant over time. For detailed information about this command, see the corresponding reference page.

To see an example of using data segmentation, run the Recursive Estimation and Data Segmentation demonstration by typing to the following command at the MATLAB prompt:

```
iddemo5
```


Estimating Nonlinear Black-Box Models

Supported Nonlinear Black-Box Models (p. 8-2)

Types of supported nonlinear black-box models.

Supported Data for Estimating Nonlinear Black-Box Models (p. 8-3)

Types of supported data for estimating nonlinear black-box models.

Estimating Nonlinear ARX Models (p. 8-4)

How to estimate nonlinear ARX models using the System Identification Tool GUI or the `nlarx` command.

Estimating Hammerstein-Wiener Models (p. 8-14)

How to estimate Hammerstein-Wiener models using the System Identification Tool GUI or the `n1hw` command.

Supported Nonlinearity Estimators (p. 8-22)

Nonlinearity estimators used in nonlinear ARX and Hammerstein-Wiener models.

Refining Nonlinear Black-Box Models (p. 8-25)

How to refine a nonlinear black-box model after estimation.

Next Steps After Estimating Nonlinear Black-Box Models (p. 8-27)

Next steps for estimated nonlinear black-box models.

Supported Nonlinear Black-Box Models

You can use System Identification Toolbox to estimate the following types of nonlinear black-box models:

- Nonlinear ARX models
- Hammerstein-Wiener models

You can estimate these models both in the System Identification Tool GUI and at the command line.

For a quick introduction, see “Tutorial: Estimating Nonlinear Black-Box Models”.

Supported Data for Estimating Nonlinear Black-Box Models

You can estimate discrete-time black-box models for data with the following characteristics:

- Time-domain input-output or time-series data.

Note Time series are supported for nonlinear ARX models only.

- Single-output or multiple-output data.

Before you begin estimating models, import your data into MATLAB and represent the data in either of the following ways:

- **In the System Identification Tool GUI.** Import the data into the GUI, as described in “Importing Data into the GUI” on page 2-14.
- **At the command line.** Represent your data as an `iddata` or `idfrd` object.

To examine the data features, plot the data on a time plot or an estimated frequency-response plot. You can preprocess your data by interpolating missing values, filtering to emphasize a specific frequency range, or resampling using a different time interval. For more information about available data plots and data-preprocessing operations, see Chapter 2, “Preparing Data”.

Note For nonlinear modeling, do not remove offsets and linear trends from the measured signals.

Estimating Nonlinear ARX Models

In this section...

“Supported Data for Nonlinear ARX Models” on page 8-4

“Definition of the Nonlinear ARX Model” on page 8-4

“Using Regressors” on page 8-5

“Nonlinearity Estimators for Nonlinear ARX Models” on page 8-8

“How to Estimate Nonlinear ARX Models in the GUI” on page 8-8

“How to Estimate Nonlinear ARX Models at the Command Line” on page 8-10

Supported Data for Nonlinear ARX Models


You can estimate both continuous-time and discrete-time nonlinear ARX models for data with the following characteristics:

- Time-domain input-output data or time-series data
- Single-output or multiple-output data

For more information about representing your data for system identification, see Chapter 2, “Preparing Data”.

Definition of the Nonlinear ARX Model

The nonlinear ARX structure models dynamic systems using a parallel combination of nonlinear and linear blocks, as shown in the following figure.



The nonlinear and linear functions are expressed in terms of variables called *regressors*, which are functions of measured input-output data. For more information about regressors, see “Using Regressors” on page 8-5.

The predicted output $\hat{y}(t)$ of a nonlinear model at time t is given by the following general equation:

$$\hat{y}(t) = F(x(t))$$

where $x(t)$ represents the regressors. F is a nonlinear regression command, which is approximated by the nonlinearity estimators. For a list of nonlinearity estimators supported by nonlinear ARX models, see “Nonlinearity Estimators for Nonlinear ARX Models” on page 8-8.

The command F can include both linear and nonlinear functions of $x(t)$, as shown in the previous diagram. You can specify which regressors to use as inputs to the nonlinear block.

The following equation provides a general description of F :

$$F(x) = \sum_{k=1}^d \alpha_k \kappa(\beta_k (x - \gamma_k))$$

where κ is the unit nonlinear command, d is the number of nonlinearity units, and α_k , β_k , and γ_k are the parameters of the nonlinearity estimator.

Using Regressors

System Identification Toolbox supports the following types of regressors for nonlinear ARX models:

- *Standard regressors* — Past input $u(t)$ and output signals $y(t)$, computed automatically as delay transformations for specified model orders.
- *Custom regressors* — Products, powers, and other MATLAB expressions of input and output variables that you specify.

Specifying Model Order and Delays

You must specify the following model orders for computing standard regressors:

- n_a — The number of past output terms used to predict the current output.
- n_b — The number of past input terms used to predict the current output.
- n_k — The delay from input to the output in terms of the number of samples. This value defines the least delayed input regressor.

The meaning of n_a and n_b is similar to the linear-ARX model parameters in the sense that n_a represents the number of output terms and n_b represents the number of input terms. n_k represents the minimum input delay from an input to an output. For more information about the linear ARX model structure, see “What Are Black-Box Polynomial Models?” on page 5-21.

Note The total number of regressors in the model must be greater than zero. If you only need to use custom regressors, set $n_a=n_b=n_k=0$ to omit creating standard regressors.

Example – Relationship Between Regressors, Model Orders, and Delays

This example describes how the model orders and delays you specify relate to computing the regressors.

Suppose that you specify a nonlinear ARX model with a minimum of a two-sample input delay and the number of input terms is $n_b=2$. System Identification Toolbox computes the following standard regressors from the input signal:

- $u(t-2)$
- $u(t-3)$

If you specify that the number of output terms is $n_a=4$, System Identification Toolbox computes the following standard regressors from the output signals:

- $y(t-1)$
- $y(t-2)$
- $y(t-3)$
- $y(t-4)$

Note You cannot modify the minimum output delay—it is set to one sample.

If you have physical insight that your current output depends on specific delayed inputs and outputs, select the appropriate model orders to compute the required regressors.

Using Custom Regressors

In general, custom regressors are nonlinear functions of the standard regressors. You can specify custom regressors, such as $\tan(u(t-1))$, $u(t-1)^2$, or $u(t-1)y(t-3)$.

In the System Identification Tool GUI. You can create custom regressors in the Model Regressors dialog box. For more information, see “How to Estimate Nonlinear ARX Models in the GUI” on page 8-8.

At the command line. Use the `customreg` or `polyreg` command to construct custom regressors in terms of input-output variables. For more information, see the corresponding reference page.

The linear block includes all standard and custom regressors. However, you can include specific standard and custom regressors in your nonlinear block to fine-tune the model structure.

To get a linear-in-the-parameters ARX model structure, you can exclude the nonlinear block from the model structure completely. When using only a linear block with custom regressors, you can create the simplest types of nonlinear models. In this case, the custom regressors capture the nonlinearities and the estimation routine computes the weights of the standard and custom regressors in the linear block to predict the output.

Nonlinearity Estimators for Nonlinear ARX Models

Nonlinear ARX models support the following nonlinearity estimators:

- Sigmoid Network
- Tree Partition
- Wavelet Network
- Custom Network
- Neural Network

Note You must have Neural Network Toolbox to specify neural network objects as nonlinearities.

You can exclude the nonlinearity command from the model structure. In this case, the model includes all standard and custom regressors and is linear in the parameters.

In the System Identification Tool GUI. You can omit the nonlinear block by selecting None for the **Nonlinearity**.

At the command line. You can omit the nonlinear block by setting the Nonlinearity property value to 'Linear'. For more information, see the `nlarx` and `idnlarx` reference pages.

For a description of each nonlinearity estimator, see “Supported Nonlinearity Estimators” on page 8-22.

How to Estimate Nonlinear ARX Models in the GUI

You must have already imported your data into the System Identification Tool GUI. For more information about preparing your data, see Chapter 2, “Preparing Data”.

To estimate a nonlinear ARX model in the System Identification Tool GUI:

1 In the System Identification Tool GUI, select **Estimate > Nonlinear models** to open the Nonlinear Models dialog box. The **Model Type** tab is selected.

2 In the **Model Structure** list, select Nonlinear ARX.

This action updates the options in the Nonlinear Models dialog box to correspond to this model structure. For information about this model structure, see “Definition of the Nonlinear ARX Model” on page 8-4.

3 In the **Model name** field, edit the name of the model, or keep the default name. The name of the model should be unique in the Model Board.

4 (Optional) If you want to try refining a previously estimated model, select the name of this model in the **Initial model** list.

Note The model structure and algorithm properties of the initial model populate the fields in the Nonlinear Models dialog box.

A model is available in the **Initial model** list under the following conditions:

- The model exists in the System Identification Tool GUI.
 - The number of model inputs and outputs matches the dimensions of the **Working Data** (estimation data) you selected in the System Identification Tool GUI.
- 5** Keep the default settings in the Nonlinear Models dialog box that specify the model structure and the algorithm, or modify the following settings:
- In the **Regressors** tab, change the input delay of the input signals.
To gain insight into possible input delay values, click **Infer Input Delay**. This action opens the Infer Input Delay dialog box.
 - In the **Regressors** tab, change the number of terms to include in the nonlinear block.
 - In the **Regressors** tab, click **Edit Regressors** to select which regressors are included in the nonlinear block. This action opens the Model

Regressors dialog box. You can also use this dialog box create custom regressors.

- In the **Model Properties** tab, select and configure the nonlinearity estimator, and choose whether to include the linear block. To use all standard and custom regressors in the linear block only, you can exclude the nonlinear block by choosing None.

For more information about the available options, click **Help** in the Nonlinear Models dialog box to open the GUI help.

- 6 Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.

The **Estimation** tab displays the estimation progress and results.

- 7 To plot this model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. If you get an inaccurate fit, try estimating a new model with different orders or nonlinearity estimator. For more information about validating models, see Chapter 10, “Validating and Analyzing Models”.

- 8 For further analysis, export the model to the MATLAB Workspace browser by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

How to Estimate Nonlinear ARX Models at the Command Line

- “General nlarx Syntax” on page 8-10
- “Example – Using nlarx to Estimate Nonlinear ARX Models ” on page 8-12

General nlarx Syntax

You can estimate nonlinear ARX models using `nlarx`. The resulting models are stored as `idnlarx` model objects.

Use the following general syntax to both configure and estimate nonlinear ARX models:

```
m = nlarx(data, 'na', na,
```

```

    'nb',nb,
    'nk',nk,
    Nonlinearity,
    'Property1',Value1,...,
    'PropertyN',ValueN)

```

where `data` is the estimation data. `na`, `nb`, and `nk` specify the model orders and delays. For more information about model orders, see “Specifying Model Order and Delays” on page 8-6.

`Nonlinearity` specifies the nonlinearity estimator object as `'sigmoidnet'`, `'wavenet'`, `'treepartition'`, `'customnet'`, `'neuralnet'`, or `'linear'`.

The property-value pairs specify any `idnlarx` model properties that configure the estimation algorithm. You can enter all model property-value pairs and top-level algorithm properties as a comma-separated list in `nlarx`.

For multiple inputs and outputs, `na`, `nb`, and `nk` are described in “Options for Multiple-Input and Multiple-Output ARX Orders” on page 5-44.

You can specify different nonlinearity estimators for different output channels by setting `Nonlinearity` to an object array. For example:

```

m = nlarx(data,[[2 1; 0 1] [2;1] [1;1]],...
           [wavenet;sigmoidnet('num',7)])

```

To specify the same nonlinearity for all outputs, set `Nonlinearity` to a single nonlinearity estimator. For example:

```

m = nlarx(data,[[2 1; 0 1] [2;1] [1;1]],...
           sigmoidnet('num',7))

```

For detailed information about the `nlarx` and `idnlarx` properties and values, see the corresponding reference page.

For more information about validating models, see Chapter 10, “Validating and Analyzing Models”.

Note You do not need to construct the model object using `idnlarx` before estimation.

You can also use `pem` to refine parameter estimates of an existing nonlinear ARX model, as described in “Refining Nonlinear Black-Box Models” on page 8-25.

Example – Using `nlarx` to Estimate Nonlinear ARX Models

This example uses `nlarx` to estimate a nonlinear ARX model for the two-tank system. The data for this system is described in “Tutorial: Estimating Nonlinear Black-Box Models”.

Prepare the data for estimation using the following commands:

```
load twotankdata
z = iddata(y, u, 0.2);
ze = z(1:1000); zv = z(1001:3000);
```

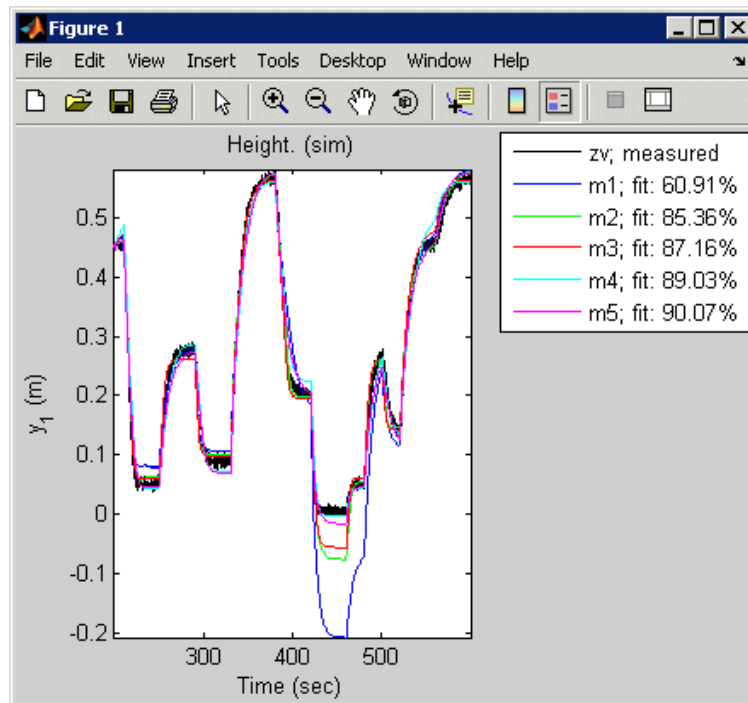
Estimate several models using different model orders, delays, and nonlinearity settings:

```
m1 = nlarx(ze,[2 2 1], 'wav');
m2 = nlarx(ze,[2 2 3], wavenet);
m3 = nlarx(ze,[2 2 3], wavenet('num',8));
m4 = nlarx(ze,[2 2 3], wavenet('num',8),...
          'nlr', [1 2]);
m5 = nlarx(ze,[2 2 3], sigmoidnet('num',14),...
          'nlr',[1 2]);
```

Compare the resulting models by plotting the model outputs on top of the measured output:

```
compare(zv, m1,m2,m3,m4,m5)
```

MATLAB responds with the following plot.



Estimating Hammerstein-Wiener Models

In this section...

“Supported Data for Estimating Hammerstein-Wiener Models” on page 8-14

“Definition of the Hammerstein-Wiener Model” on page 8-14

“Nonlinearity Estimators for Hammerstein-Wiener Models” on page 8-16

“How to Estimate Hammerstein-Wiener Models in the GUI” on page 8-17

“How to Estimate Hammerstein-Wiener Models at the Command Line” on page 8-18

Supported Data for Estimating Hammerstein-Wiener Models

You can estimate both continuous-time and discrete-time Hammerstein-Wiener models for data with the following characteristics:

- Time-domain input-output data

Note Hammerstein-Wiener models do not support time-series data, where there is no input.

- Single-output or multiple-output data

Definition of the Hammerstein-Wiener Model

The Hammerstein-Wiener structure models dynamic systems using one or two static nonlinear blocks in series with a linear block. For example, you can use a static nonlinearity to represent saturation or dead-zone behavior.

The input signal passes through the first nonlinear block, a linear block, and a second nonlinear block to produce the output signal, as shown in the following figure.



This model structure represents a nonlinear system as a linear system that is modified by static input and output nonlinearities. Thus, the linear model provides a reference point for estimating the nonlinear contributions in a system.

The following general equation describes the Hammerstein-Wiener structure:

$$\begin{aligned} w(t) &= f(u(t)) \\ x(t) &= \frac{B_{j,i}(q)}{F_{j,i}(q)} w(t) \\ y(t) &= h(x(t)) \end{aligned}$$

which contains the following variables:

- $u(t)$ and $y(t)$ are the inputs and outputs for the system, respectively.
- f and h are nonlinear functions that corresponding to the input and output nonlinearities, respectively.

For multiple inputs and multiple outputs, f and h are defined component-wise.

- $w(t)$ and $x(t)$ are internal variables.
 $w(t)$ has the same dimension as $u(t)$. $x(t)$ has the same dimension as $y(t)$.
- $B(q)$ and $F(q)$ in the linear dynamic block are polynomials in the backward shift operator, as described in “What Are Black-Box Polynomial Models?” on page 5-21.

For n_y outputs and n_u inputs, the linear block is a transfer command matrix containing entries in the following form:

$$\frac{B_{j,i}(q)}{F_{j,i}(q)}$$

where $j = 1, 2, \dots, n_y$ and $i = 1, 2, \dots, n_u$.

If only the input nonlinearity is present, the model is called the Hammerstein model. If only the output nonlinearity is present, the model is called the Wiener model.

You must specify the following model orders for the linear block:

- n_b — The number of zeros plus one.
- n_f — The number of poles.
- n_k — The delay from input to the output in terms of the number of samples.

For n_y outputs and n_u inputs, n_b , n_f , and n_k are n_y -by- n_u matrices. You can specify a nonlinearity for only certain inputs and outputs, and exclude the nonlinearity for other inputs and outputs.

Nonlinearity Estimators for Hammerstein-Wiener Models

Hammerstein-Wiener models support the following nonlinearity estimators for estimating the parameters of its input and output nonlinear blocks:

- Dead Zone
- Piecewise Linear
- Saturation
- Sigmoid Network
- Wavelet Network

You can exclude either the input nonlinearity or the output nonlinear from the model structure.

In the System Identification Tool GUI. Exclude a nonlinearity for a specific channel by selecting None.

At the command line. Exclude a nonlinearity for a specific channel by specifying the `unitgain` value for the `InputNonlinearity` or `OutputNonlinearity` properties. For more information about estimation objects and their properties, see `nlhw` and `idnlhw` reference page.

For a description of each nonlinearity estimator, see “Supported Nonlinearity Estimators” on page 8-22.

How to Estimate Hammerstein-Wiener Models in the GUI

You must have already imported your data into the System Identification Tool GUI, as described in Chapter 2, “Preparing Data”.

To estimate a Hammerstein-Wiener model in the System Identification Tool GUI:

- 1 In the System Identification Tool GUI, select **Estimate > Nonlinear models** to open the Nonlinear Models dialog box. The **Model Type** tab is shown.
- 2 In the **Model Structure** list, select Hammerstein-Wiener.

This action updates the options in the Nonlinear Models dialog box to correspond to this model structure. For information about this model structure, see “Definition of the Hammerstein-Wiener Model” on page 8-14.

- 3 In the **Model name** field, edit the name of the model, or keep the default name. The name of the model should be unique in the Model Board.
- 4 (Optional) If you want to try refining a previously estimated model, select the name of this model in the **Initial model** list.

Note The model structure and algorithm properties of the initial model populate the fields in the Nonlinear Models dialog box.

A model is available in the **Initial model** list under the following conditions:

- The model exists in the System Identification Tool GUI.
 - The number of model inputs and outputs matches the dimensions of the **Working Data** (estimation data) you selected in the System Identification Tool GUI.
- 5 Keep the default settings in the Nonlinear Models dialog box that specify the model structure and the algorithm, or modify the following settings:

- In the **I/O Nonlinearity** tab, specify whether to include or exclude the input and output nonlinearities. For multiple-input and multiple-output systems, you can choose to apply nonlinearities only to specific input and output channels.
- In the **I/O Nonlinearity** tab, change the input and output nonlinearity types and configure the nonlinearity settings.
- In the **Linear Block** tab, specify the model orders and delays. To gain insight into possible delays, click **Infer Input Delay**.

For more information about the available options, click **Help** in the Nonlinear Models dialog box to open the GUI help.

- 6 Click **Estimate** to add this model to the Model Board in the System Identification Tool GUI.

The **Estimation** tab displays the estimation progress and results.

- 7 To plot this model, select the appropriate check box in the **Model Views** area of the System Identification Tool GUI. For more information about working with plots and validating models, see Chapter 10, “Validating and Analyzing Models”.

If you get an inaccurate fit, try estimating a new model with different orders or nonlinearity estimator.

You can export the model to the MATLAB Workspace browser for further analysis by dragging it to the **To Workspace** rectangle in the System Identification Tool GUI.

How to Estimate Hammerstein-Wiener Models at the Command Line

- “General nlhw Syntax” on page 8-19
- “Example – Using nlhw to Estimate Hammerstein-Wiener Models ” on page 8-20

General nlhw Syntax

You can estimate Hammerstein-Wiener models using the `nlhw` command. The resulting models are stored as `idnlhw` model objects.

Use the following general syntax to both configure and estimate Hammerstein-Wiener models:

```
m = nlhw(data, 'nb', nb,
          'nf', nf,
          'nk', nk,
          InputNonlinearity,
          OutputNonlinearity,
          'Property1', Value1, ...,
          'PropertyN', ValueN)
```

where `data` is the estimation data. `nb`, `nf`, and `nk` specify the orders and delays of the linear OE model. For more information about model orders, see “Definition of the Hammerstein-Wiener Model” on page 8-14.

`InputNonlinearity` specifies the input static nonlinearity estimator object as `'pwnlinear'`, `'deadzone'`, `'saturation'`, `'sigmoidnet'`, `'wavenet'`, `'customnet'`, or `'unitgain'`. Similarly, `OutputNonlinearity` specifies the output static nonlinearity estimator object.

The property-value pairs specify any `idnlhw` model properties that configure the estimation algorithm. You can enter all model property-value pairs and top-level algorithm properties as a comma-separated list in `nlhw`. For example, you can control the iterative search for a best fit using the following properties:

```
m = nlhw(data, 'nb', nb,
          'nf', nf,
          'nk', nk,
          InputNonlinearity,
          OutputNonlinearity,
          'MaxIter', N,
          'Tolerance', tol,
          'LimitError', lim,
          'Trace', 'on')
```

For n_u inputs and n_y outputs, n_a , n_b , and n_k are n_y -by- n_u matrices whose i - j th entry specifies the order and delay of the transfer command from the j th input to the i th output.

You can specify different nonlinearity estimators for different output channels by setting `InputNonlinearity` or `OutputNonlinearity` to an object array. For example:

```
m = nlarx(data,[nb,nf,nk],...
           [sigmoidnet;pwlinear],...
           [])
```

For detailed information about `nlhw` and `idnlhw`, see the corresponding reference page.

For more information about validating your models, see Chapter 10, “Validating and Analyzing Models”.

Note You do not need to construct the model object using `idnlhw` before estimation.

You can also use `pem` to refine parameter estimates of an existing Hammerstein-Wiener model, as described in “Refining Nonlinear Black-Box Models” on page 8-25.

Example – Using `nlhw` to Estimate Hammerstein-Wiener Models

This example uses `nlhw` to estimate a Hammerstein-Wiener model for the two-tank system. The data for this system is described in “Tutorial: Estimating Nonlinear Black-Box Models”.

Prepare the data for estimation using the following commands:

```
load twotankdata
z = iddata(y, u, 0.2);
ze = z(1:1000); zv = z(1001:3000);
```

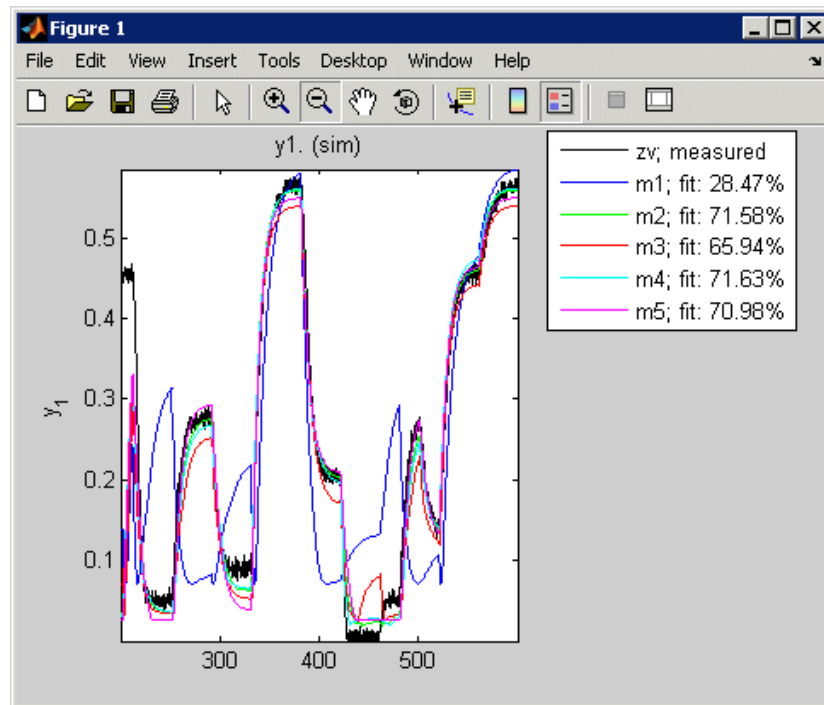
Estimate several models using different model orders, delays, and nonlinearity settings:

```
m1 = nlhw(ze,[2 3 1],'pw1','pw1');
m2 = nlhw(ze,[2 2 3],'pw1','pw1');
m3 = nlhw(ze,[2 2 3], pwlinear('num',13),...
        pwlinear('num',10));
m4 = nlhw(ze,[2 2 3], sigmoidnet('num',2),...
        pwlinear('num',10));
m5 = nlhw(ze,[2 2 3], 'dead','sat');
```

Compare the resulting models by plotting the model outputs on top of the measured output:

```
compare(zv,m1,m2,m3,m4,m5)
```

MATLAB responds with the following plot.



Supported Nonlinearity Estimators

In this section...
“Types of Nonlinearity Estimators” on page 8-22
“Creating Custom Nonlinearities” on page 8-23

Types of Nonlinearity Estimators

When configuring the nonlinear ARX and Hammerstein-Wiener models for estimation, you must specify a mathematical structure for the nonlinear portion of the model.

If you are working in the System Identification Tool GUI, specify the nonlinearity type by name when you configure the nonlinear model structure. If you are estimating or constructing a nonlinear model at the command line instead, specify the nonlinearity as an argument in the `nlarx` or `nlhw` estimation command.

The following table summarizes supported nonlinearities in System Identification Toolbox for each type of nonlinear model. For a description of each nonlinearity, see the references page for the corresponding nonlinearity object.

Nonlinearity	Object Name	Supported Model Type	Supports Multiple Inputs?
Custom Network (user-defined)	<code>customnet</code>	Hammerstein-Wiener and Nonlinear ARX	Yes
Dead Zone	<code>deadzone</code>	Hammerstein-Wiener	No
Neural Network	<code>neuralnet</code>	Nonlinear ARX	Yes
Piecewise Linear	<code>pwlinear</code>	Hammerstein-Wiener	No
Polynomial	<code>poly1d</code>	Hammerstein-Wiener	
Saturation	<code>saturation</code>	Hammerstein-Wiener	No
Sigmoid Network	<code>sigmoidnet</code>	Hammerstein-Wiener and Nonlinear ARX	Yes

Nonlinearity	Object Name	Supported Model Type	Supports Multiple Inputs?
Tree Partition	treepartition	Nonlinear ARX	Yes
Wavelet Network	wavenet	Hammerstein-Wiener and Nonlinear ARX	Yes

The Neural Network nonlinearity lets you import a network object you created in Neural Network Toolbox.

The nonlinearity estimators `deadzone`, `poly1d`, `pwnlinear`, and `saturation` are optimized for estimating Hammerstein-Wiener models.

Creating Custom Nonlinearities

System Identification Toolbox lets you create custom nonlinearities for nonlinear ARX and Hammerstein-Wiener models.

A custom nonlinearity uses a unit function that you define. This custom unit command uses a weighted sum of inputs to compute a scalar output.

You can use a combination of these unit functions to approximate the nonlinearity.

Note Hammerstein-Wiener models require that your custom nonlinearity have one input and one output.

```

command [f, g, a] = gaussunit(x)
%GAUSSUNIT example of customnet unit command
%
%[f, g, a] = GAUSSUNIT(x)
%
% x: unit command variable
% f: unit command value
% g: df/dx
% a: unit active range (g(x) is significantly
% nonzero in the interval [-a a])

```

```
%  
% The unit command must be vectorized:  
% for a vector or matrix x, the output  
% arguments f and g must have the same size as x  
% f and g are computed element by element.  
  
f = exp(-x.*x);  
  
if nargin>1  
    g = - 2*x .* f;  
    a = 0.2;  
end
```


Refining Nonlinear Black-Box Models

In this section...

“How to Refine Nonlinear Black-Box Models in the GUI” on page 8-25

“How to Refine Nonlinear Black-Box Models at the Command Line” on page 8-26

How to Refine Nonlinear Black-Box Models in the GUI

The following procedure assumes that the model you want to refine is already in the System Identification Tool GUI. You might have estimated this model in the current session or imported the model into the GUI from the MATLAB Workspace browser. For more information about estimating nonlinear black-box models, see Chapter 8, “Estimating Nonlinear Black-Box Models”.

To refine your model:

- 1 In the System Identification Tool GUI, verify that you have the correct data set in the **Working Data** area for refining your model.

If you are using a different data set than the one you used to estimate the model, drag the correct data set into the **Working Data** area. For more information about specifying estimation data, see “Specifying Estimation and Validation Data” on page 1-11.

- 2 Select **Estimate > Nonlinear models** to open the Nonlinear Models dialog box, if this dialog box is not already open.

- 3 In the Nonlinear Models dialog box, select the model you want to refine in the **Initial model** list.

The list includes only those models that have the selected **Model structure** and the same number of inputs and outputs as the estimation data in **Working Data** area.

Any settings in the Nonlinear Models dialog box related to model structure and estimation algorithms are overridden by the selected initial model.

- 4 Click **Estimate** to refine the model.
- 5 Validate the new model, as described in Chapter 10, “Validating and Analyzing Models”.

Tip To continue refining the model directly from the **Estimation** tab, select the **Use last estimated model as initial model for the next estimation** check box, and click **Estimate**. This action automatically selects the most recent model in the **Initial model** list in the **Model Type** tab.

How to Refine Nonlinear Black-Box Models at the Command Line

If you are working at the command line, you can use `pem` to refine nonlinear black-box models.

The general syntax for refining initial models is as follows:

```
m = pem(data,init_model)
```

`pem` uses the properties of the initial model unless you specify different properties. For more information about specifying model properties directly in the estimator, see “Specifying Model Properties in the Estimator” on page 3-15.

Next Steps After Estimating Nonlinear Black-Box Models

After estimating nonlinear black-box models, you can perform the following operations:

- View parameter values, standard deviations of the parameters, loss command, and Akaike's Final Prediction Error (FPE) Criterion at the command line using the `present` command.
- Simulate the model using the `sim` command.
- Predict the model output using the `predict` command.
- Linearize nonlinear ARX and Hammerstein-Wiener models using `lintan` or `linapp`. `lintan` provides a small-signal tangent linearization about a specific operating point. `linapp` computes a linear approximation for a nonlinear model of a given input. For more information about these commands, see the corresponding reference page.

After linearization, you can perform linear analysis on your models and use the models with Control System Toolbox. For more information, see "Using Models with Control System Toolbox" on page 13-2.

Estimating Grey-Box Models

Supported Grey-Box Models (p. 9-2)	Types of supported grey-box models.
Data Supported by Grey-Box Models (p. 9-3)	Types of supported data for estimating grey-box models.
Before You Begin (p. 9-4)	Preparatory steps for estimating grey-box models.
Estimating Linear Grey-Box Models (p. 9-5)	How to define and estimate linear grey-box models at the command line.
Estimating Nonlinear Grey-Box Models (p. 9-13)	How to define and estimate nonlinear grey-box models at the command line.
After Estimating Grey-Box Models (p. 9-19)	Next steps for estimated grey-box models.

Supported Grey-Box Models

If you understand the physics of your system and can represent the system using ordinary differential or difference equations (ODEs) with unknown parameters, then you can use System Identification Toolbox to perform linear or nonlinear grey-box modeling. *Grey-box model* ODEs specify the mathematical structure of the model explicitly, including couplings between parameters and known parameter values. Grey-box modeling is useful when you know the relationships between variables, constraints on model behavior, or explicit equations representing system dynamics.

System Identification Toolbox supports both continuous-time and discrete-time models. However, because most laws of physics are expressed in continuous time, it is easier to construct models with physical insight in continuous time, rather than in discrete time.

In addition to dynamic input-output models, you can also create time-series models that have no inputs and static models that have no states.

If it is too difficult to describe your system using known physical laws, you can use System Identification Toolbox to perform black-box modeling.

Data Supported by Grey-Box Models

You can estimate both continuous-time or discrete-time grey-box models for data with the following characteristics:

- Time-domain or frequency-domain data

Note Nonlinear grey-box models support only time-domain data.

- Single-output or multiple-output data

Before You Begin

Before you begin estimating models, you must import your data into the MATLAB Workspace browser. If you are using the System Identification Tool GUI, then import the data into the GUI to make the data available to System Identification Toolbox. However, if you prefer to work at the command line, then represent your data as an `iddata` or `idfrd` object. For more information about preparing data for identification, see Chapter 2, “Preparing Data”.

The first step in grey-box modeling is to define the model structure in an M-file or MEX-file.

For linear models. Write an M-file that returns state-space matrices as a command of user-defined parameters and information about the model. For examples, see “Estimating Linear Grey-Box Models” on page 9-5.

Create an `idgrey` object based on this M-file. For more information, see “Specifying the Linear Grey-Box Model Structure” on page 9-5.

For nonlinear models. Write an M-file or MEX-file to return the first-order derivatives of the states and output values as a command of the states, inputs, time, parameters, and auxiliary variables. For information about the file structure, see “Estimating Nonlinear Grey-Box Models” on page 9-13

Create an `idn1grey` object based on this M-file or MEX-file.

Estimating Linear Grey-Box Models

In this section...

“Specifying the Linear Grey-Box Model Structure” on page 9-5

“Example – Representing a Grey-Box Model in an M-File” on page 9-6

“Example – Estimating a Continuous-Time Grey-Box Model for Heat Diffusion” on page 9-8

“Example – Estimating a Discrete-Time Grey-Box Model with Parameterized Disturbance” on page 9-11

Specifying the Linear Grey-Box Model Structure

You can estimate linear discrete-time and continuous-time grey-box models for arbitrary ordinary differential or difference equations using single-output and multiple-output time-domain data, or time-series data that has no measured inputs.

You must represent your system equations in state-space form. *State-space models* use state variables $x(t)$ to describe a system as a set of first-order differential equations, rather than by one or more n th-order differential equations.

In continuous-time, the state-space description has the following form:

$$\begin{aligned}\dot{x}(t) &= Fx(t) + Gu(t) + \tilde{K}w(t) \\ y(t) &= Hx(t) + Du(t) + w(t) \\ x(0) &= x_0\end{aligned}$$

The discrete-time state-space model structure is often written in the *innovations form*:

$$\begin{aligned}x(kT + T) &= Ax(kT) + Bu(kT) + Ke(kT) \\ y(kT) &= Cx(kT) + Du(kT) + e(kT) \\ x(0) &= x_0\end{aligned}$$

Use the following format to implement the linear grey-box model in an M-file:

```
[A,B,C,D,K,x0] = myfunc(par,T,CDmfile,aux)
```

where the matrices A, B, C, D, K, and x0 represent both continuous-time and discrete-time descriptions, myfunc is the name of the M-file, par contains the parameters as a column vector, and T is the sampling interval. aux contains auxiliary variables in your system. You use auxiliary variables to vary system parameters at the input to the command, and avoid editing the M-file.

CDmfile is an optional argument that describes whether the resulting state-space matrices are in discrete time or continuous time. By default, CDmfile='cd', which means that the sampling interval property of the model Ts determines whether the model is continuous or discrete in time.

For more information about these arguments, see the idgrey reference page.

For more information about validating your models, see Chapter 10, “Validating and Analyzing Models”.

Example – Representing a Grey-Box Model in an M-File

In this example, you represent the structure of the following continuous-time model:

$$\begin{aligned}\dot{x} &= \begin{bmatrix} 0 & 1 \\ 0 & \theta_1 \end{bmatrix} x(t) + \begin{bmatrix} 0 \\ \theta_2 \end{bmatrix} u(t) \\ y(t) &= \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} x(t) + e(t) \\ x(0) &= \begin{bmatrix} \theta_3 \\ 0 \end{bmatrix}\end{aligned}$$

This equation represents an electrical motor, where $y_1(t) = x_1(t)$ is the angular position of the motor shaft, and $y_2(t) = x_2(t)$ is the angular velocity.

The parameter $-\theta_1$ is the inverse time constant of the motor, and $-\theta_2/\theta_1$ is the static gain from the input to the angular velocity.

The motor is at rest at $t=0$, but its angular position θ_3 is unknown. Suppose that the approximate nominal values of the unknown parameters are $\theta_1 = -1$ and $\theta_2 = 0.25$. The variance of the errors in the position measurement is 0.01, and the variance in the angular velocity measurements is 0.1. For more information about this example, see the section on state-space models in *System Identification: Theory for the User*, Second Edition, by Lennart Ljung, Prentice Hall PTR, 1999.

The continuous-time state-space model structure is defined by the following equation:

$$\begin{aligned}\dot{x}(t) &= Fx(t) + Gu(t) + \tilde{K}w(t) \\ y(t) &= Hx(t) + Du(t) + w(t) \\ x(0) &= x_0\end{aligned}$$

To prepare this model for identification:

- 1 Create the following M-file to represent the model structure in this example:

```
command [A,B,C,D,K,x0] = myfunc(par,T,aux)
A = [0 1; 0 par(1)];
B = [0;par(2)];
C = eye(2);
D = zeros(2,2);
K = zeros(2,1);
x0 =[par(3);0];
```

- 2** Use the following syntax to define an `idgrey` model object based on the `myfunc` M-file:

```
m = idgrey('myfunc',par,'c',T,aux)
```

where `par` represents user-defined parameters and contains their nominal (initial) values. `'c'` specifies that the underlying parameterization is in continuous time. `aux` contains the values of the auxiliary parameters.

Note You must specify `T` and `aux` even if they are not used by the `myfunc` code.

Use `pem` to estimate the grey-box parameter values:

```
m = pem(data,m)
```

where `data` is the estimation data and `m` is the `idgrey` object with unknown parameters.

Note Compare this example to “Example – Estimating Structured Continuous-Time State-Space Models” on page 5-72, where the same problem is solved using a structured state-space representation.

Example – Estimating a Continuous-Time Grey-Box Model for Heat Diffusion

In this example, you estimate the heat conductivity and the heat-transfer coefficient of a continuous-time grey-box model for a heated-rod system.

This system consists of a well-insulated metal rod of length L and a heat-diffusion coefficient κ . The input to the system is the heating power $u(t)$ and the measured output $y(t)$ is the temperature at the other end.

Under ideal conditions, this system is described by the heat-diffusion equation—which is a partial differential equation in space and time.

$$\frac{\partial x(t, \xi)}{\partial t} = \kappa \frac{\partial^2 x(t, \xi)}{\partial \xi^2}$$

To get a continuous-time state-space model, you can represent the second-derivative using the following difference approximation:

$$\frac{\partial^2 x(t, \xi)}{\partial \xi^2} = \frac{x(t, \xi + \Delta L) - 2x(t, \xi) + x(t, \xi - \Delta L)}{(\Delta L)^2}$$

$$\text{where } \xi = k \cdot \Delta L$$

This transformation produces a state-space model of order $n = \frac{L}{\Delta L}$, where the state variables $x(t, k \cdot \Delta L)$ are lumped representations for $x(t, \xi)$ for the following range of values:

$$k \cdot \Delta L \leq \xi < (k + 1) \Delta L$$

The dimension of x depends on the spatial grid size ΔL in the approximation.

The heat-diffusion equation is mapped to the following continuous-time state-space model structure to identify the state-space matrices:

$$\dot{x}(t) = Fx(t) + Gu(t) + \tilde{K}w(t)$$

$$y(t) = Hx(t) + Du(t) + w(t)$$

$$x(0) = x_0$$

The following M-file describes the state-space equation for this model. In this case, the auxiliary variables specify grid-size variables, so that you can modify the grid size without the M-file.

```
command [A,B,C,D,K,x0] = heatd(pars,T,aux)
% Number of points in the space-discretization
Ngrid = aux(1);
% Length of the rod
L = aux(2);
% Initial rod temperature (uniform)
temp = aux(3);
% Space interval
deltaL = L/Ngrid;
% Heat-diffusion coefficient
kappa = pars(1);
% Heat transfer coefficient at far end of rod
htf = pars(2);
A = zeros(Ngrid,Ngrid);
for kk = 2:Ngrid-1
    A(kk,kk-1) = 1;
    A(kk,kk) = -2;
    A(kk,kk+1) = 1;
end
% Boundary condition on insulated end
A(1,1) = -1; A(1,2) = 1;
A(Ngrid,Ngrid-1) = 1;
A(Ngrid,Ngrid) = -1;
A = A*kappa/deltaL/deltaL;
B = zeros(Ngrid,1);
B(Ngrid,1) = htf/deltaL;
C = zeros(1,Ngrid);
C(1,1) = 1;
D = 0;
K = zeros(Ngrid,1);
x0 = temp*ones(Ngrid,1);
```

Use the following syntax to define an `idgrey` model object based on the `myfunc` M-file:

```
m = idgrey('heatd',[0.27 1],'c',[10,1,22])
```

This command specifies the auxiliary parameters as inputs to the command, include the model order 10, the rod length of 1 meter, and an initial temperature of 22 degrees Celsius. The command also specifies the initial values for heat conductivity as 0.27, and for the heat transfer coefficient as 1.

For given data, you can use `pem` to estimate the grey-box parameter values:

```
me = pem(data,m)
```

The following command shows how you can specify to estimate a new model with different auxiliary variables directly in the estimator command:

```
me = pem(data,m,'FileArgument',[20,1,22])
```

This syntax uses the `FileArgument` model property to specify a finer grid using a larger value for `Ngrid`. For more information about linear grey-box model properties, see the `idgrey` reference page.

Example – Estimating a Discrete-Time Grey-Box Model with Parameterized Disturbance

This example shows how to create a grey-box model structure when you know the variance of the measurement noise. The code in this example uses the Control System Toolbox command `d1qr` for computing the Kalman gain from the known and estimated noise variance.

Consider the following discrete-time state-space equation:

$$\begin{aligned}x(kT + T) &= \begin{bmatrix} \text{par1} & \text{par2} \\ 1 & 0 \end{bmatrix} x(kT) + \begin{bmatrix} 1 \\ 0 \end{bmatrix} u(kT) + w(kT) \\ y(kT) &= [\text{par3} \quad \text{par4}] x(kT) + e(kT) \\ x(0) &= x0\end{aligned}$$

where w and e are independent white noises with covariance matrices $R1$ and $R2$, respectively. par1 , par2 , par3 , and par4 represent the unknown parameter values to be estimated.

Suppose that you know the variance of the measurement noise $R2$, and that only the first component of $w(t)$ is nonzero. The following M-file shows how to capture this information in an M-file:

```
command [A,B,C,D,K,x0] = mynoise(par,T,aux)
R2 = aux(1); % Known measurement noise variance
A = [par(1) par(2);1 0];
B = [1;0];
C = [par(3) par(4)];
D = 0;
R1 = [par(5) 0;0 0];
K = A*d1qr(A,eye(2),C,R1,R2); % Uses Control System Toolbox
x0 = [0;0];
```

The Kalman gain is computed using the `d1qr` command in Control System Toolbox.

Estimating Nonlinear Grey-Box Models

In this section...

“Supported Nonlinear Grey-Box Models” on page 9-13

“Nonlinear Grey-Box Demos and Examples” on page 9-13

“Specifying the Nonlinear Grey-Box Model Structure” on page 9-13

“Constructing the `idnlgrey` Object” on page 9-15

“Using `pem` to Estimate Nonlinear Grey-Box Models” on page 9-16

“Options for the Estimation Algorithm” on page 9-16

Supported Nonlinear Grey-Box Models

You can estimate nonlinear discrete-time and continuous-time grey-box models for arbitrary nonlinear ordinary differential equations using single-output and multiple-output time-domain data, or time-series data that has no inputs. Your grey-box models can be static or dynamic.

Grey-box models describe the system behavior as a set of nonlinear differential or difference equations with unknown parameters.

Nonlinear Grey-Box Demos and Examples

System Identification Toolbox provides several demos and case studies on creating, manipulating and estimating nonlinear grey-box models. You can access these demos by typing the following command at the MATLAB prompt:

```
iddemo
```

For examples of m-files and MEX-files that specify model structure, see the `toolbox/ident/iddemos/examples` directory. For example, the model of a DC motor—used in the demo `idnlgreydemo1`—is described in files `dcmotor_m` and `dcmotor_c`.

Specifying the Nonlinear Grey-Box Model Structure

You must represent your system as a set of first-order nonlinear difference or differential equations:

$$\begin{aligned}x^{\dagger}(t) &= F(t, x(t), u(t), \text{par1}, \text{par2}, \dots, \text{parN}) \\y(t) &= H(t, x(t), u(t), \text{par1}, \text{par2}, \dots, \text{parN}) + e(t) \\x(0) &= x_0\end{aligned}$$

where $x^{\dagger}(t) = \frac{d}{dt}x(t)$ for continuous-time representation and $x^{\dagger}(t) = x(t + T_s)$ for discrete-time representation with T_s as the sampling interval. F and H are arbitrary linear or nonlinear functions with N_x and N_y components, respectively. N_x is the number of states and N_y is the number of outputs.

After you establish the equations for your system, create an M-file or MEX-file. MEX-files, which can be created in C or Fortran, are dynamically-linked subroutines that can be loaded and executed by the MATLAB interpreter. For more information about MEX-files, see the MATLAB documentation.

The purpose of the model file is to return the state derivatives and model outputs as a command of time, states, inputs, and model parameters, as follows:

```
[dx,y] = MODFILENAME(t,x,u,p1,p2, ...,pN,FileArgument)
```

Tip The template file for writing the C MEX-file, `IDNLGREY_MODEL_TEMPLATE.c`, is located in `matlab/toolbox/ident/nlident`.

The output variables are:

- dx — Represents the right side(s) of the state-space equation(s). A column vector with N_x entries. For static models, $dx = []$.
For discrete-time models. dx is the value of the states at the next time step $x(t+T_s)$.
For continuous-time models. dx is the state derivatives at time t , or $\frac{dx}{dt}$.
- y — Represents the right side(s) of the output equation(s). A column vector with N_y entries.

The file inputs are:

- t — Current time.
- x — State vector at time t . For static models, equals $[]$.
- u — Input vector at time t . For time-series models, equals $[]$.
- p_1, p_2, \dots, p_N — Parameters, which can be real scalars, column vectors or two-dimensional matrices. N is the number of parameter object. For scalar parameters, N is the total number of parameter elements.
- `FileArgument` — Contains auxiliary variables that might be required for updating the constants in the state equations.

Tip After creating a model file, call it directly from MATLAB with reasonable inputs and verify the output values.

For an example of creating grey-box model files, see the demo `Creating idnlgrey Model Files`.

Constructing the `idnlgrey` Object

After you create the M-file or MEX-file with your model structure, you must define an `idnlgrey` object. This object shares many of the properties of the linear `idgrey` model object.

Use the following syntax to define the `idnlgrey` model object:

```
m = idnlgrey('filename',Order,Parameters,InitialStates)
```

The `idnlgrey` arguments are defined as follows:

- `'filename'` — Name of the M-file or MEX-file storing the model structure. This file must be on the MATLAB path.
- `Order` — Vector with three entries $[N_y \ N_u \ N_x]$, specifying the number of model outputs N_y , the number of inputs N_u , and the number of states N_x .
- `Parameters` — Parameters, specified as struct arrays, cell arrays, or double arrays.
- `InitialStates` — Specified in the same way as parameters. Must be the fourth input to the `idnlgrey` constructor.

For detailed information about this object and its properties, see the `idnlgrey` reference page.

Using `pem` to Estimate Nonlinear Grey-Box Models

You can use the `pem` command to estimate the unknown `idnlgrey` model parameters and initial states using measured data.

The input-output dimensions of the data must be compatible with the input and output orders you specified for the `idnlgrey` model.

Use the following general estimation syntax:

```
m = pem(data,m)
```

where `data` is the estimation data and `m` is the `idnlgrey` model object you constructed.

You can pass additional property-value pairs to `pem` to specify the properties of the model or the estimation algorithm. Assignable properties include the ones returned by the `get(idnlgrey)` command and the algorithm properties returned by the `get(idnlgrey, 'Algorithm')`, such as `MaxIter` and `Tolerance`. For detailed information about these model properties, see the `idnlgrey` reference page.

For more information about validating your models, see Chapter 10, “Validating and Analyzing Models”.

Options for the Estimation Algorithm

The `Algorithm` property of the model specifies the estimation algorithm, which simulates the model several times by trying various parameter values to reduce the prediction error.

The following algorithm properties can affect the quality of the results:

- “Simulation Method” on page 9-17
- “Search Method” on page 9-17
- “Gradient Options” on page 9-18

- “Example – Specifying Algorithm Properties” on page 9-18

For detailed information about these and other model properties, see the `idnlgrey` reference page.

Simulation Method

You can specify the simulation method using the `SimulationOptions` (struct) fields of the model `Algorithm` property.

MATLAB provides several variable-step and fixed-step solvers for simulating `idnlgrey` models. To view a list of available solvers and their properties, type the following command at the MATLAB prompt:

```
idprops idnlgrey algorithm.simulationoptions
```

For discrete-time systems, the default solver is `'FixedStepDiscrete'`. For continuous-time systems, the default solver is `'ode45'`.

By default, `SimulationOptions.Solver` is set to `'Auto'`, which automatically selects either `'ode45'` or `'FixedStepDiscrete'` during estimation and simulation—depending on whether the system is continuous or discrete in time.

Search Method

You can specify the search method for generating maximum likelihood estimates of model parameters using the `SearchMethod` field of the `Algorithm` property. Two categories of methods are available for nonlinear grey-box modeling.

One category of methods consists of the minimization schemes that are based on line-search methods, including Gauss-Newton type methods, steepest-descent methods, and Levenberg-Marquardt methods.

(Requires Optimization Toolbox) The Trust-Region Reflective Newton method of nonlinear least-squares (`'lsqnonlin'`), where the cost is the sum of squares of errors between the measured and simulated outputs. When the parameter bounds differ from the default $\pm \text{Inf}$, this search method handles the bounds better than the schemes based on a line search, but the results might not be optimal in the maximum-likelihood sense.

By default, `SearchMethod` is set to `'Auto'`, automatically selects a method from the available minimizers. If Optimization Toolbox is available, `SearchMethod` is set to `'lsqnonlin'`. Otherwise, `SearchMethod` is a combination of line-search based schemes.

Gradient Options

You can specify the method for calculating gradients using the `GradientOptions` field of the `Algorithm` property. *Gradients* are the derivatives of errors with respect to unknown parameters and initial states.

Gradients are calculated by numerically perturbing unknown quantities and measuring their effects on the simulation error.

Options for gradient computation include the choice of the differencing scheme (forward, backward or central), the size of minimum perturbation of the unknown quantities, and whether the gradients are calculated simultaneously or individually.

Example – Specifying Algorithm Properties

You can specify the `Algorithm` fields directly in the estimation syntax, as property-value pairs.

For example, if you want to use `SearchMethod = 'gn'`, `MaxIter = 5`, and `Trace = 'on'`, use the following syntax in the `pem` command:

```
m = pem(data,init_model,'Search','gn',...
        'MaxIter',5,...
        'Trace','On')
```

After Estimating Grey-Box Models

After estimating linear and nonlinear grey-box models, you can simulate the model output using the `sim` command. For more information, see Chapter 12, “Simulating and Predicting Model Output”.

System Identification Toolbox represents linear grey-box models as `idgrey` model objects. You can convert these models to state-space form using the `idss` command and analyze the model behavior using transient- and frequency-response plots and other linear analysis plots, as described in Chapter 11, “Transforming Model Representation”.

System Identification Toolbox represents nonlinear grey-box models as `idnlgrey` model objects. These model objects store the parameter values resulting from the estimation. You can access these parameters from the model objects to use these variables in computation in the MATLAB Workspace browser.

Note Linearization of nonlinear grey-box models is not supported.

Validating and Analyzing Models

Overview of Model Validation and Plots (p. 10-3)

Introduction to validating models and supported model plots.

Using Model Output Plots to Validate Models (p. 10-7)

Plotting simulated or predicted model output and comparing model output to measured output for all linear parametric and nonlinear models.

Using Residual Analysis Plots to Validate Models (p. 10-15)

Plotting residuals and performing residual analysis tests for all linear parametric and nonlinear models.

Using Impulse- and Step-Response Plots to Validate Models (p. 10-21)

Plotting transient response plots for models, including impulse response and step response, for all linear parametric models and correlation analysis models.

Using Frequency-Response Plots to Validate Models (p. 10-29)

Plotting Bode and Nyquist plots for models.

Creating Noise-Spectrum Plots (p. 10-37)

Plotting the frequency-response of the estimated noise model for a linear system.

Using Pole-Zero Plots to Validate Models (p. 10-44)

Plotting pole-zero plots for linear parametric models and using pole-zero plots to gain insight into model-order reduction.

Using Nonlinear ARX Plots to Validate Models (p. 10-49)	Plotting nonlinearity characteristics of a nonlinear ARX model.
Using Hammerstein-Wiener Plots to Validate Models (p. 10-53)	Plotting characteristics of linear and nonlinear blocks in a Hammerstein-Wiener model.
Using Akaike's Criteria to Validate Models (p. 10-58)	Validating models using Akaike's Final Prediction Error (FPE) and Akaike's Information Criterion (AIC).
Computing Model Uncertainty (p. 10-61)	Computing model parameter uncertainty of linear models.
Extracting Parameter Values from Linear Models (p. 10-64)	Extracting numerical parameter values and uncertainties from linear models.
Extracting Dynamic Model and Noise Model Separately (p. 10-66)	Extracting the dynamic model and noise model separately from a model that includes both dynamics and noise.
Troubleshooting Models (p. 10-68)	Adjusting your modeling strategy based on model-validation plots.
Next Steps After Getting an Accurate Model (p. 10-73)	How you can work with models identified using System Identification Toolbox.

Overview of Model Validation and Plots

In this section...
“When to Validate Models” on page 10-3
“Ways to Validate Models” on page 10-3
“Data for Validating Models” on page 10-4
“Supported Model Plots” on page 10-5
“Getting Advice About Models” on page 10-6

When to Validate Models

After estimating each model, you can validate whether the model reproduces system behavior within acceptable bounds. You iterate between estimation and validation until you find the simplest model that best captures the system dynamics.

For ideas on how to adjust your modeling strategy based on validation results, see “Troubleshooting Models” on page 10-68.

Tip If you have installed Control System Toolbox, you can also view models using the LTI Viewer. For more information, see “Viewing Model Response Using the LTI Viewer” on page 13-5.

Ways to Validate Models

The following validation techniques are available in System Identification Toolbox:

- Comparing simulated or predicted model output to measured output. See “Using Model Output Plots to Validate Models” on page 10-7.
- Analyzing autocorrelation and cross-correlation of the residuals with input. See “Using Residual Analysis Plots to Validate Models” on page 10-15.
- Analyzing model response. For more information, see the following:

- “Using Impulse- and Step-Response Plots to Validate Models” on page 10-21
- “Using Frequency-Response Plots to Validate Models” on page 10-29

For information about the response of the noise model, see “Creating Noise-Spectrum Plots” on page 10-37.

- Plotting the poles and zeros of the linear parametric model. For more information, see “Using Pole-Zero Plots to Validate Models” on page 10-44.
- Comparing the response of nonparametric models, such as impulse-, step-, and frequency-response models, to parametric models, such as linear polynomial models, state-space model, and nonlinear parametric models.

Note Do not use this comparison when feedback is present in the system because feedback makes nonparametric models unreliable. To test if feedback is present in the system, use the `advce` command on the data.

- Compare models using Akaike Information Criterion or Akaike Final Prediction Error. For more information, see the `aic` and `fpe` reference page.
- Plotting linear and nonlinear blocks of Hammerstein-Wiener and nonlinear ARX models. For more information, see “Using Hammerstein-Wiener Plots to Validate Models” on page 10-53 and “Using Nonlinear ARX Plots to Validate Models” on page 10-49.

Displaying confidence intervals on supported plots helps you assess the uncertainty of model parameters. For more information, see “Computing Model Uncertainty” on page 10-61.

Data for Validating Models

For plots that compare model response to measured response, such as model output and residual analysis plots, you designate two types of data sets: one for estimating the models (*estimation data*), and the other for validating the models (*validation data*). When you validate a model using a fresh data set, this process is called *cross-validation*. Although you can designate the same data set to be used for estimating and validating the model, you risk overfitting your data.

Note Validation data should be the same in frequency content as the estimation data. If you detrended the estimation data, you must remove the same trend from the validation data. For more information about detrending, see “Detrending Data” on page 2-75.

Supported Model Plots

System Identification Toolbox provides several plots for validating model dynamics. The following table summarizes the types of supported model plots.

Plot Type	Supported Models	Learn More
Model Output	All linear and nonlinear models	“Using Model Output Plots to Validate Models” on page 10-7
Residual Analysis	All linear and nonlinear models	“Using Residual Analysis Plots to Validate Models” on page 10-15
Transient Response	<ul style="list-style-type: none"> • All linear parametric models • Correlation analysis (nonparametric) models • For nonlinear models, only step response. 	“Using Impulse- and Step-Response Plots to Validate Models” on page 10-21
Frequency Response	<ul style="list-style-type: none"> • All linear parametric models • Spectral analysis (nonparametric) models 	“Using Frequency-Response Plots to Validate Models” on page 10-29

Plot Type	Supported Models	Learn More
Noise Spectrum	<ul style="list-style-type: none"> • All linear parametric models • Spectral analysis (nonparametric) models 	“Creating Noise-Spectrum Plots” on page 10-37
Poles and Zeros	All linear parametric models	“Using Pole-Zero Plots to Validate Models” on page 10-44
Nonlinear ARX	Nonlinear ARX models only	“Using Nonlinear ARX Plots to Validate Models” on page 10-49
Hammerstein-Wiener	Hammerstein-Wiener models only	“Using Hammerstein-Wiener Plots to Validate Models” on page 10-53

For general information about working with plots in System Identification Toolbox, see “Plots Data and Models in the GUI” on page 1-18.

Getting Advice About Models

Use the `advise` command on an estimated model to answer the following questions about the model:

- Should I increase or decrease the model order?
- Should I estimate a noise model?
- Is feedback present?

Using Model Output Plots to Validate Models

In this section...

“Supported Model Types” on page 10-7

“What Does a Model Output Plot Show?” on page 10-7

“Choosing Simulated or Predicted Output” on page 10-8

“How to Plot Model Output Using the GUI” on page 10-10

“Displaying the Confidence Interval” on page 10-12

“How to Plot Model Output at the Command Line” on page 10-13

Supported Model Types

System Identification Toolbox lets you validate all linear parametric models and nonlinear models by checking how well the simulated or predicted output of the model matches the measured output.

Note For nonparametric models, including impulse-response, step-response, and frequency-response models, model output plots are not available. For time-series models, you can only generate model-output plots for parametric models using time-domain time-series (no input) measured data.

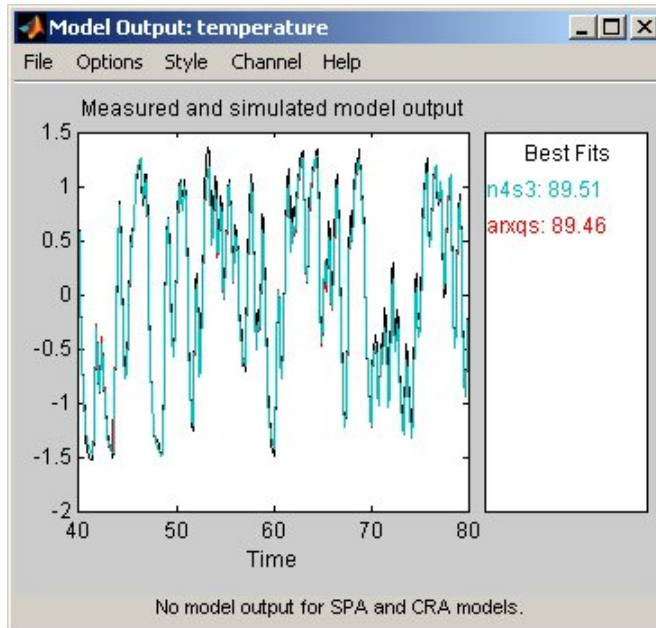
What Does a Model Output Plot Show?

The model output plot shows different information depending on the domain of the input-output validation data, as follows:

- For time-domain validation data, the plot shows simulated or predicted model output.
- For frequency-domain data, the plot shows the simulated complex-valued amplitude of the model output. The complex-valued amplitude is equal to the product of the Fourier transform of the input and the model frequency command.
- For frequency-response data, the plot shows the simulated amplitude of the model frequency response.

For linear models, you can estimate a model using time-domain data, and then validate the model using frequency domain data. For nonlinear models, you can only use time-domain data for both estimation and validation.

The following figure shows a sample Model Output plot, created in the System Identification Tool GUI.



Choosing Simulated or Predicted Output

How you validate the model output should match how you plan to use the model. If you plan to use the model for simulation applications, validate the model by comparing simulated output to the validation data. However, if you plan to use the model for prediction, compare the k -step-ahead predicted output to the validation data. For example, if you are modeling a plant for a control system, your model must perform for prediction over a horizon that corresponds to the time-constant of the system.

The main difference between simulation and prediction is whether System Identification Toolbox uses measured or computed previous outputs for computing the next output.

Simulating models uses the input-data values from a data set to compute the output values. When you *simulate* the model output, System Identification Toolbox computes the first output value using the initial conditions and the inputs. Then, System Identification Toolbox feeds this computed output into the differential (continuous-time) or difference (discrete-time) equation for calculating the next output value. In this way, the simulation progresses using previously calculated outputs in the difference equation to produce the next output; with an *infinite prediction horizon* ($k=\infty$), the simulation has no limit on how far out in time it computes output values. Simulating models uses the input-data values from the validation data set to compute the output values.

Using a model for prediction is common in controls applications where you want to predict output a certain number of steps in advance. When you use System Identification Toolbox to *predict* model output, the algorithm uses both the measured and the calculated output data values in the difference equation for computing the next output.

The predicted value $y(t)$ is computed from all available inputs $u(s)$, where $s \leq (t - k)$, and all available outputs $y(s)$, where $s \leq (t - k)$.

To check whether the model has picked up interesting dynamic properties, let the predicted time horizon kT be larger than the important time constants, where T is the sampling interval.

Note Output-error models, obtained by fixing K to zero for state-space models and setting $na=nc=nd=0$ for polynomial models, do not use past outputs. Therefore, for these models, the simulated and the predicted outputs are the same for any value of k .

To learn how to display simulated or predicted output, see the description of the plot settings in “How to Plot Model Output Using the GUI” on page 10-10.

How to Plot Model Output Using the GUI

To create a model output plot for parametric linear and nonlinear models in the System Identification Tool GUI, select the **Model output** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

The right side of the plot displays the percentage of the output that the model reproduces (**Best Fit**), computed using the following equation:

$$\text{Best Fit} = \left(1 - \frac{|y - \hat{y}|}{|y - \bar{y}|} \right) \times 100$$

In this equation, y is the measured output, \hat{y} is the simulated or predicted model output, and \bar{y} is the mean of y . 100% corresponds to a perfect fit, and 0% indicates that the fit is no better than guessing the output to be a constant ($\hat{y} = \bar{y}$).

Because of the definition of **Best Fit**, it is possible for this value to be negative. A negative best fit is worse than 0% and can occur for the following reasons:

- The estimation algorithm failed to converge.
- The model was not estimated by minimizing $|y - \hat{y}|$. **Best Fit** can be negative when you minimized 1-step-ahead prediction during the estimation, but validate using the simulated output \hat{y} .
- The validation data set was not preprocessed in the same way as the estimation data set.

The following table summarizes the Model Output plot settings.

Model Output Plot Settings

Action	Command
<p>Display confidence intervals.</p> <hr/> <p>Note Confidence intervals are only available for simulated model output of linear models. Confidence intervals are not available for nonlinear ARX and Hammerstein-Wiener models.</p>	<ul style="list-style-type: none"> • To display the dashed lines on either side of the nominal model curve, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. • To change the confidence value, select Options > Set % confidence level, and choose a value from the list. • To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.
<p>Change between simulated output or predicted output.</p> <hr/> <p>Note Prediction is only available for time-domain validation data.</p>	<ul style="list-style-type: none"> • Select Options > Simulated output or Options > k step ahead predicted output. • To change the prediction horizon, select Options > Set prediction horizon, and select the number of samples. • To enter your own prediction horizon, select Options > Set prediction horizon > Other. Enter the value in terms of the number of samples.
<p>Display the actual output values (Signal plot), or the difference between model output and measured output (Error plot).</p>	<p>Select Options > Signal plot or Options > Error plot.</p>

Model Output Plot Settings (Continued)

Action	Command
(Time-domain validation data only) Set the time range for model output and the time interval for which the Best Fit value is computed.	Select Options > Customized time span for fit and enter the minimum and maximum time values. For example: [1 20]
(Multiple-output system only) Select a different output.	Select the output by name in the Channel menu.

Displaying the Confidence Interval

In the GUI, you can display a confidence interval on the plot to gain insight into the quality of a linear model. To learn how to show or hide confidence interval, see the description of the plot settings in “How to Plot Model Output Using the GUI” on page 10-10.

The *confidence interval* corresponds to the range of output values with a specific probability of being the actual output of the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around the nominal curve represents the range of values that have a 95% probability of being the true system response. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

Note The calculation of the confidence interval assumes that the model sufficiently describes the system dynamics and the model residuals pass independence tests.

How to Plot Model Output at the Command Line

You can plot simulated and predicted model output using the `compare`, `sim`, and `predict` commands.

Simulation and prediction requires input data, a model, and the values of the initial states. If you estimated the model using one data set, but want to simulate the model using a different data set, the initial states of your simulation must be consistent with the latter data set.

By default, `sim` and `predict` use the initial states that were derived from the data you used to estimate the model. These initial states are not appropriate if you are simulating or predicting output using new data.

To use `sim` or `predict` with a data set that differs from the data you used to estimate the model, first estimate the new initial states `X0est` using `pe`:

```
[E,X0est]=pe(model,data)
```

Next, specify the estimated initial states `X0est` as an argument in `sim` or `predict`. For example:

```
y=sim(model,data,'InitialState',X0est)
```

Note The `compare` command automatically estimates the initial states from the data and ensures consistency.

Command	Description	Example
compare	Plots simulated or predicted model output on top of the measured output. You should use an independent validation data set as input to the model.	<p>To plot five-step-ahead predicted output of the model <code>mod</code> against the validation data <code>data</code>, use the following command:</p> <pre>compare(data,mod,5)</pre> <hr/> <p>Note Omitting the third argument assumes an infinite horizon and results in simulation.</p>
sim	Plots simulated model output only.	<p>To simulate the response of the model <code>model</code> using input data <code>data</code>, use the following command:</p> <pre>sim(model,data)</pre>
predict	Plots predicted model output only.	<p>To perform one-step-ahead prediction of the response for the model <code>model</code> and input data <code>data</code>, use the following command:</p> <pre>predict(model,data,1)</pre>

Using Residual Analysis Plots to Validate Models

In this section...

“What Is Residual Analysis?” on page 10-15

“Supported Model Types” on page 10-16

“What Does the Residuals Plot Show?” on page 10-16

“Displaying the Confidence Interval” on page 10-17

“How to Plot Residuals Using the GUI” on page 10-18

“How to Plot Residuals at the Command Line” on page 10-20

What Is Residual Analysis?

Residuals are differences between the one-step-predicted output from the model and the measured output from the validation data set. Thus, residuals represent the portion of the validation data not explained by the model.

Residual analysis consists of two tests: the whiteness test and the independence test.

According to the *whiteness test* criteria, a good model has the residual autocorrelation command inside the model confidence interval, indicating that the residuals are uncorrelated.

According to the *independence test* criteria, a good model has residuals uncorrelated with past inputs. Evidence of correlation indicates that the model does not describe how part of the output relates to the corresponding input. For example, a peak outside the confidence interval for lag k means that the output $y(t)$ that originates from the input $u(t-k)$ is not properly described by the model.

Your model should pass both the whiteness and the independence tests, except in the following cases:

- For Output-Error models and when using instrumental-variable (IV) methods, make sure that your model shows independence of e and u , and pay less attention to the results of the whiteness of e .

In this case, the modeling focus is on the dynamics G and not the disturbance properties H .

- Correlation between residuals and input for negative lags, is not necessarily an indication of an inaccurate model.

When current residuals at time t affect future input values, there might be feedback in your system. In the case of feedback, concentrate on the positive lags in the cross-correlation plot during model validation.

Supported Model Types

You can validate parametric linear and nonlinear models by checking the behavior of the model residuals. For a description of residual analysis, see “What Does the Residuals Plot Show?” on page 10-16.

Note For nonparametric models, including impulse-response, step-response, and frequency-response models, residual analysis plots are not available. For time-series models, you can only generate model-output plots for parametric models using time-domain time-series (no input) measured data.

What Does the Residuals Plot Show?

Residual analysis plots show different information depending on whether you use time-domain or frequency-domain input-output validation data.

For time-domain validation data, the plot shows the following two axes:

- Autocorrelation command of the residuals for each output
- Cross-correlation between the input and the residuals for each input-output pair

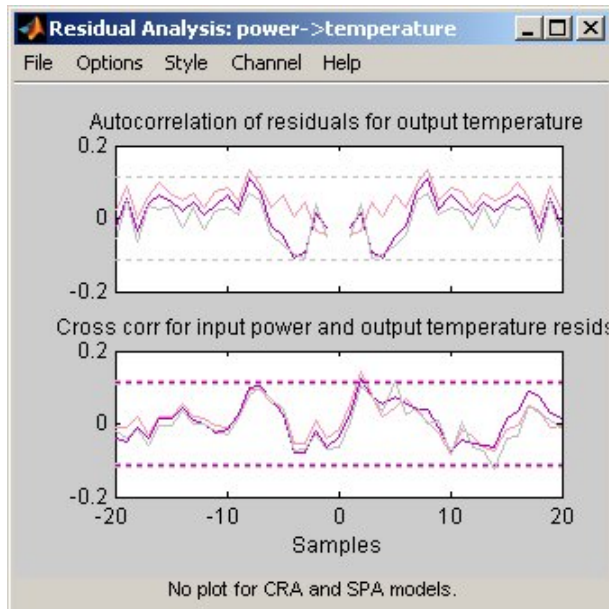
Note For time-series models, the residual analysis plot does not provide any input-residual correlation plots.

For frequency-domain validation data, the plot shows the following two axes:

- Estimated power spectrum of the residuals for each output
- Transfer-command amplitude from the input to the residuals for each input-output pair

For linear models, you can estimate a model using time-domain data, and then validate the model using frequency domain data. For nonlinear models, System Identification Toolbox supports only time-domain data.

The following figure shows a sample Residual Analysis plot, created in the System Identification Tool GUI.



Displaying the Confidence Interval

You can display a confidence interval on the plot in the GUI to gain insight into the quality of the model. To learn how to show or hide confidence interval, see the description of the plot settings in “How to Plot Residuals Using the GUI” on page 10-18.

Note If you are working in the System Identification Tool GUI, you can specify a custom confidence interval. If you are using the `resid` command, the confidence interface is fixed at 99%.

The *confidence interval* corresponds to the range of residual values with a specific probability of being statistically insignificant for the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around zero represents the range of residual values that have a 95% probability of being statistically insignificant. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

How to Plot Residuals Using the GUI

To create a residual analysis plot for parametric linear and nonlinear models in the System Identification Tool GUI, select the **Model resids** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

The following table summarizes the Residual Analysis plot settings.

Residual Analysis Plot Settings

Action	Command
<p>Display confidence intervals around zero.</p> <hr/> <p>Note Confidence intervals are not available for nonlinear ARX and Hammerstein-Wiener models.</p> <hr/>	<ul style="list-style-type: none"> • To display the dashed lines on either side of the nominal model curve, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. • To change the confidence value, select Options > Set % confidence level and choose a value from the list. • To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.
<p>Change the number of lags (data samples) for which to compute autocorrelation and cross-correlation functions.</p> <hr/> <p>Note For frequency-domain validation data, increasing the number of lags increases the frequency resolution of the residual spectrum and the transfer command.</p> <hr/>	<ul style="list-style-type: none"> • Select Options > Number of lags and choose the value from the list. • To enter your own lag value, select Options > Set confidence level > Other. Enter the value as the number of data samples.
<p>(Multiple-output system only) Select a different input-output pair.</p>	<p>Select the input-output by name in the Channel menu.</p>

How to Plot Residuals at the Command Line

The following table summarizes commands that generate residual-analysis plots for linear and nonlinear models. For detailed information about this command, see the corresponding reference page.

Note Apply `pe` and `resid` to one model at a time.

Command	Description	Example
<code>pe</code>	Computes and plots model prediction errors.	To plot the prediction errors for the model <code>model</code> using data <code>data</code> , type the following command: <code>pe(model,data)</code>
<code>resid</code>	Performs whiteness and independence tests on model residuals, or prediction errors. Uses validation data input as model input.	To plot residual correlations for the model <code>model</code> using data <code>data</code> , type the following command: <code>resid(model,data)</code>

Using Impulse- and Step-Response Plots to Validate Models

In this section...

“Supported Models” on page 10-21

“How Transient Response Helps to Validate Models” on page 10-21

“What Does a Transient Response Plot Show?” on page 10-22

“How to Plot Impulse and Step Response Using the GUI” on page 10-23

“Displaying the Confidence Interval” on page 10-26

“How to Plot Impulse and Step Response at the Command Line” on page 10-27

Supported Models

You can plot the simulated response of a model using impulse and step signals as the input for all linear parametric models and correlation analysis (nonparametric) models.

You can also create step-response plots for nonlinear models. These step and impulse response plots, also called *transient response* plots, provide insight into the characteristics of model dynamics, including peak response and settling time.

Note For frequency-response models, impulse- and step-response plots are not available. For nonlinear models, only step-response plots are available.

How Transient Response Helps to Validate Models

Transient response plots provide insight into the basic dynamic properties of the model, such as response times, static gains, and delays.

Transient response plots also help you validate how well a linear parametric model, such as a linear ARX model or a state-space model, captures the dynamics. For example, you can estimate an impulse or step response from the data using correlation analysis (nonparametric model), and then plot the

correlation analysis result on top of the transient responses of the parametric models.

Because nonparametric and parametric models are derived using different algorithms, agreement between these models increases confidence in the parametric model results.

What Does a Transient Response Plot Show?

Transient response plots show the value of the impulse or step response on the vertical axis. The horizontal axis is in units of time you specified for the data used to estimate the model.

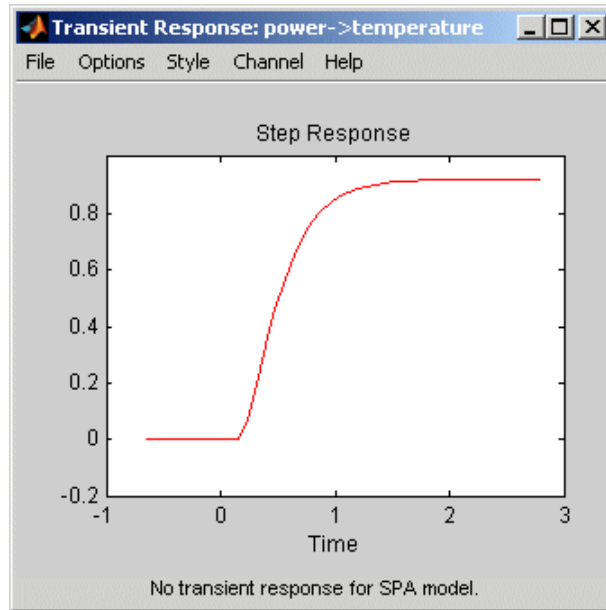
The impulse response of a dynamic model is the output signal that results when the input is an impulse. That is, $u(t)$ is zero for all values of t except at $t=0$, where $u(0)=1$. In the following difference equation, you can compute the impulse response by setting $y(-T)=y(-2T)=0$, $u(0)=1$, and $u(t>0)=0$.

$$y(t) - 1.5y(t - T) + 0.7y(t - 2T) = 0.9u(t) + 0.5u(t - T)$$

The step response is the output signal that results from a step input, where $u(t<0)=0$ and $u(t>0)=1$.

If your model includes a noise model, you can display the transient response of the noise model associated with each output channel. For more information about how to display the transient response of the noise model, see “How to Plot Impulse and Step Response Using the GUI” on page 10-23.

The following figure shows a sample Transient Response plot, created in the System Identification Tool GUI.



How to Plot Impulse and Step Response Using the GUI

To create a transient analysis plot in the System Identification Tool GUI, select the **Transient resp** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

The following table summarizes the Transient Response plot settings.

Transient Response Plot Settings

Action	Command
Display step response for linear or nonlinear model.	Select Options > Step response .
Display impulse response for linear model.	Select Options > Impulse response . <hr/> Note Not available for nonlinear models.
Display the confidence interval. <hr/> Note Only available for linear models.	<ul style="list-style-type: none"> • To display the dashed lines on either side of the nominal model curve, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. • To change the confidence value, select Options > Set % confidence level, and choose a value from the list. • To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.

Transient Response Plot Settings (Continued)

Action	Command
<p>Change time span over which the impulse or step response is calculated. For a scalar time span T, the resulting response is plotted from $-T/4$ to T.</p> <hr/> <p>Note To change the time span of models you estimated using correlation analysis models, select Estimate > Correlation models and reestimate the model using a new time span.</p>	<ul style="list-style-type: none"> • Select Options > Time span (time units), and choose a new time span in units of time you specified for the model. • To enter your own time span, select Options > Time span (time units) > Other, and enter the total response duration. • To use the time span based on model dynamics, type [] or default. <p>The default time span is computed based on the model dynamics and might be different for different models. For nonlinear models, the default time span is 10.</p>
<p>Toggle between line plot or stem plot.</p> <hr/> <p>Tip Use a stem plot for displaying impulse response.</p>	<p>Select Style > Line plot or Style > Stem plot.</p>

Transient Response Plot Settings (Continued)

Action	Command
<p>(Multiple-output system only) Select an input-output pair to view the noise spectrum corresponding to those channels.</p>	<p>Select the output by name in the Channel menu.</p> <p>If the plotted models include a noise model, you can display the transient response properties associated with each output channel. The name of the channel has the format e@OutputName, where OutputName is the name of the output channel corresponding to the noise model.</p>
<p>(Step response for nonlinear models only) Set level of the input step.</p> <hr/> <p>Note For multiple-input models, the input-step level applies only to the input channel you selected to display in the plot.</p> <hr/>	<p>Select Options > Step Size, and then chose from two options:</p> <ul style="list-style-type: none"> • 0->1 sets the lower level to 0 and the upper level to 1. • Other opens the Step Level dialog box, where you enter the values for the lower and upper level values.

Displaying the Confidence Interval

In addition to the transient-response curve, you can display a confidence interval on the plot. To learn how to show or hide confidence interval, see the description of the plot settings in “How to Plot Impulse and Step Response Using the GUI” on page 10-23.

The *confidence interval* corresponds to the range of response values with a specific probability of being the actual response of the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around the nominal curve represents the range of values that have a 95% probability of being the true system response. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

Note The calculation of the confidence interval assumes that the model sufficiently describes the system dynamics and the model residuals pass independence tests.

How to Plot Impulse and Step Response at the Command Line

You can plot impulse- and step-response plots using the `impulse` and `step` commands, respectively.

All plot commands have the same basic syntax, as follows:

- To plot one model, use the syntax `command(model)`.
- To plot several models, use the syntax `command(model1,model2,...,modelN)`.

In this case, `command` represents any of the plotting commands.

To display confidence intervals for a specified number of standard deviations, use the following syntax:

```
command(model, 'sd', sd)
```

where `sd` is the number of standard deviations of a Gaussian distribution. For example, a confidence value of 99% for the nominal model curve corresponds to 2.58 standard deviations.

To display a filled confidence region, use the following syntax:

```
command(model, 'sd', sd, 'fill')
```

The following table summarizes commands that generate impulse- and step-response plots. For detailed information about each command, see the corresponding reference page.

Command	Description	Example
impulse	<p>Plots impulse response for idpoly, idproc, idarx, idss, and idgrey model objects. Estimates and plots impulse response models for iddata objects.</p> <hr/> <p>Note Does not support nonlinear models.</p> <hr/>	<p>To plot the impulse response of the model mod, type the following command:</p> <pre>impulse(mod)</pre>
step	<p>Plots the step response of all linear and nonlinear models. Estimates and plots step response models for iddata objects.</p>	<p>To plot the step response of the model mod, type the following command:</p> <pre>step(mod)</pre> <p>To specify step levels for a nonlinear model, type the following command:</p> <pre>step(mod, 'InputLevel', [u1;u2])</pre>

Using Frequency-Response Plots to Validate Models

In this section...

“What Is Frequency Response?” on page 10-29

“How Frequency Response Helps to Validate Models” on page 10-30

“What Does a Frequency-Response Plot Show?” on page 10-31

“How to Plot Bode Plots Using the GUI” on page 10-32

“How to Plot Bode and Nyquist Plots at the Command Line” on page 10-35

What Is Frequency Response?

Frequency response plots show the complex values of a transfer command as a command of frequency.

In the case of linear dynamic systems, the transfer command G is essentially an operator that takes the input u of a linear system to the output y :

$$y = Gu$$

For a continuous-time system, the transfer command relates the Laplace transforms of the input $U(s)$ and output $Y(s)$:

$$Y(s) = G(s)U(s)$$

In this case, the frequency command $G(i\omega)$ is the transfer command evaluated on the imaginary axis $s=i\omega$.

For a discrete-time system sampled with a time interval T , the transfer command relates the Z-transforms of the input $U(z)$ and output $Y(z)$:

$$Y(z) = G(z)U(z)$$

In this case, the frequency command $G(e^{i\omega T})$ is the transfer command $G(z)$ evaluated on the unit circle. The argument of frequency command $G(e^{i\omega T})$ is

scaled by the sampling interval T to make the frequency command periodic with the sampling frequency $2\pi/T$.

How Frequency Response Helps to Validate Models

You can plot the frequency response of a model to gain insight into the characteristics of linear model dynamics, including the frequency of the peak response and stability margins. Frequency-response plots are available for all linear parametric models and spectral analysis (nonparametric) models.

Note Frequency-response plots are not available for nonlinear models. In addition, Nyquist plots do not support time-series models that have no input.

The frequency response of a linear dynamic model describes how the model reacts to sinusoidal inputs. If the input $u(t)$ is a sinusoid of a certain frequency, then the output $y(t)$ is also a sinusoid of the same frequency. However, the magnitude of the response is different from the magnitude of the input signal, and the phase of the response is shifted relative to the input signal.

Frequency response plots provide insight into linear systems dynamics, such as frequency-dependent gains, resonances, and phase shifts. Frequency response plots also contain information about controller requirements and achievable bandwidths. Finally, frequency response plots can also help you validate how well a linear parametric model, such as a linear ARX model or a state-space model, captures the dynamics.


One example of how frequency-response plots help validate other models is that you can estimate a frequency response from the data using spectral analysis (nonparametric model), and then plot the spectral analysis result on top of the frequency response of the parametric models. Because nonparametric and parametric models are derived using different algorithms, agreement between these models increases confidence in the parametric model results.

What Does a Frequency-Response Plot Show?

System Identification Tool GUI supports the following types of frequency-response plots for linear parametric models, linear state-space models, and nonparametric frequency-response models:

- Bode plot of the model response. A Bode plot consists of two plots. The top plot shows the magnitude $|G|$ by which the transfer command G magnifies the amplitude of the sinusoidal input. The bottom plot shows the phase $\phi = \arg G$ by which the transfer command shifts the input. The input to the system is a sinusoid, and the output is also a sinusoid with the same frequency.
- Bode plot of the disturbance model, called *noise spectrum*. This plot is the same as a Bode plot of the model response, but it shows the frequency response of the noise model instead. For more information, see “Creating Noise-Spectrum Plots” on page 10-37.
- (Only in MATLAB Command Window)
Nyquist plot. Plots the imaginary versus the real part of the transfer command.

The following figure shows a sample Bode plot of the model dynamics, created in the System Identification Tool GUI.



How to Plot Bode Plots Using the GUI

To create a frequency-response plot for parametric linear models in the System Identification Tool GUI, select the **Frequency resp** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

In addition to the frequency-response curve, you can display a confidence interval on the plot. The *confidence interval* corresponds to the range of response values with a specific probability of being the actual response of the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around the nominal curve represents the range of values that have a 95% probability of being the

true system response. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

The following table summarizes the Frequency Function plot settings.

Frequency Function Plot Settings

Action	Command
Display the confidence interval.	<ul style="list-style-type: none"> <li data-bbox="765 583 1307 739">• To display the dashed lines on either side of the nominal model curve, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. <li data-bbox="765 756 1307 852">• To change the confidence value, select Options > Set % confidence level, and choose a value from the list. <li data-bbox="765 869 1307 1065">• To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.

Frequency Function Plot Settings (Continued)

Action	Command
<p>Change the frequency values for computing the noise spectrum.</p> <p>The default frequency vector is 128 linearly distributed values, greater than zero and less than or equal to the Nyquist frequency.</p>	<p>Select Options > Frequency range and specify a new frequency vector in units of rad/s.</p> <p>Enter the frequency vector using any one of following methods:</p> <ul style="list-style-type: none"> • MATLAB expression, such as <code>[1:100]*pi/100</code> or <code>logspace(-3,-1,200)</code>. Cannot contain variables in the MATLAB Workspace browser. • Row vector of values, such as <code>[1:.1:100]</code> <hr/> <p>Note To restore the default frequency vector, enter <code>[]</code>.</p> <hr/>
<p>Change frequency units between hertz and radians per second.</p>	<p>Select Style > Frequency (Hz) or Style > Frequency (rad/s).</p>
<p>Change frequency scale between linear and logarithmic.</p>	<p>Select Style > Linear frequency scale or Style > Log frequency scale.</p>

Frequency Function Plot Settings (Continued)

Action	Command
Change amplitude scale between linear and logarithmic.	Select Style > Linear amplitude scale or Style > Log amplitude scale .
(Multiple-output system only) Select an input-output pair to view the noise spectrum corresponding to those channels.	Select the output by name in the Channel menu.
Note You cannot view cross spectra between different outputs.	

How to Plot Bode and Nyquist Plots at the Command Line

You can plot Bode and Nyquist plots for linear models using the `bode`, `ffplot`, and `nyquist` commands.

All plot commands have the same basic syntax, as follows:

- To plot one model, use the syntax `command(model)`.
- To plot several models, use the syntax `command(model1,model2,...,modelN)`.

In this case, `command` represents any of the plotting commands.

To display confidence intervals for a specified number of standard deviations, use the following syntax:

```
command(model, 'sd', sd)
```

where *sd* is the number of standard deviations of a Gaussian distribution. For example, a confidence value of 99% for the nominal model curve corresponds to 2.58 standard deviations.

To display a filled confidence region, use the following syntax:

```
command(model, 'sd', sd, 'fill')
```

The following table summarizes commands that generate Bode and Nyquist plots for linear models. For detailed information about each command and how to specify the frequency values for computing the response, see the corresponding reference page.

Command	Description	Example
bode	Plots the magnitude and phase of the frequency response on a logarithmic frequency scale.	To create the bode plot of the model <i>mod</i> , use the following command: <code>bode(mod)</code>
ffplot	Plots the magnitude and phase of the frequency response on a linear frequency scale (hertz).	To create the bode plot of the model <i>mod</i> , use the following command: <code>ffplot(mod)</code>
nyquist	Plots the imaginary versus real part of the transfer command. <hr/> Note Does not support time-series models. <hr/>	To plot the frequency response of the model <i>mod</i> , use the following command: <code>nyquist(mod)</code>

Creating Noise-Spectrum Plots

In this section...

“Supported Models” on page 10-37

“What Does a Noise Spectrum Plot Show?” on page 10-37

“Displaying the Confidence Interval” on page 10-38

“How to Plot the Noise Spectrum Using the GUI” on page 10-39

“How to Plot the Noise Spectrum at the Command Line” on page 10-42

Supported Models

When you estimate the noise model of your linear system, you can plot the spectrum of the estimated noise model. Noise-spectrum plots are available for all linear parametric models and spectral analysis (nonparametric) models.

Note For nonlinear models and correlation analysis models, noise-spectrum plots are not available. For time-series models, you can only generate noise-spectrum plots for parametric and spectral-analysis models.

What Does a Noise Spectrum Plot Show?

The general equation of a linear dynamic system is given by:

$$y(t) = G(z)u(t) + v(t)$$

In this equation, G is an operator that takes the input to the output and captures the system dynamics, and v is the additive noise term. In System Identification Toolbox, the noise term is treated as filtered white noise, as follows:

$$v(t) = H(z)e(t)$$

System Identification Toolbox computes both H and λ during the estimation of the noise model and stores these quantities as model properties. The

$H(z)$ operator represents the noise model. $e(t)$ is a white-noise source with variance λ .

Whereas the frequency-response plot in System Identification Toolbox shows the response of G , the noise-spectrum plot shows the frequency-response of the noise model H .

For input-output models, the noise spectrum is given by the following equation:

$$\Phi_v(\omega) = \lambda \left| H(e^{i\omega}) \right|^2$$

For time-series models (no input), the vertical axis of the noise-spectrum plot is the same as the dynamic model spectrum. These axes are the same because there is no input for time series and $y = He$.

Note You can avoid estimating the noise model by selecting the Output-Error model structure or by setting the `DisturbanceModel` property value to 'None' for a state space model. If you choose to not estimate a noise model for your system, then H and the noise spectrum amplitude are equal to 1 at all frequencies.

Displaying the Confidence Interval

In addition to the noise-spectrum curve, you can display a confidence interval on the plot. To learn how to show or hide confidence interval, see the description of the plot settings in “How to Plot the Noise Spectrum Using the GUI” on page 10-39.

The *confidence interval* corresponds to the range of power-spectrum values with a specific probability of being the actual noise spectrum of the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around the nominal curve represents the range of values that have a 95% probability of being the true system noise spectrum. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

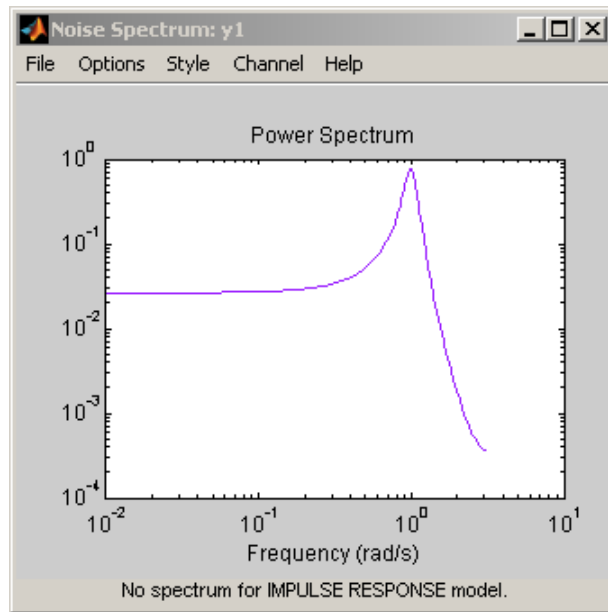
Note The calculation of the confidence interval assumes that the model sufficiently describes the system dynamics and the model residuals pass independence tests.

How to Plot the Noise Spectrum Using the GUI

To create a noise spectrum plot for parametric linear models in the GUI, select the **Noise spectrum** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

The following figure shows a sample Noise Spectrum plot.



The following table summarizes the Noise Spectrum plot settings.

Noise Spectrum Plot Settings

Action	Command
Display the confidence interval.	<ul style="list-style-type: none"> To display the dashed lines on either side of the nominal model curve, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. To change the confidence value, select Options > Set % confidence level, and choose a value from the list. To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.
<p>Change the frequency values for computing the noise spectrum.</p> <p>The default frequency vector is 128 linearly distributed values, greater than zero and less than or equal to the Nyquist frequency.</p>	<p>Select Options > Frequency range and specify a new frequency vector in units of radians per second.</p> <p>Enter the frequency vector using any one of following methods:</p> <ul style="list-style-type: none"> MATLAB expression, such as <code>[1:100]*pi/100</code> or <code>logspace(-3,-1,200)</code>. Cannot contain variables in the MATLAB Workspace browser. Row vector of values, such as <code>[1:.1:100]</code> <hr/> <p>Tip To restore the default frequency vector, enter <code>[]</code>.</p>
Change frequency units between hertz and radians per second.	Select Style > Frequency (Hz) or Style > Frequency (rad/s) .
Change frequency scale between linear and logarithmic.	Select Style > Linear frequency scale or Style > Log frequency scale .

Noise Spectrum Plot Settings (Continued)

Action	Command
Change amplitude scale between linear and logarithmic.	Select Style > Linear amplitude scale or Style > Log amplitude scale .
(Multiple-output system only) Select an input-output pair to view the noise spectrum corresponding to those channels. Note You cannot view cross spectra between different outputs.	Select the output by name in the Channel menu.

How to Plot the Noise Spectrum at the Command Line

To plot the frequency-response of the noise model, use a combination of System Identification Toolbox commands.

First, select the portion of the model object that corresponds to the noise model H . For example, to select the noise model in the model object m , type the following command:

```
m_noise=m('noise')
```

Tip You can abbreviate the command to `m_noise=m('n')`.

To plot the frequency-response of the noise model, use the `bode` command:

```
bode(m_noise)
```

To determine if your estimated noise model is good enough, you can compare the frequency-response of the estimated noise-model H to the estimated frequency response of $v(t)$. To compute $v(t)$, which represents the actual noise term in the system, use the following commands:

```
ysimulated = sim(m,data);
v = ymeasured-ysimulated;
```

`ymeasured` is `data.y`. `v` is the noise term $v(t)$, as described in “What Does a Noise Spectrum Plot Show?” on page 10-37 and corresponds to the difference between the simulated response `ysimulated` and the actual response `ymeasured`.

To compute the frequency-response model of the actual noise, use `spa`:

```
V = spa(v);
```

System Identification Toolbox uses the following equation to compute the noise spectrum of the actual noise:

$$\Phi_v(\omega) = \sum_{\tau=-\infty}^{\infty} R_v(\tau) e^{-i\omega\tau}$$

The covariance command R_v is given in terms of E , which denotes the mathematical expectation, as follows:

$$R_v(\tau) = E v(t) v(t - \tau)$$

To compare the parametric noise-model H to the (nonparametric) frequency-response estimate of the actual noise $v(t)$, use `bode`:

```
bode(V,m('noise'))
```

If the parametric and the nonparametric estimates of the noise spectra are different, then you might need a higher-order noise model.

Using Pole-Zero Plots to Validate Models

In this section...

“Supported Models” on page 10-44

“What Does a Pole-Zero Plot Show?” on page 10-44

“How to Plot Model Poles and Zeros Using the GUI” on page 10-45

“How to Plot Poles and Zeros at the Command Line” on page 10-47

“Reducing Model Order Using Pole-Zero Plots” on page 10-48

Supported Models

You can create pole-zero plots of linear polynomial, state-space, and grey-box models.

What Does a Pole-Zero Plot Show?

The general equation of a linear dynamic system is given by:

$$y(t) = G(z)u(t) + v(t)$$

In this equation, G is an operator that takes the input to the output and captures the system dynamics, and v is the additive noise term.


The *poles* of a linear system are the roots of the denominator of the transfer command G . The poles have a direct influence on the dynamic properties of the system. The *zeros* are the roots of the numerator of G . If you estimated a noise model H in addition to the dynamic model G , you can also view the poles and zeros of the noise model.

Zeros and the poles are equivalent ways of describing the coefficients of a linear difference equation, such as the ARX model. Poles are associated with the output side of the difference equation, and zeros are associated with the input side of the equation. The number of poles is equal to the number of sampling intervals between the most-delayed and least-delayed output. The number of zeros) is equal to the number of sampling intervals between the

most-delayed and least-delayed input. For example, there two poles and one zero in the following ARX model:

$$y(t) - 1.5y(t - T) + 0.7y(t - 2T) = 0.9u(t) + 0.5u(t - T)$$

The following figure shows a sample pole-zero plot of the model with confidence intervals. x indicate poles and o indicate zeros.



How to Plot Model Poles and Zeros Using the GUI

To create a pole-zero plot for parametric linear models in the System Identification Tool GUI, select the **Zeros and poles** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

In addition, you can display a confidence interval for each pole and zero on the plot. The *confidence interval* corresponds to the range of pole or zero values with a specific probability of being the actual pole or zero of the system. System Identification Toolbox uses the estimated uncertainty in the model parameters to calculate confidence intervals and assumes the estimates have a Gaussian distribution.

For example, for a 95% confidence interval, the region around the nominal pole or zero value represents the range of values that have a 95% probability of being the true system pole or zero value. You can specify the confidence interval as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution. For example, a probability of 0.99 (99%) corresponds to 2.58 standard deviations.

The following table summarizes the Zeros and Poles plot settings.

Zeros and Poles Plot Settings

Action	Command
Display the confidence interval.	<ul style="list-style-type: none"> • To display the dashed lines on either side of the nominal pole and zero values, select Options > Show confidence intervals. Select this option again to hide the confidence intervals. • To change the confidence value, select Options > Set % confidence level, and choose a value from the list. • To enter your own confidence level, select Options > Set confidence level > Other. Enter the value as a probability (between 0 and 1) or as the number of standard deviations of a Gaussian distribution.
Show real and imaginary axes.	Select Style > Re/Im-axes . Select this option again to hide the axes.

Zeros and Poles Plot Settings (Continued)

Action	Command
Show the unit circle.	Select Style > Unit circle . Select this option again to hide the unit circle.
(Multiple-output system only) Select an input-output pair to view the poles and zeros corresponding to those channels.	Select the output by name in the Channel menu.

How to Plot Poles and Zeros at the Command Line

You can create a pole-zero plot for linear polynomial, linear state-space, and linear grey-box models using the `pzmap` command. `pzmap` lets you include several models on a plot.

To display confidence intervals for a specified number of standard deviations, use the following syntax:

```
pzmap(model, 'sd', sd)
```

where `sd` is the number of standard deviations of a Gaussian distribution. For example, a confidence value of 99% for the nominal model curve corresponds to 2.58 standard deviations.

Command	Description	Example
<code>pzmap</code>	Plots zeros and poles of the model on the S-plane or Z-plane for continuous-time or discrete-time model, respectively.	To plot the poles and zeros of the model <code>mod</code> , use the following command: <code>pzmap(mod)</code>

For detailed information about `pzmap`, see the corresponding reference page.

Reducing Model Order Using Pole-Zero Plots

You can use pole-zero plots to evaluate whether it might be useful to reduce model order. When confidence intervals for a pole-zero pair overlap, this overlap indicates a possible pole-zero cancelation.

For example, you can use the following syntax to plot a 1-standard-deviation confidence interval around model poles and zeros.

```
pzmap(model, 'sd', 1)
```

If poles and zeros overlap, try estimating a lower order model.

Always validate model output and residuals to see if the quality of the fit changes after reducing model order. If the plot indicates pole-zero cancellations, but reducing model order degrades the fit, then the extra poles probably describe noise. In this case, you can choose a different model structure that decouples system dynamics and noise. For example, try ARMAX, Output-Error, or Box-Jenkins polynomial model structures with an A or F polynomial of an order equal to that of the number of uncanceled poles. For more information about estimating linear polynomial models, see “Estimating Polynomial Models” on page 5-21.

Using Nonlinear ARX Plots to Validate Models

In this section...

“About Nonlinear ARX Plots” on page 10-49

“How to Plot Nonlinear ARX Plots Using the GUI” on page 10-49

“Configuring the Nonlinear ARX Plot” on page 10-50

“Axis Limits, Legend, and 3-D Rotation” on page 10-51

“How to plot Nonlinear ARX Plots at the Command Line” on page 10-52

About Nonlinear ARX Plots

The Nonlinear ARX plot displays the characteristics of model nonlinearities as a command of one or two regressors. For more information about estimating nonlinear ARX models, see “Estimating Nonlinear ARX Models” on page 8-4.

Examining a nonlinear ARX plot can help you gain insight into which regressors have the strongest effect on the model output. Understanding the relative importance of the regressors on the output can help you decide which regressors should be included in the nonlinear command.

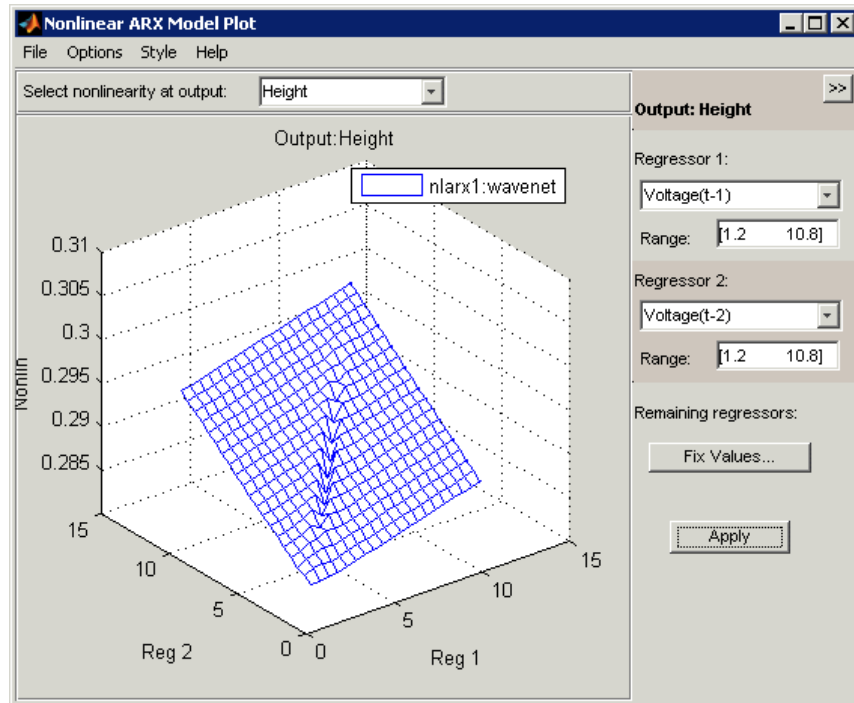
Furthermore, you can create several nonlinear models for the same data set using different nonlinearity estimators, such a wavelet network and tree partition, and then compare the nonlinear surfaces of these models. Agreement between nonlinear surfaces increases the confidence that these nonlinear models capture the true dynamics of the system.

How to Plot Nonlinear ARX Plots Using the GUI

To create a nonlinear ARX plot in the System Identification Tool GUI, select the **Nonlinear ARX** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

Note The **Nonlinear ARX** check box is unavailable if you do not have a nonlinear ARX model in the Model Board.


The following figure shows a sample nonlinear ARX plot.



Configuring the Nonlinear ARX Plot

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the Model Board icon.

To configure the plot:

- 1 If your model contains multiple output, select the output channel in the **Select nonlinearity at output** list. Selecting the output channel displays the nonlinearity values that correspond to this output channel.
- 2 If the regressor selection options are not visible, click  to expand the Nonlinear ARX Model Plot window.

- 3** Select **Regressor 1** from the list of available regressors. In the **Range** field, enter the range of values to include on the plot for this regressor. The regressor values are plotted on the **Reg1** axis.
- 4** Specify a second regressor for a 3-D plot by selecting one of the following types of options:
 - Select **Regressor 2** to display three axes. In the **Range** field, enter the range of values to include on the plot for this regressor. The regressor values are plotted on the **Reg2** axis.
 - Select <none> in the **Regressor 2** list to display only two axes.
- 5** To fix the values of the regressor that are not displayed, click **Fix Values**. In the Fix Regressor Values dialog box, double-click the **Value** cell to edit the constant value of the corresponding regressor. The default values are determined during model estimation. Click **OK**.
- 6** In the Nonlinear ARX Model Plot window, click **Apply** to update the plot.
- 7** To change the grid of the regressor space along each axis, **Options > Set number of samples**, and enter the number of samples to use for each regressor. Click **Apply** and then **Close**.

For example, if the number of samples is 20, each regressor variable contains 20 points in its specified range. For a 3-D plots, this results in evaluating the nonlinearity at $20 \times 20 = 400$ points.

Axis Limits, Legend, and 3-D Rotation

The following table summarizes the commands to modify the appearance of the Nonlinear ARX plot.

Changing Appearance of the Nonlinear ARX Plot

Action	Command
Change axis limits.	Select Options > Set axis limits to open the Axis Limits dialog box, and edit the limits. Click Apply .

Changing Appearance of the Nonlinear ARX Plot (Continued)

Action	Command
Hide or show the legend.	Select Style > Legend . Select this option again to show the legend.
(Three axes only) Rotate in three dimensions.	Select Style > 3D Rotate and drag the axes on the plot to a new orientation. To disable three-dimensional rotation, select Style > 3D Rotate again.
Note Available only when you have selected two regressors as independent variables.	

How to plot Nonlinear ARX Plots at the Command Line

You can plot the nonlinearity shape of nonlinear ARX models using the following syntax:

```
plot(model)
```

`model` must be an `idnlarx` model object. You can use additional plot arguments to specify the following information:

- Include multiple nonlinear ARX models on the plot.
- Configure the regressor values for computing the nonlinearity values.

The plot command opens the Nonlinear ARX Model Plot window. For more information about working with this plot window, see “Configuring the Nonlinear ARX Plot” on page 10-50 and “Axis Limits, Legend, and 3-D Rotation” on page 10-51.

For detailed information about `plot`, type the following command at the MATLAB prompt:

```
help idnlarx/plot
```

Using Hammerstein-Wiener Plots to Validate Models

In this section...

“About Hammerstein-Wiener Plots” on page 10-53

“How to Create Hammerstein-Wiener Plots in the GUI” on page 10-53

“How to Plot Hammerstein-Wiener Plots at the Command Line” on page 10-55

“Plotting Nonlinear Block Characteristics” on page 10-55

“Plotting Linear Block Characteristics” on page 10-56

About Hammerstein-Wiener Plots

Hammerstein-Wiener model plot lets you explore the characteristics of the linear block and the static nonlinearities of the Hammerstein-Wiener model. For more information about estimating nonlinear Hammerstein-Wiener models, see “Estimating Hammerstein-Wiener Models” on page 8-14.

Examining a Hammerstein-Wiener plot can help you determine whether you chose an unnecessarily complicated nonlinearity for modeling your system. For example, if you chose a piece-wise-linear nonlinearity (which is very general), but the plot indicates saturation behavior, then you can estimate a new model using the simpler saturation nonlinearity instead.

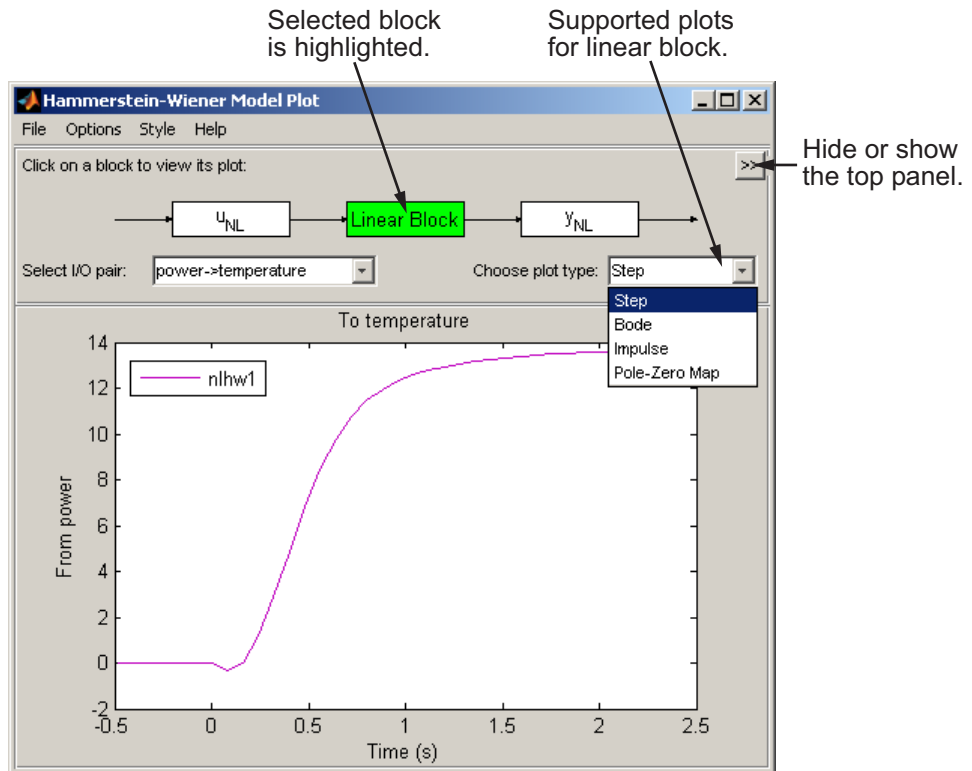
For multivariable systems, you can use the Hammerstein-Wiener plot to determine whether to exclude nonlinearities for specific channels. If the nonlinearity for a specific input or output channel does not exhibit strong nonlinear behavior, you can estimate a new model after setting the nonlinearity at that channel to unit gain.

How to Create Hammerstein-Wiener Plots in the GUI

To create a Hammerstein-Wiener plot in the System Identification Tool GUI, select the **Hamm-Wiener** check box in the **Model Views** area. For general information about creating and working with plots, see “Plots Data and Models in the GUI” on page 1-18.

Note The **Hamm-Wiener** check box is unavailable if you do not have a Hammerstein-Wiener model in the Model Board.

To include or exclude a model on the plot, click the corresponding model icon in the System Identification Tool GUI. Active models display a thick line inside the model icon, as shown in the following figure.



After you generate a plot, you can learn more about your model by:

- “Plotting Nonlinear Block Characteristics” on page 10-55
- “Plotting Linear Block Characteristics” on page 10-56

How to Plot Hammerstein-Wiener Plots at the Command Line

You can plot input and output nonlinearity and linear responses for Hammerstein-Wiener models using the following syntax:

```
plot(model)
```

`model` must be an `idnlhw` model object. You can use additional plot arguments to specify the following information:

- Include several Hammerstein-Wiener models on the plot.
- Configure how to evaluate the nonlinearity at each input and output channel.
- Specify the time or frequency values for computing transient and frequency response plots of the linear block.

The plot command opens the Hammerstein-Wiener Model Plot window. For more information about working with this plot window, see “Plotting Nonlinear Block Characteristics” on page 10-55 and “Plotting Linear Block Characteristics” on page 10-56.


For detailed information about `plot`, type the following command at the MATLAB prompt:

```
help idnlhw/plot
```

Plotting Nonlinear Block Characteristics

The Hammerstein-Wiener model can contain up to two nonlinear blocks. The nonlinearity at the input to the Linear Block is labeled u_{NL} and is called the *input nonlinearity*. The nonlinearity at the output of the Linear Block is labeled y_{NL} and is called the *output nonlinearity*.

To configure the plot, perform the following steps:

- 1 If the top pane is not visible, click  to expand the Hammerstein-Wiener Model Plot window.
- 2 Select the nonlinear block you want to plot:

- To plot u_{NL} as a command of the input data, click the u_{NL} block.
- To plot y_{NL} as a command of its inputs, click the y_{NL} block.

The selected block is highlighted in green.


Note An input to the output nonlinearity block y_{NL} is the output from the Linear Block and not the measured input data.

- 3** If your model contains multiple variables, select the channel in the **Select nonlinearity at channel** list. Selecting the channel updates the plot and displays the nonlinearity values versus the corresponding input to this nonlinear block.
- 4** To change the range of the horizontal axis, select **Options > Set input range** to open the Range for Input to Nonlinearity dialog box. Enter the range using the format [MinValue MaxValue]. Click **Apply** and then **Close** to update the plot.

Plotting Linear Block Characteristics

The Hammerstein-Wiener model contains one Linear Block that represents the embedded linear model.

To configure the plot:

- 1** If the top pane is not visible, click  to expand the Hammerstein-Wiener Model Plot window.
- 2** Click the Linear Block to select it. The Linear Block is highlighted in green.
- 3** In the **Select I/O pair** list, select the input and output data pair for which to view the response.
- 4** In the **Choose plot type** list, select the linear plot from the following options:
 - Step
 - Impulse

- Bode
- Pole-Zero Map

5 If you selected to plot step or impulse response, you can set the time span. Select **Options > Time span** and enter a new time span in units of time you specified for the model.

For a time span T , the resulting response is plotted from $-T/4$ to T . The default time span is 10.

Click **Apply** and then **Close**.

6 If you selected to plot a Bode plot, you can set the frequency range.

The default frequency vector is 128 linearly distributed values, greater than zero and less than or equal to the Nyquist frequency. To change the range, select **Options > Frequency range**, and specify a new frequency vector in units of rad per model time units.

Enter the frequency vector using any one of following methods:

- MATLAB expression, such as `[1:100]*pi/100` or `logspace(-3,-1,200)`. Cannot contain variables in the MATLAB Workspace browser.
- Row vector of values, such as `[1:.1:100]`.

Click **Apply** and then **Close**.

Using Akaike's Criteria to Validate Models

In this section...

"Definition of FPE" on page 10-58

"Computing FPE" on page 10-59

"Definition of AIC" on page 10-59

"Computing AIC" on page 10-60

Definition of FPE

Akaike's Final Prediction Error (FPE) criterion and his closely related Information Criterion (AIC) provide a measure of model quality by simulating the situation where the model is tested on a different data set.

If you use the same data set to both model estimation and validation, the fit always improves you increase the model order and the flexibility of the model structure increases.

After computing several different models, you can compare them using these criteria. According to Akaike's theory, the most accurate model has the smallest FPE and AIC.

Akaike's Final Prediction Error (FPE) is defined by the following equation:

$$FPE = V \left(\frac{1 + d/N}{1 - d/N} \right)$$

where V is the loss command, d is the number of estimated parameters, and N is the number of estimation data.

The loss command V is defined by the following equation:

$$V = \det \left(\frac{1}{N} \sum_1^N \varepsilon(t, \theta_N) (\varepsilon(t, \theta_N))^T \right)$$

where $\hat{\theta}_N$ represents the estimated parameters.

Computing FPE

Use the `fpe` command to compute Akaike's Final Prediction Error (FPE) criterion for one or more linear or nonlinear models, as follows:

$$\text{FPE} = \text{fpe}(m1, m2, m3, \dots, mN)$$

According to Akaike's theory, the most accurate model has the smallest FPE.

You can also access the FPE value of an estimated model by accessing the FPE field of the `EstimationInfo` property of this model. For example, if you estimated the model `m`, you can access its FPE using the following command:

```
m.EstimationInfo.FPE
```

Definition of AIC

Akaike's Information Criterion (AIC) is defined by the following equation:

$$AIC = \log V + \frac{2d}{N}$$

where V is the loss command, d is the number of estimated parameters, and N is the number of values in the estimation data set.

The loss command V is defined by the following equation:

$$V = \det \left(\frac{1}{N} \sum_1^N \varepsilon(t, \theta_N) (\varepsilon(t, \theta_N))^T \right)$$

where $\hat{\theta}_N$ represents the estimated parameters.

For $d \ll N$

$$AIC = \log \left(V + \left(1 + \frac{2d}{N} \right) \right)$$

Computing AIC

Use the `aic` command to compute Akaike's Information Criterion (AIC) for one or more linear or nonlinear models, as follows:

```
AIC = aic(m1,m2,m3,...,mN)
```

According to Akaike's theory, the most accurate model has the smallest AIC.

Computing Model Uncertainty

In this section...

“Why Analyze Model Uncertainty” on page 10-61

“What Is Model Covariance?” on page 10-61

“Viewing Model Uncertainty Information” on page 10-62

Why Analyze Model Uncertainty

In addition to estimating model parameters, the System Identification Toolbox estimation algorithms also estimate variability of the model parameters that result from random disturbances in the output.

Understanding model variability helps you to understand how different your model parameters would be if you repeated the estimation using a different data set (with the same input sequence as the original data set) and the same model structure.

When validating your parametric models, check the uncertainty values. Large uncertainties in the parameters might be caused by high model orders, inadequate excitation, and poor signal-to-noise ratio in the data.

Note You can get model uncertainty data for linear parametric black-box models, and both linear and nonlinear grey-box models. Supported model objects include `idproc`, `idpoly`, `idss`, `idarx`, `idgrey`, `idfrd`, and `idnlgrey`.

What Is Model Covariance?

Uncertainty in the model is called *model covariance*.

If you estimate model uncertainty data, this information is stored in the `Model.CovarianceMatrix` model property. The covariance matrix is used to compute all uncertainties in model output, Bode plots, residual plots, and pole-zero plots.

Computing the covariance matrix is based on the assumption that the model structure gives the correct description of the system dynamics. For models that include a disturbance model H , a correct uncertainty estimate assumes that the model produces white residuals. To determine whether you can trust the estimated model uncertainty values, perform residual analysis tests on your model, as described in “Using Residual Analysis Plots to Validate Models” on page 10-15. If your model passes residual analysis tests, there is a good chance that the true system lies within the confidence interval and any parameter uncertainties results from random disturbances in the output.

In the case of output-error models, where the noise model H is fixed to 1, computing the covariance matrix does not assume that the residuals are white. Instead, the covariance is estimated based on the estimated color of the residual correlations. This estimation of the noise color is also performed for state-space models with $K=0$, which is equivalent to an output-error model.

Viewing Model Uncertainty Information

You can view the following uncertainty information from linear and nonlinear grey-box models:

- Uncertainties of estimated parameters.

Type `present(model)` at the MATLAB prompt, where `model` represents the name of a linear or nonlinear model.

- Confidence intervals on the linear model plots, including step-response, impulse-response, Bode, and pole-zero plots.

Confidence intervals are computed based on the variability in the model parameters. For information about displaying confidence intervals, see the corresponding plot section.

- Covariance matrix of the estimated parameters in linear and nonlinear grey-box models.

Type `model.CovarianceMatrix` at the MATLAB prompt, where `model` represents the name of the model object.

- Estimated standard deviations of polynomial coefficients or state-space matrices

Type `model.dA` at the MATLAB prompt to access the estimated standard deviations of the `model.A` estimated property, where `model` represents the name of the model object, and `A` represents any estimated model property. In general, you prefix the name of the estimated property with a `d` to get the standard deviation estimate for that property. For example, to get the standard deviation value of the `A` polynomial in an estimated ARX model, type `model.da`.

Note State-space models, estimated with free parameterization, do not have well-defined standard deviations of the matrix elements. To display matrix parameter uncertainties in this case, first transform the model to a canonical parameterization by setting the `ss` model property to `model.ss = 'canon'`. For more information about free and canonical parameterizations, see “Estimating State-Space Models” on page 5-48.

- Simulated output values for linear models with standard deviations using the `sim` command.

Call the `sim` command with output arguments, where the second output argument is the estimated standard deviation of each output value. For example, type `[ysim,ysimsd]=sim(model,data)`, where `ysim` is the simulated output, `ysimsd` contains the standard deviations on the simulated output, and `data` is the simulation data.

Extracting Parameter Values from Linear Models

You can use System Identification Toolbox to extract the numerical parameter values and uncertainties of model objects and store these values using double data format.

For example, you can extract state-space matrices for state-space models, and extract polynomials for polynomial models. You can operate on extracted model data as you would on any other MATLAB vectors and matrices. You can also pass these numerical values to Control System Toolbox, for example, or Simulink blocks.

If you specified to estimate model uncertainty data, this information is stored in the property `Model.CovarianceMatrix` in the estimated model. The covariance matrix is used to compute uncertainties in parameter estimates, model output plots, Bode plots, residual plots, and pole-zero plots.

The following table summarizes commands for extracting numerical data from models. All of these commands have the following syntax form:

$$[G, dG] = \text{command}(\text{model})$$

where `G` stores model parameters and `dG` stores standard deviation of parameters or covariance.

Commands for Extracting Numerical Model Data

Command	Description	Syntax
<code>arxdata</code>	Extracts ARX parameters from multiple-output <code>idarx</code> or single-output <code>idpoly</code> objects that represent ARX models.	<code>[A,B,dA,dB] = arxdata(m)</code>
<code>freqresp</code>	Extracts frequency-response data from any <code>idmodel</code> or <code>idfrd</code> object.	<code>[H,w,CovH] = freqresp(m)</code>

Commands for Extracting Numerical Model Data (Continued)

Command	Description	Syntax
polydata	Extracts polynomials from any single-output idmodel object.	$[A, B, C, D, F, dA, dB, dC, dD, dF] = \dots$ polydata(m)
ssdata	Extracts state-space matrices from any idmodel object.	$[A, B, C, D, K, X0, \dots$ $dA, dB, dC, dD, dK, dX0] = \dots$ ssdata(Model)
tfdata	Extracts numerator and denominator polynomials from any idmodel object.	$[Num, Den, dNum, dDen] = \dots$ tfdata(Model)
zpkdata	Extracts zeros, poles, and transfer command gains from any idmodel object.	$[Z, P, K, covZ, covP, covK] = \dots$ zpkdata(m)

Extracting Dynamic Model and Noise Model Separately

You can extract the numerical data associated with a dynamic model and the noise model separately.

For linear models, the general symbolic model description is given by:

$$y = Gu + He$$

G is an operator that takes the measured inputs u to the outputs and captures the system dynamics. H is an operator that describes the properties of the additive output disturbance and takes the hypothetical (unmeasured) noise source inputs e to the outputs, also called the *noise model*. When you estimate a noise model, System Identification Toolbox includes one noise channel at the input e for each output in your system.

The following table summarizes the results of `ssdata`, `tfdata`, and `zpkdata` commands for extracting the numerical values of the dynamic model and noise model separately. fcn represents `ssdata`, `tfdata`, and `zpkdata`, and m is a model object. L represents the covariance matrix e , as defined in “Subreferencing Measured and Noise Models” on page 3-21.

For information about subreferencing noise channels or treating noise channels as measured input, see “Subreferencing Models” on page 3-20.

Note The syntax `fcn(m('noise'))` is equivalent to `fcn(m('n'))`.

Syntax for Extracting Transfer-Function Data

Command	Syntax
<code>fcn(m)</code>	Returns the properties of G for ny outputs and nu inputs.
<code>fcn(m('noise'))</code>	Returns the properties of H for ny outputs and ny inputs.
<code>fcn(noisecnv(m))</code>	Returns the properties of $[G \ H]$ ny outputs and $ny+nu$ inputs.

Syntax for Extracting Transfer-Function Data (Continued)

Command	Syntax
<code>fcn(noiseconv(m, 'Norm'))</code>	Returns the properties of $[G \ HL]$ ny outputs and $ny+nu$ inputs.
<code>fcn(noiseconv(m('noise'), 'Norm'))</code>	Returns the properties of HL ny outputs and ny inputs.
<code>fcn(m)</code>	If m is a time-series model, returns the properties of H .
<code>fcn(noiseconv(m, 'Norm'))</code>	If m is a time-series model, returns the properties of HL .

Note The estimated covariance matrix `NoiseVariance` is uncertain. Thus, the uncertainty of H differs from the uncertainty of HL .

Troubleshooting Models

In this section...
“About Troubleshooting Models” on page 10-68
“Model Order Is Too High or Too Low” on page 10-68
“Nonlinearity Estimator Produces a Poor Fit” on page 10-69
“Substantial Noise in the System” on page 10-70
“Unstable Models” on page 10-70
“Missing Input Variables” on page 10-71
“Complicated Nonlinearities” on page 10-72

About Troubleshooting Models

During validation, you might find that your model output fits the validation data poorly. You might also find some unexpected or undesirable model characteristics.

If the tips suggested in these sections do not help improve your models, then a good model might not be possible for this data. For example, your data might have poor signal-to-noise ratio, large and nonstationary disturbances, or varying system properties.

Model Order Is Too High or Too Low

When the Model Output plot does not show a good fit, there is a good chance that you need to try a different model order. System identification is largely a trial-and-error process when selecting model structure and model order. Ideally, you want the lowest-order model that adequately captures the system dynamics.

System Identification Toolbox provides assistance in finding the approximate model order, as described in “Preliminary Step – Estimating Model Orders and Input Delays ” on page 5-29. Typically, you use the suggested order as a starting point to estimate the lowest possible order with different model structures. After each estimation, you monitor the Model Output and the Residual Analysis plots, and then adjust your settings for the next estimation.

When a low-order model fits the validation data poorly, try estimating a higher-order model to see if the fit improves. For example, if a Model Output plot shows that a fourth-order model gives poor results, try estimating an eighth-order model. When a higher-order model improves the fit, you can conclude that higher-order models might be required and linear models might be sufficient.

You should use an independent data set to validate your models. If you use the same data set to both estimate and validate a model, the fit always improves as you increase model order, and you risk overfitting. However, if you use an independent data set to validate your models, the fit eventually deteriorates if your model orders are too high.

High-order models are more expensive to compute and result in greater parameter uncertainty.

Nonlinearity Estimator Produces a Poor Fit

In the case of nonlinear ARX and Hammerstein-Wiener models, the Model Output plot does not show a good fit when the nonlinearity estimator has incorrect complexity.

You specify the complexity of piece-wise-linear, wavelet, sigmoid, and custom networks using the number of units (NumberOfUnits nonlinearity estimator property). A high number of units indicates a complex nonlinearity estimator. In the case of neural networks, you specify the complexity using the parameters of the network object. For more information, see the Neural Network Toolbox documentation.

To select the appropriate complexity of the nonlinearity estimator, start with a low complexity and validate the model output. Next, increase the complexity and validate the model output again. The model fit degrades when the nonlinearity estimator becomes too complex.

Note To see the model fit degrade when the nonlinearity estimator becomes too complex, you must use an independent data set to validate the data that is different from the estimation data set.

Substantial Noise in the System

There are a couple of indications that you might have substantial noise in your system and might need to use linear model structures that are better equipped to model noise.

One indication of noise is when a state-space model is better than an ARX model at reproducing the measured output; whereas the state-space structure has sufficient flexibility to model noise, the ARX model structure is less able to model noise because the A polynomial must account for both the system dynamics and the noise. The following equation represents the ARX model and shows that A couples the dynamics and the noise by appearing in the denominator of both the dynamics term and the noise terms:

$$y = \frac{B}{A}u + \frac{1}{A}e$$

Another indication that a noise model is needed appears in residual analysis plots when you see significant autocorrelation of residuals at nonzero lags. For more information about residual analysis, see “Using Residual Analysis Plots to Validate Models” on page 10-15.

To model noise more carefully, use the ARMAX or the Box-Jenkins model structure, where the dynamics term and the noise term are modeled by different polynomials.

Unstable Models

One of the most conclusive approaches to determining whether a linear model is unstable is by examining the pole-zero plot of the model, which is described in “Using Pole-Zero Plots to Validate Models” on page 10-44. The stability threshold for pole values differs for discrete-time and continuous-time models, as follows:

- For stable continuous-time models, the real part of the pole is less than 0.
- For stable discrete-time models, the magnitude of the pole is less than 1.

In some cases, an unstable model is still a useful model. For example, your system might be unstable without a controller, and you plan to use your

model for control design. In this case, you can import your unstable model into Simulink or Control System Toolbox.

One way to check if a nonlinear model is unstable is to plot the simulated model output on top of the validation data. If the simulated output diverges from measured output, the model is unstable. However, agreement between model output and measured output does not guarantee stability.

In the case of linear models, if you believe that your system is stable, but your model is unstable, then you can estimate the model again with a Focus setting that guarantees stability. For example, set Focus to Stability to find the best stable model. This setting might result in a reduced model quality. For more information about Focus, see the Algorithm Properties reference page.

A more advanced approach to achieving a stable model is by setting the stability threshold property to allow a margin of error. The threshold model property is accessed as a field in the algorithm structure:

- For continuous-time models, set the value of `model.algorithm.advanced.sstability`. The model is considered stable if the pole on the far right is to the left of `sstability` threshold.
- For discrete-time models, set the value of `model.algorithm.advanced.zstability`. The model is considered stable if all poles inside the circle centered at the origin and with a radius `zstability`.

For more information about Threshold fields for linear models, see the Algorithm Properties reference page.

Missing Input Variables

If the Model Output plot and Residual Analysis plot shows a poor fit and you have already tried different structures and orders and modeled noise, it might be that there are one or more missing inputs that have a significant effect on the output.

Try including other measured signals in your input data, and then estimating the models again.

Inputs need not be control signals. Any measurable signal can be considered an input, including measurable disturbances.

Complicated Nonlinearities

If the Model Output plot and Residual Analysis plot shows a poor fit, consider if nonlinear effects are present in the system.

You can model the nonlinearities by performing a simple transformation on the signals to make the problem linear in the new variables. For example, if electrical power is the driving stimulus in a heating process and temperature is the output, you can form a simple product of voltage and current measurements.

If your problem is sufficiently complex and you do not have physical insight into the problem, you might try fitting nonlinear black-box models. For more information, see Chapter 8, “Estimating Nonlinear Black-Box Models”.

Next Steps After Getting an Accurate Model

After you get an accurate model, you can simulate or predict model output.

For linear parametric models (`idmodel` objects), you can perform the following operations:

- Transform between continuous-time and discrete-time representation.
- Transform between linear model representations, such as between polynomial, state-space, and zero-pole representations.
- Extract numerical data from transfer functions, pole-zero models, and state-space matrices.

System Identification Toolbox models in the MATLAB Workspace browser are immediately available to other MathWorks products. However, if you used the System Identification Tool GUI to estimate models, you must first export the models to the MATLAB Workspace Browser.

Tip To export a model from the GUI, drag the model icon to the **To Workspace** rectangle. For more information about working with the GUI, see Chapter 1, “Working with the System Identification Tool GUI”.

If you have Control System Toolbox, you can import your linear plant model for control-system design. For more information, see “Using Models with Control System Toolbox” on page 13-2.

Finally, if you have Simulink, you can exchange data between System Identification Toolbox and the Simulink environment. For more information, see “Using Models with Simulink” on page 13-8.

Transforming Model Representation

Transforming Between
Discrete-Time and Continuous-Time
Representations (p. 11-2)

Converting linear polynomial
and state-space models between
discrete-time and continuous-time
representations.

Transforming Between Linear Model
Representations (p. 11-7)

Converting between state-space,
polynomial, and frequency-response
representations.

Transforming Between Discrete-Time and Continuous-Time Representations

In this section...

“Why Transform Between Continuous and Discrete Time?” on page 11-2

“Using the c2d, d2c, and d2d Commands” on page 11-2

“Specifying Intersample Behavior” on page 11-4

“How d2c Handles Input Delays” on page 11-4

“Effects on the Noise Model” on page 11-5

Why Transform Between Continuous and Discrete Time?

Transforming between continuous-time and discrete-time representations is useful, for example, if you have estimated a discrete-time linear model and require a continuous-time model instead.

d2d is useful if you want to change the sampling interval of a discrete model. All of these operations change the sampling interval, which is called *resampling* the model.

Using the c2d, d2c, and d2d Commands

You can use c2d and d2c to transform any `idmodel` object between continuous-time and discrete-time representations.

The following table summarizes the commands for transforming between continuous-time and discrete-time model representations. These commands also transform the estimated model uncertainty, which corresponds to the estimated covariance matrix of the parameters. For detailed information about these commands, see the corresponding references page.

Note `c2d` and `d2d` correctly approximate the transformation of the noise model when the sampling interval T is small compared to the bandwidth of the noise.

Command	Description	Usage Example
<code>c2d</code>	Converts continuous-time models to discrete-time models.	To transform a continuous-time model <code>mod_c</code> to a discrete-time form, use the following command: $\text{mod_d} = \text{c2d}(\text{mod_c}, T)$ where T is the sampling interval of the discrete-time model.
<code>d2c</code>	Converts parametric discrete-time models to continuous-time models.	To transform a discrete-time model <code>mod_d</code> to a continuous-time form, use the following command: $\text{mod_c} = \text{d2c}(\text{mod_d})$
<code>d2d</code>	Resample a linear discrete-time model and produce an equivalent discrete-time model with a new sampling interval. You can use the resampled model to simulate or predict output with a specified time interval.	To resample a discrete-time model <code>mod_d1</code> to a discrete-time form with a new sampling interval T_s , use the following command: $\text{mod_d2} = \text{d2d}(\text{mod_d1}, T_s)$

The following commands compare estimated model `m` and its continuous-time counterpart `mc` on a Bode plot:

```
% Estimate discrete-time ARMAX model
% from the data
m = armax(data,[2 3 1 2]);
% Convert to continuous-time form
mc = d2c(m);
% Plot bode plot for both models
bode(m,mc)
```

Specifying Intersample Behavior

A sampled signal is characterized only by its values at the sampling instants. However, when you apply a continuous-time input to a continuous-time system, the output values at the sampling instants depend on the inputs at the sampling instants and on the inputs between these points. Thus, the `InterSample` data property describes how the algorithms should handle the input between samples. For example, you can specify the behavior between the samples to be piece-wise constant (zero-order hold, `zoh`) or linearly interpolated between the samples (first order hold, `foh`). The transformation formulas for `c2d` and `d2c` are affected by the intersample behavior of the input.

By default, `c2d` and `d2c` use the intersample behavior you assigned to the estimation data. To override this setting during transformation, add an extra argument in the syntax. For example:

```
% Set first-order hold intersample behavior
mod_d = c2d(mod_c,T,'foh')
```

How `d2c` Handles Input Delays

The discrete-to-continuous-time conversion `d2c` properly handles any input delays in the discrete-time model, and stores this information in the continuous-time model. An *input delay* is the delay in the response of the output to the input signal.

The relationship between discrete-time and continuous-time delays depends on the input intersample behavior. For example, a continuous-time system without a delay shows a delay when sampled with a zero-order-hold input.

A delay in the discrete-time model that corresponds to an actual delay in the continuous-time model is stored in the `InputDelay` property of the resulting continuous-time model. Typically, this `InputDelay` is $(n_k - 1) / T_s$, where n_k is the delay of the discrete-time system and T_s is the sampling interval.

Note Unlike for discrete-time models, the `nk` property of continuous-time model is only used to flag when immediate response to step changes is present; `nk` is not used to store input delays greater than or equal to 1. When `nk(i)=0`, then there is an immediate response to a step change in the input *i*th. When `nk(i)=1`, then there is no immediate response to the input.

Effects on the Noise Model

`c2d`, `d2c`, and `d2d` change the sampling interval of both the dynamic model and the noise model. Resampling a model affects the variance of its noise model.

A parametric noise model is a time-series model with the following mathematical description:

$$y(t) = H(q)e(t)$$

$$Ee^2 = \lambda$$

The noise spectrum is computed by the following discrete-time equation:

$$\Phi_v(\omega) = \lambda T \left| H(e^{i\omega T}) \right|^2$$

where λ is the variance of the white noise $e(t)$, and λT represents the spectral density of $e(t)$. Resampling the noise model preserves the spectral density λT . The spectral density λT is invariant up to the Nyquist frequency. For more information about spectrum normalization, see “Understanding Spectrum Normalization” on page 4-19.

`d2d` resampling of the noise model affects simulations with noise using `sim`. If you resample a model to a faster sampling rate, simulating this model results in higher noise level. This higher noise level results from the

underlying continuous-time model being subject to continuous-time white noise disturbances, which have infinite, instantaneous variance. In this case, the *underlying continuous-time model* is the unique representation for discrete-time models. To maintain the same level of noise after interpolating

the noise signal, scale the noise spectrum by $\sqrt{T_{New}/T_{Old}}$, where T_{new} is the new sampling interval and T_{old} is the original sampling interval. before applying sim.

c2d and d2c transformations produce warnings when the continuous-time disturbance model does not have the required white-noise component. These warnings occur because the underlying state-space model, which is formed and used by these transformations, is ill-defined. In this case, modify the C -polynomial such that the degree of the monic C -polynomial in continuous-time equals the sum of the degrees of the monic A - and D -polynomials in continuous-time. For example:

$$\text{length}(C) - 1 = (\text{length}(A) - 1) + (\text{length}(D) - 1)$$

Transforming Between Linear Model Representations

You can transform linear models between state-space and polynomial forms. You can also transform between frequency-response, state-space, and polynomial forms.

If you used the System Identification Tool GUI to estimate models, you must export the models to the MATLAB Workspace Browser before converting models.

For detailed information about each command in the following table, see the corresponding reference page.

Commands for Transforming Model Representations

Command	Model Type to Convert	Usage Example
idfrd	Converts any single- or multiple-output <code>idmodel</code> object to <code>idfrd</code> model. If you have Control System Toolbox, this command converts any <code>LTI</code> object.	<p>To get frequency response of <code>m</code> at default frequencies, use the following command:</p> <pre>m_f = idfrd(m)</pre> <p>To get frequency response at specific frequencies, use the following command:</p> <pre>m_f = idfrd(m,f)</pre> <p>To get frequency response for a submodel from input 2 to output 3, use the following command:</p> <pre>m_f = idfrd(m(2,3))</pre>

Commands for Transforming Model Representations (Continued)

Command	Model Type to Convert	Usage Example
idpoly	Converts single-output idmodel object to ARMAX representation. If you have Control System Toolbox, this command converts any single-output LTI object except frd.	To get an ARMAX model from state-space model <code>m_ss</code> , use the following command: <code>m_p = idpoly(m_ss)</code>
idss	Converts any single- or multiple-output idmodel object to state-space representation. If you have Control System Toolbox, this command converts any LTI object except frd.	To get a state-space model from an ARX model <code>m_arx</code> , use the following command: <code>m_ss = idss(m_arx)</code>

Note The `idss` conversion produces warnings when the continuous-time disturbance model does not have the required white-noise component. These warnings occur because the underlying state-space model, which is formed and used by these transformations, is ill-defined. In this case, modify the *C*-polynomial such that the degree of the monic *C*-polynomial in continuous-time equals the sum of the degrees of the monic *A*- and *D*-polynomials in continuous-time. For example:

$$\text{length}(C) - 1 = (\text{length}(A) - 1) + (\text{length}(D) - 1)$$

Simulating and Predicting Model Output

Simulating Versus Predicting Output (p. 12-2)

Understanding the difference between simulated and predicted output.

Simulation and Prediction in the GUI (p. 12-4)

How to simulate and predict output using the System Identification Tool GUI.

Example – Simulating Model Output with Noise at the Command Line (p. 12-5)

How to create input data and a model, and then use the data and the model to simulate output data.

Example – Simulating a Continuous-Time State-Space Model at the Command Line (p. 12-6)

How to simulate a continuous-time state-space model using a random binary input.

Predicting Model Output at the Command Line (p. 12-7)

How to predict model output using the predict command.

Specifying Initial States (p. 12-8)

When to specify initial states for simulation and prediction, setting initial states to zero or equilibrium values, and estimating initial states from the data.

Simulating Versus Predicting Output

Simulating a model means that you compute the response of a model to a particular input. Simulation does not involve the noise model unless you explicitly specify to compute the response to the noise source input. *Predicting* future outputs of a model from previous data over a time horizon of k samples or kT s time units—where T s is the sampling interval—requires both past inputs and past outputs.

The main difference between simulation and prediction is whether System Identification Toolbox uses measured or computed previous outputs for calculating the next output.

Simulating models uses only past input values to compute the output values. If the model-output expression includes past outputs, System Identification Toolbox computes the first output value using the initial conditions and the inputs. Then, System Identification Toolbox feeds this computed output into the difference equation or differential equation for calculating the next output value. Thus, no past outputs are used in the computation of output at the current time.

Using a model for prediction is common in controls applications where you want to predict output for a specific number of steps in advance. When you use System Identification Toolbox to *predict* model output, the algorithm uses both the measured and the calculated output data values in the difference equation for computing the next output. The k -step-ahead-predicted values of $y(t)$ is computed from all available inputs $u(s)$ and relevant previous outputs $y(s)$ —where $s \leq (t - k)$. The argument s represents the data sample number.

To make sure that the model picks up important dynamic properties, let the predicted time horizon kT be larger than the system time constants, where T is the sampling interval. Prediction with $k=\infty$ means that no previous inputs are used in the computation and prediction matches simulation.

Note Output-error models, state-space models with K set to zero, polynomial models with $n_a=n_c=n_d=0$, and nonlinear grey-box models do not use past outputs. In these cases, the simulated and predicted outputs are the same for any prediction horizon k .

Simulation and Prediction in the GUI

To learn how to display simulated or predicted output using the System Identification Tool GUI, see the description of the plot settings in “How to Plot Model Output Using the GUI” on page 10-10.

For information about simulating identified models in the Simulink environment, see “Using Models with Simulink” on page 13-8.

Example – Simulating Model Output with Noise at the Command Line

This example demonstrates how you can create input data and a model, and then use the data and the model to simulate output data. In this case, you use the following ARMAX model with Gaussian noise e :

$$y(t) - 1.5y(t-1) + 0.7y(t-2) = u(t-1) + 0.5u(t-2) + e(t) - e(t-1) + 0.2e(t-1)$$

Create the ARMAX model and simulate output data with random binary input u using the following commands:

```
% Create an ARMAX model
m_armax = idpoly([1 -1.5 0.7],[0 1 0.5],[1 -1 0.2]);
% Create a random binary input
u = idinput(400,'rbs',[0 0.3]);
% Simulate the output data
y = sim(m_armax,u,'noise');
```

Note The argument 'noise' specifies to include in the simulation the Gaussian noise e present in the model. Omit this argument to simulate the noise-free response to the input u , which is equivalent to setting e to zero.

Example – Simulating a Continuous-Time State-Space Model at the Command Line

This example demonstrates how to simulate a continuous-time state-space model using a random binary input u and a sampling interval of 0.1 s.

Consider the following state-space model:

$$\dot{x} = \begin{bmatrix} -1 & 1 \\ -0.5 & 0 \end{bmatrix} x + \begin{bmatrix} 1 \\ 0.5 \end{bmatrix} u + \begin{bmatrix} 0.5 \\ 0.5 \end{bmatrix} e$$

$$y = [1 \ 0] x + e$$

where e is Gaussian white noise with variance 7.

Use the following commands to simulate the model:

```
% Set up the model matrices
A = [-1 1; -0.5 0]; B = [1; 0.5];
C = [1 0]; D = 0; K = [0.5; 0.5];
% Create a continuous-time state-space model
% Ts = 0 indicates continuous time
model_ss = idss(A,B,C,D,K,'Ts',0,'NoiseVariance',7)
% Create a random binary input
u = idinput(400,'rbs',[0 0.3]);
% Create an iddata object with empty output
data = iddata([],u);
data.ts = 0.1
% Simulate the output using the model
y=sim(model_ss,data,'noise');
```

Note The argument 'noise' specifies to simulate with the Gaussian noise e present in the model. Omit this argument to simulate the noise-free response to the input u , which is equivalent to setting e to zero.

Predicting Model Output at the Command Line

You can use System Identification Toolbox to predict model output.

Use the following syntax to compute k -step-ahead prediction of the output signal using model m :

```
yhat = predict(m,[y u],k)
```

The predicted value $\hat{y}(t|t-k)$ is computed using information in $u(s)$ up to time $s=t$, and then information in $y(s)$ up to time $s=t-kT$, where T is the sampling interval.

The way information in past outputs is used depends on the disturbance model of m . For example, because $H = 1$ in the output-error model, there is no information in past outputs. In this case, predictions and simulations coincide.

The following example demonstrates commands you can use to evaluate how well a time-series model predicts future values. In this case, y is the original series of monthly sales figures. The first half of the measured data is used to estimate the time-series model, and then the second half of the data is used to predict half a year ahead.

```
% Split time-series data into  
% two halves  
y1 = y(1:48),  
y2 = y(49:96)  
% Estimate a fourth-order autoregressive model  
% using the first half of the data.  
m = ar(y1,4)  
% Predict time-series output  
yhat = predict(m4,y2,6)  
% Plot predicted output  
plot(y2,yhat)
```

Specifying Initial States

In this section...
“When to Specify Initial States” on page 12-8
“Setting Initial States to Zero” on page 12-8
“Setting Initial States to Equilibrium Values” on page 12-9
“Estimating Initial States from the Data” on page 12-9

When to Specify Initial States

The `sim` and `predict` commands require initial states to start the computations. If you prefer to use different initial states for state-space models, or if you are working with other model types, you must specify the initial states for simulation or prediction.

By default, `simulating` or `predicting` output for state-space models uses the initial states stored in the `X0` model property after estimation. For multiexperiment state-space models, the stored initial states correspond to the data in the last experiment. These stored initial states might not be appropriate when you simulate or predict model output using new data.

Use the following general syntax for specifying initial states for simulation or prediction:

```
y=sim(model,data,'InitialState',S)
y=predict(model,data,'InitialState',S)
```

where `S` represents a vector of states.

Note The `compare` command automatically estimates the initial states from the data and ensures consistency.

Setting Initial States to Zero

If the system starts at rest, or if transient effects are not important, then you can set the initial states to zero.

You can use the following shortcut syntax for setting initial states to zero:

```
y=sim(model,data,'InitialState','z')
y=predict(model,data,'InitialState','z')
```

Setting Initial States to Equilibrium Values

If you have physical insight about the starting point of the system, create a vector of specific initial states at the command line.

Use the following syntax to specify initial states for simulation or prediction:

```
y=sim(model,data,'InitialState',S)
y=predict(model,data,'InitialState',S)
```

where S represents a vector of initial states.

If you are working with multiexperiment data, specify S as a matrix containing as many columns as there are experiments.

Estimating Initial States from the Data

- “How to Estimate States for Simulation” on page 12-9
- “How to Estimate States for Prediction” on page 12-10

How to Estimate States for Simulation

Simulation using data from a different experiment than the data used to estimate the model requires the corresponding initial states.

To use `sim` with a data set that differs from the data you used to estimate the model, first estimate the new initial states `Sest` using `pe`:

```
[E,XOest]=pe(model,data)
```

Next, specify the estimated initial states `Sest` as an argument in `sim`. For example:

```
y=sim(model,data,'InitialState',Sest)
```

When you simulate a multiexperiment model, use the `pe` command to estimate initial states for the data from that specific experiment. For example, suppose you estimate a three-state model `M` using a merged data set `Z`, which contains data from five experiments—`z1`, `z2`, `z3`, `z4`, and `z5`:

```
Z = merge(z1,z2,z3,z4,z5);  
M = n4sid(Z,3);
```

If you want to simulate using data from `z2`, you must estimate the initial states for the second experiment `Z(z2.u)`, as follows:

```
[E,X0est] = pe(M,getexp(Z,2))
```

where `getexp(Z,2)` gets the data in `z2`. The estimated states matrix `Sest` contains one column of initial-state values for each experiment.

To simulate with these initial states, specify the estimated initial states `Sest` as an argument in `sim`. For example:

```
y=sim(M,getexp(Z,2),'InitialState',Sest)
```

How to Estimate States for Prediction

Prediction using data from a different experiment than the data used to estimate the model requires the corresponding initial states.

Unlike for `sim`, you can specify to estimate the initial states directly in the `predict` command.

To estimate the initial states that correspond to the data set you use for prediction, use the following syntax:

```
y=predict(M,data,'InitialState','Estimate')
```

Using Models with Other Products

Using Models with Control System
Toolbox (p. 13-2)

Preparing System Identification
Toolbox plant models for control
design in Control System Toolbox.

Using Models with Simulink (p. 13-8)

Using the System Identification
block library in Simulink.

Using Models with Control System Toolbox

In this section...
“How Control System Toolbox Works with Estimated Models” on page 13-2
“Using balred to Reduce Model Order” on page 13-3
“Compensator Design Using Control System Toolbox” on page 13-3
“Converting Models to LTI Objects” on page 13-4
“Viewing Model Response Using the LTI Viewer” on page 13-5
“Combining Model Objects” on page 13-6
“Example – Using System Identification Toolbox and Control System Toolbox” on page 13-6

How Control System Toolbox Works with Estimated Models

System Identification Toolbox integrates with Control System Toolbox by providing linear plant identification for control-system design. If you have Control System Toolbox, you can use System Identification Toolbox to identify a linear, time-invariant plant system, and then use Control System Toolbox to design a controller for this plant.

Control System Toolbox also provides the LTI Viewer to extend System Identification Toolbox functionality for linear model analysis.

Only linear models are supported in Control System Toolbox. If you identified a nonlinear plant model using System Identification Toolbox, you must linearize it using `linapp` or `lintan` before you can work with this model in Control System Toolbox.

Note You can only use System Identification Toolbox to linearize nonlinear ARX (`idnlarx`) and Hammerstein-Wiener (`idnlhw`) models.

Linearization of nonlinear grey-box (`idnlgrey`) models is not supported.

For information about using Control System Toolbox, see the Control System Toolbox documentation.

Using `balred` to Reduce Model Order

In some cases, your identified model order might be higher than necessary to capture the dynamics. If you have Control System Toolbox, you can use `balred` to compute a state-space model approximation with a reduced model order for any `idmodel` object, including `idarx`, `idpoly`, `idss`, and `idgrey`.

For more information about using `balred`, see the corresponding reference page. To learn how you can reduce model order using pole-zero plots, see “Reducing Model Order Using Pole-Zero Plots” on page 10-48.

Compensator Design Using Control System Toolbox

After you estimate a plant model in System Identification Toolbox, you can use Control System Toolbox to design a controller for this plant.

System Identification Toolbox models in the MATLAB Workspace browser are immediately available to Control System Toolbox commands. However, if you used the System Identification Tool GUI to estimate models, you must first export the models to the MATLAB Workspace Browser. To export a model from the GUI, drag the model icon to the **To Workspace** rectangle.

Control System Toolbox provides both the SISO Design Tool GUI and commands for working at the command line. You can import polynomial and state-space models directly into SISO Design Tool using the following command:

```
sisotool(model('measured'))
```

where you use only the dynamic model and not the noise model. For more information about subreferencing the dynamic or the noise model, see “Subreferencing Measured and Noise Models” on page 3-21. To design a controller using Control System Toolbox commands and methods at the command line, you must convert the plant model to an LTI object. For more information, see “Converting Models to LTI Objects” on page 13-4.

Note The syntax `sisotool(model('m'))` is equivalent to `sisotool(model('measured'))`.

For more information about controller design using SISO Design Tool and Control System Toolbox commands, see the Control System Toolbox documentation.

Converting Models to LTI Objects

Control System Toolbox commands operate on Control System Toolbox LTI objects. To design a controller for a plant model you estimated in System Identification Toolbox, you must first convert the plant model to an LTI object.

You can convert linear polynomial, state-space, and grey-box model objects, including `idarx`, `idpoly`, `idproc`, `idss`, or `idgrey`, to LTI model objects.

The following table summarizes the commands for transforming linear state-space and polynomial models to LTI model object.

Commands for Converting Models to LTI Objects

Command	Description	Example
<code>frd</code>	Convert to frequency-response representation.	<code>ss_sys = frd(model)</code>
<code>ss</code>	Convert to state-space representation.	<code>ss_sys = ss(model)</code>
<code>tf</code>	Convert to transfer-command form.	<code>tf_sys = tf(model)</code>
<code>zpk</code>	Convert to zero-pole form.	<code>zpk_sys = zpk(model)</code>

The following command transforms an `idmodel` object to an LTI state-space object:

```
% Extract the measured model
```



```
% and ignore the noise model
model = model('measured')
% Convert to LTI object
LTI_sys = idss(model)
```

The LTI object includes only the dynamic model and not the noise model, which is estimated for every linear model in System Identification Toolbox.

Note To include noise channels in the LTI models, first use `noisecnv` to convert the noise in the `idmodel` object to measured channels, and then convert to an LTI object.

For more information about subreferencing the dynamic or the noise model, see “Subreferencing Measured and Noise Models” on page 3-21.

Viewing Model Response Using the LTI Viewer

- “What Is the LTI Viewer?” on page 13-5
- “Displaying Identified Models in the LTI Viewer” on page 13-6

What Is the LTI Viewer?

If you have Control System Toolbox, you can plot models in the LTI Viewer from either the System Identification Tool GUI or the MATLAB Command Window.

The LTI Viewer is a graphical user interface for viewing and manipulating the response plots of linear models.

Note The LTI Viewer does not display model uncertainty.

For more information about working with plots in the LTI Viewer, see the Control System Toolbox documentation.

Displaying Identified Models in the LTI Viewer

When Control System Toolbox is available, the System Identification Tool GUI contains the **To LTI Viewer** rectangle. To plot models in the LTI Viewer, drag and drop the corresponding icon to the **To LTI Viewer** rectangle in the System Identification Tool GUI.

Alternatively, use the following syntax when working at the command line to view a model in the LTI Viewer:

```
view(model)
```

Combining Model Objects

If you have Control System Toolbox, you can combine linear model objects, such as `idarx`, `idgrey`, `idpoly`, `idproc`, and `idss` model objects, similar to the way you combine LTI objects.

For example, you can perform the following operations on estimated models:

- `G1+G2`
- `G1*G2`
- `append(G1,G2)`
- `feedback(G1,G2)`

Note These operations lose covariance information.

Example – Using System Identification Toolbox and Control System Toolbox

This example demonstrates how to use both System Identification Toolbox commands and Control System Toolbox commands to create and plot models:

```
% Construct model using Control System Toolbox
m0 = drss(4,3,2)
% Convert model to an idss object
m0 = idss(m0,'NoiseVar',0.1*eye(3))
% Generate input data for simulating output
```

```
u = iddata([], idinput([800 2], 'rbs'));
% Simulate model output using System Identification Toolbox
% with added noise
y = sim(m0,u,'noise')
% Form an input-output iddata object
Data = [y u];
% Estimate state-space model from the generated data
% using System Identification Toolbox command pem
m = pem(Data(1:400))
% Convert the model to a Control System Toolbox transfer command
tf(m)
% Plot model output for model m using System Identification Toolbox
compare(Data(401:800),m)
% Display identified model m in LTI Viewer
view(m)
```

Using Models with Simulink

In this section...

“What You Can Do with Simulink” on page 13-8

“Opening the System Identification Library” on page 13-9

“Working with System Identification Blocks” on page 13-9

“Blocks for Importing Data and Models into Simulink” on page 13-9

“Blocks for Model Estimation” on page 13-10

“Example – Using Simulink to Simulate a Multiexperiment Model” on page 13-11

What You Can Do with Simulink

System Identification Toolbox provides the System Identification block for exchanging information between System Identification Toolbox and the Simulink environment.

You can use the Simulink library to perform the following tasks:

- Import estimated models into a Simulink model and simulate the model with or without noise.

The model you import might be a subsystem or the main model. For example, if you estimated a plant model using System Identification Toolbox, you can import this plant into Simulink for controller design.

- Stream time-domain data sources (iddata objects) into a Simulink model.
- Save data from a simulation in Simulink as a System Identification Toolbox data sink object (iddata objects).
- Estimate linear polynomial and state-space models during simulation for single-output data.

Note Simulink supports linear polynomial, linear state-space, linear grey-box, and nonlinear grey-box System Identification Toolbox model objects (idarx, idpoly, idproc, idss, idgrey, and idnlgrey model objects). If you estimated a nonlinear ARX or Hammerstein-Wiener model in System Identification Toolbox, you must linearize this model using `lintan` or `linapp` before importing it into Simulink.

Opening the System Identification Library

To open the System Identification Library window, type the following command at the MATLAB prompt:

```
slident
```

The System Identification Library contains blocks for exchanging information between Simulink and System Identification Toolbox.

Working with System Identification Blocks

You can exchange data between Simulink and System Identification Toolbox by adding System Identification blocks to your Simulink model and specifying block parameters. For more information about creating and running Simulink models, see the discussion on building a model in the *Simulink User's Guide*.

To add a block from the System Identification Library, drag the block into the model window. To specify block parameters, double-click the block icon to open the Block Parameters dialog box.

You can get help on a specific block by right-clicking the block in the Library window and selecting **Help**.

Blocks for Importing Data and Models into Simulink

System Identification blocks let you import and simulate estimated models in Simulink, stream System Identification Toolbox data to Simulink, and save data from a simulation as a System Identification Toolbox object.

After you add a block to the model, double-click the block to specify block parameters. For an example of bringing data into a Simulink model, see

the tutorial on estimating process models in *Getting Started with System Identification Toolbox*. To get help on a specific block, right-click the block in the Library window, and select **Help**.

Block	Description
Iddata Sink	Saves model input and output signals to the MATLAB Workspace Browser as an iddata object.
Iddata Source	Imports iddata object from the MATLAB Workspace Browser. Input and output ports of the block correspond to input and output channels of the data object. These inputs and outputs provide signals to the next block in a Simulink model.
Idmodel	Imports linear polynomial, state-space, and grey-box models (idarx, idpoly, idproc, idss, and idgrey model objects) into a Simulink model.
Idnlgrey Model	Imports a nonlinear grey-box model (idnlgrey model object) into a Simulink model.

Blocks for Model Estimation

System Identification blocks let you estimate linear polynomial and state-space models during simulation. For information about AR, ARX, ARMAX, Box-Jenkins, and Output-Error models, see “Estimating Polynomial Models” on page 5-21.

After you add a block to the model, double-click the block to specify block parameters. To get help on a specific block, right-click the block in the Library window, and select **Help**.

Block	Description
AR	Estimates an AutoRegressive model as an idpoly model object for time-series data, which has one output and no input. Use system output as the input to this block.
ARX	Estimates an AutoRegressive model with an eXternal input as an idpoly object for input-output data. Use system input and output as inputs to this block.
ARMAX	Estimates an AutoRegressive Moving Average model with an eXternal input as an idpoly object for input-output data. Use system input and output as inputs to this block.
BJ	Estimates a Box-Jenkins model as an idpoly object for input-output data. Use system input and output as inputs to this block.
OE	Estimates an Output-Error model as an idpoly object for input-output data. Use system input and output as inputs to this block.
PEM	Uses a general prediction-error method to estimate any single-input and single output idpoly object, such as ARX, ARMAX, Box-Jenkins, and Output-Error models. Use system input and output as inputs to this block.

Example – Using Simulink to Simulate a Multiexperiment Model

This example demonstrates how to set initial states before simulating a multiexperiment model. For multiexperiment data, Simulink uses the initial states of the last experiment unless you specify otherwise.

Suppose you estimate a three-state model *M* using a merged data set *Z*, which contains data from five experiments—*z1*, *z2*, *z3*, *z4*, and *z5*:

```
Z = merge(z1, z2, z3, z4, z5);  
M = n4sid(Z, 3);
```

When a model uses several data sets, the `initial-states` property stores only the estimated states corresponding to the last data set. In this example, `M.X0` is a vector of length 3 (corresponding to the three states of the model). The values of `M.X0` are the estimated state values corresponding to *z5*.

The following procedure describes how to access the initial states that correspond to *z2* for the simulation, where *z2* is a portion of the estimation data *Z*.

To compare the measured output from experiment *z2* with the simulated output:

- 1 Estimate the initial states using the second experiment as input, that is `Z(z2.u)`, as follows:

```
[E, X0est] = pe(M, getexp(Z, 2))
```

In this case, `getexp(Z, 2)` gets the data in *z2*.

- 2 In Simulink, open the Function Block Parameters dialog box for the `idmodel` block.
- 3 In the **idmodel variable** field, type `M` to specify the estimated model.
- 4 In the **Initial states** field, type `X0est` to specify the estimated initial states.
- 5 Click **OK**.

Run the simulation to compare the measured output `z2.y` to the simulated output.

A

- active
 - data in GUI 1-18
 - model in GUI 1-19
- advice
 - for data 2-60
 - for models 10-6
- AIC 10-58
 - definition 10-59
- Akaike's Final Prediction Error (FPE) 10-58
- Akaike's Information Criterion (AIC) 10-58
- Algorithm property 3-13
- algorithms for estimation
 - recursive 7-6
 - spectral models 4-13
- aliasing effects 2-83
- AR 6-7
- ARMA 6-7
- ARMAX 5-24
- ARX 5-24
- ARX Model Structure Selection window 5-35

B

- best fit
 - definition 10-10
 - negative value 10-10
- BJ model. *See* Box-Jenkins model
- Bode plot 10-31
- Box-Jenkins model 5-24
- Burg's method 6-11

C

- c2d 11-2
- canonical parameterization 5-66
- complex data 2-115
- concatenating
 - iddata objects 2-50
 - idfrd objects 2-56

- models 3-25
- confidence interval
 - impulse response plot 10-26
 - model output plot 10-12
 - noise spectrum plot 10-38
 - residual plot 10-17
 - step response plot 10-26
- confidence interval on plots 10-62
- constructor 3-3
- continuous-time process models 5-2
- Control System Toolbox
 - combining model objects 13-6
 - converting models to LTI objects 13-4
 - for compensator design 13-3
 - LTI Viewer 13-5
 - reducing model order 13-3
- correlation analysis 4-2
- covariance 10-61
- CovarianceMatrix 10-61
- cra 4-4
- cross-validation 10-4
- custom network 8-22

D

- D matrix 5-64
- d2c 11-2
- d2d 11-2
- data
 - creating iddata object 2-32
 - creating idfrd object 2-52
 - creating subsets 2-20
 - detrending 2-75
 - exporting to MATLAB workspace 1-16
 - filter 2-87
 - frequency-domain 2-9
 - frequency-response 2-10
 - importing into System Identification Tool GUI 2-14
 - managing in GUI 1-9

- merging 2-23
 - missing data 2-71
 - multiexperiment data 2-23
 - outliers 2-72
 - plotting 2-62
 - renaming in GUI 1-14 2-30
 - resampling 2-80
 - sampling interval 2-19
 - segmentation 7-14
 - selecting 2-95
 - simulating 2-99
 - supported types 2-6
 - time-domain 2-7
 - time-series 2-8
 - transforming domain 2-103
 - viewing properties in GUI 1-13
- Data Board 1-6
- arranging icons 1-14
 - deleting icons 1-15
- data tip 1-20
- dead time 5-21
- dead zone 8-22
- delay
- estimating for polynomial models 5-29
- detrending data 2-75
- E**
- estimating models
- black-box polynomial 5-21
 - commands 3-5
 - frequency response 4-10
 - Hammerstein-Wiener 8-14
 - linear grey-box 9-5
 - nonlinear ARX 8-4
 - nonlinear grey-box 9-13
 - process models 5-2
 - recursive estimation 7-2
 - state-space 5-48
 - time-series 6-1
- transient response 4-2
 - uncertainty 10-61
- EstimationInfo property 3-14
- etfe
- algorithm 4-13
- export
- data to MATLAB workspace 1-16
 - model to MATLAB workspace 1-16
- F**
- filtering data 2-87
- forgetting factor algorithm 7-10
- FPE 10-58
- free parameterization 5-59
- frequency resolution 4-14
- frequency response
- estimating in the GUI 4-11
 - etfe 4-13
 - spa 4-13
 - spafdr 4-13
- frequency-domain data 2-9
- frequency-response command 4-17
- frequency-response data 2-10
- frequency-response plot 10-29
- Bode plot 10-32
 - Nyquist plot 10-35
- H**
- Hammerstein-Wiener models 8-14
- Hammerstein-Wiener plot 10-53
- I**
- idarx 3-4
- iddata
- concatenating 2-50
 - creating 2-32
 - subreferencing 2-40
- ident 1-4

- idfrd
 - concatenating 2-56
 - creating 2-52
 - model 3-4
 - subreferencing 2-55
- idgrey 3-4
- idlayout.m 1-27
- idnlarx 3-4
- idnlgrey 3-4
- idnlhw 3-4
- idpoly 3-4
- idproc 3-4
- idss 3-4
- importing
 - data into MATLAB 2-13
 - data into System Identification Tool GUI 2-14
- impulse response
 - computing values 4-6
 - confidence interval 10-26
 - definition 4-2
 - estimating in the GUI 4-3
 - impulse 4-4
- impulse-response plot 10-21
- independence test 10-15
- initial states for simulation and prediction 12-8

K

- K matrix 5-64
- Kalman filter algorithm 7-8

L

- linear grey-box models 9-5
- linear models
 - extracting numerical data 10-64
 - transforming between continuous and discrete time 11-2
 - transforming between structures 11-7

- LTI Viewer 13-5

M

- MDL 5-34
- merging
 - data 2-23
 - models 3-29
- methods 3-3
- missing data 2-71
- model
 - black-box polynomial 5-21
 - estimating frequency response 4-10
 - estimating process model 5-2
 - estimating transient response 4-2
 - exporting to MATLAB workspace 1-16
 - gray-box estimation 9-1
 - Hammerstein-Wiener estimation 8-14
 - importing into GUI 1-11
 - linear grey-box estimation 9-5
 - managing in GUI 1-9
 - nonlinear ARX estimation 8-4
 - nonlinear black-box estimation 8-1
 - nonlinear grey-box estimation 9-13
 - plotting 10-5
 - predict output 12-7
 - properties 3-13
 - recursive estimation 7-2
 - reducing order using balred 13-3
 - reducing order using pole-zero plot 10-48
 - refining linear parametric 5-78
 - refining nonlinear black-box 8-25
 - renaming in GUI 1-14
 - state-space 5-48
 - time-series 6-1
 - uncertainty 10-61
 - validating 10-3
 - viewing properties in GUI 1-13
- Model Board 1-6
 - arranging icons 1-14

- deleting icons 1-15
 - model object
 - concatenating 3-25
 - definition 3-2
 - instantiating 3-3
 - merging 3-29
 - methods 3-3
 - properties 3-3 3-13
 - types of 3-4
 - model order
 - definition 5-21
 - estimating for polynomial models 5-29
 - estimating for state-space 5-54
 - too high or too low 10-68
 - Model Order Selection window 5-58
 - model output
 - confidence interval 10-12
 - model output plot 10-7
 - model properties
 - accessing 3-16
 - help on 3-18
 - specifying 3-15
 - multiexperiment data 2-23
- N**
- neural network 8-22
 - noise
 - converting to measured channels 3-23
 - evidence in estimated model 10-70
 - subreferencing 3-21
 - noise spectrum
 - confidence interval 10-38
 - noise spectrum plot 10-37
 - nonlinear ARX models 8-4
 - nonlinear ARX plot 10-49
 - nonlinear grey-box models 9-13
 - nonlinear models 8-1
 - nonlinearities 8-22
 - nonlinearity estimators 8-22
 - troubleshooting 10-69
 - normalized gradient algorithm 7-11
- O**
- OE model. *See* Output-Error model
 - offset levels 2-75
 - order. *See* model order
 - outliers 2-72
 - Output-Error model 5-24
- P**
- pem
 - for polynomial models 5-42
 - for process models 5-9
 - for state-space models 5-62
 - periodogram
 - etfe for time series 6-5
 - physical equilibrium 2-75
 - piece-wise linear 8-22
 - plot
 - copy to MATLAB Figure window 1-25
 - data 2-62
 - data in the GUI 1-18
 - data tip 1-20
 - in LTI Viewer 13-5
 - models 10-5
 - models in the GUI 1-19
 - print 1-25
 - selecting noise channels 1-23
 - pole-zero cancelation 10-48
 - pole-zero plot 10-44
 - polynomial models 5-21
 - estimating order 5-29
 - for time-series 6-7
 - polynomial nonlinearity 8-22
 - predicting model output 12-1
 - initial states 12-8
 - prediction 12-7

- print plot 1-25
- process model 5-2
 - definition 5-2
- properties
 - for models 3-13

R

- recursive estimation 7-2
- reducing model order
 - using balred 13-3
 - using pole-zero plot 10-48
- refining models
 - linear 5-79
 - linear parametric 5-78
 - nonlinear 8-25
 - nonlinear black-box 8-25
 - using pem 5-80 8-26
- renaming data 2-30
- resampling data 2-80
 - avoiding aliasing 2-83
- residual analysis
 - confidence interval 10-17
 - plot 10-15
- Rissanen's Minimum Description Length (MDL) 5-34
- robust criterion
 - for outliers 2-72

S

- sampling interval 2-19
- saving session preferences 1-26
- segmentation of data 7-14
- selecting data 2-95
- session
 - definition 1-3
 - managing in GUI 1-3
 - preferences 1-26
 - starting 1-4

- sigmoid network 8-22
- simulating data 2-99
- simulating model output 12-1
 - initial states 12-8
- Simulink 13-8
- slident 13-9
- spa
 - algorithm 4-13
- spafdr
 - algorithm 4-13
- spectral analysis 4-10
 - algorithm 4-13
 - frequency resolution 4-14
 - spectrum normalization 4-19
- spectrum normalization 4-19
- state-space models 5-48
 - canonical parameterization 5-66
 - estimating order 5-54
 - for time series 6-12
 - free parameterization 5-59
 - structured parameterization 5-68
 - supported parameterization 5-53
- step response
 - computing values 4-6
 - confidence interval 10-26
 - definition 4-2
 - estimating in the GUI 4-3
 - step 4-4
- step-response plot 10-21
- structured parameterization 5-68
- subreferencing
 - iddata objects 2-40
 - idfrd objects 2-55
 - model channels 3-20
 - model noise channels 3-21
 - models 3-20
- System Identification blocks 13-8
 - for data and models 13-9
 - for model estimation 13-10
 - open 13-9

System Identification Tool GUI

- customizing 1-26
- exiting 1-8
- help 1-7
- open 1-4
- organizing icons 1-14
- plots 1-18
- window 1-5
- workflow 1-2

T

- time-domain data 2-7
- time-series data 2-8
- time-series models 6-1
- transforming data domain 2-103
- Trash 1-15
- tree partition 8-22
- troubleshooting models 10-68
 - complicated nonlinearities 10-72
 - high noise content 10-70
 - missing inputs 10-71
 - model order 10-68
 - nonlinearity estimators 10-69
 - unstable models 10-70

U

- uncertainty of models 10-61

- confidence interval on plots 10-62

- covariance 10-61

- unnormalized gradient algorithm 7-11

- unstable models 10-70

V

- validating models 10-3

- comparing model output 10-7

- residual analysis 10-15

- troubleshooting 10-68

- Validation Data 1-11

W

- warnings 1-26

- wavelet network 8-22

- whiteness test 10-15

- Working Data 1-11

X

- X0 matrix 5-64

Y

- Yule-Walker approach 6-11