# MATLAB Machine Learning Recipes

## A Problem-Solution Approach

*Second Edition*

Michael Paluszek
Stephanie Thomas

Apress®

# MATLAB Machine Learning Recipes

## A Problem-Solution Approach

## Second Edition

Michael Paluszek
Stephanie Thomas

Apress®

***MATLAB Machine Learning Recipes: A Problem-Solution Approach***

Michael Paluszek
Plainsboro, NJ
USA

Stephanie Thomas
Plainsboro, NJ
USA

ISBN-13 (pbk): 978-1-4842-3915-5
ISBN-13 (electronic): 978-1-4842-3916-2
https://doi.org/10.1007/978-1-4842-3916-2

Library of Congress Control Number: 2018967208

Copyright © 2019 by Michael Paluszek and Stephanie Thomas

This work is subject to copyright. All rights are reserved by the Publisher, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilms or in any other physical way, and transmission or information storage and retrieval, electronic adaptation, computer software, or by similar or dissimilar methodology now known or hereafter developed.

Trademarked names, logos, and images may appear in this book. Rather than use a trademark symbol with every occurrence of a trademarked name, logo, or image we use the names, logos, and images only in an editorial fashion and to the benefit of the trademark owner, with no intention of infringement of the trademark.

The use in this publication of trade names, trademarks, service marks, and similar terms, even if they are not identified as such, is not to be taken as an expression of opinion as to whether or not they are subject to proprietary rights.

While the advice and information in this book are believed to be true and accurate at the date of publication, neither the authors nor the editors nor the publisher can accept any legal responsibility for any errors or omissions that may be made. The publisher makes no warranty, express or implied, with respect to the material contained herein.

Managing Director, Apress Media LLC: Welmoed Spahr
Acquisitions Editor: Steve Anglin
Development Editor: Matthew Moodie
Coordinating Editor: MarkPowers

Cover designed by eStudioCalamar
Cover image designed by Freepik (http://www.freepik.com)

Distributed to the book trade worldwide by Springer Science+Business Media New York, 233 Spring Street, 6th Floor, New York, NY 10013. Phone 1-800-SPRINGER, fax (201) 348-4505, e-mail orders-ny@springer-sbm.com, or visit www.springeronline.com. Apress Media, LLC is a California LLC and the sole member (owner) is Springer Science + Business Media Finance Inc (SSBM Finance Inc). SSBM Finance Inc is a Delaware corporation.

For information on translations, please e-mail rights@apress.com, or visit www.apress.com.

Apress and friends of ED books may be purchased in bulk for academic, corporate, or promotional use. eBook versions and licenses are also available for most titles. For more information, reference our Special Bulk Sales-eBook Licensing web page at www.apress.com/bulk-sales.

Any source code or other supplementary materials referenced by the author in this text are available to readers at www.apress.com. For detailed information about how to locate your book's source code, go to www.apress.com/source-code/. Readers can also access source code at SpringerLink in the Supplementary Material section for each chapter.

Printed on acid-free paper

# Contents

**About the Authors** XV

**Introduction** XVII

**1 An Overview of Machine Learning** 1

1.1 Introduction 1
1.2 Elements of Machine Learning 2
1.2.1 Data 2
1.2.2 Models 3
1.2.3 Training 3
1.2.3.1 Supervised Learning 3
1.2.3.2 Unsupervised Learning 4
1.2.3.3 Semi-Supervised Learning 4
1.2.3.4 Online Learning 4
1.3 The Learning Machine 4
1.4 Taxonomy of Machine Learning 6
1.5 Control 8
1.5.1 Kalman Filters 8
1.5.2 Adaptive Control 9
1.6 Autonomous Learning Methods 9
1.6.1 Regression 10
1.6.2 Decision Trees 13
1.6.3 Neural Networks 14
1.6.4 Support Vector Machines 15
1.7 Artificial Intelligence 16
1.7.1 What is Artificial Intelligence? 16
1.7.2 Intelligent Cars 16
1.7.3 Expert Systems 17
1.8 Summary 18

**2 Representation of Data for Machine Learning in MATLAB** **19**
2.1 Introduction to MATLAB Data Types . . . . . . . . . . . . . . . . . . . . 19
2.1.1 Matrices . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 19
2.1.2 Cell Arrays . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 20
2.1.3 Data Structures . . . . . . . . . . . . . . . . . . . . . . . . . . . 21
2.1.4 Numerics . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
2.1.5 Images . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 23
2.1.6 Datastore . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 25
2.1.7 Tall Arrays . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 26
2.1.8 Sparse Matrices . . . . . . . . . . . . . . . . . . . . . . . . . . . 27
2.1.9 Tables and Categoricals . . . . . . . . . . . . . . . . . . . . . . . 27
2.1.10 Large MAT-Files . . . . . . . . . . . . . . . . . . . . . . . . . . 29
2.2 Initializing a Data Structure Using Parameters . . . . . . . . . . . . . . . 30
2.2.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 30
2.2.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 30
2.2.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . . . . . 30
2.3 Performing MapReduce on an Image Datastore . . . . . . . . . . . . . . . 33
2.3.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 33
2.3.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 33
2.3.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . . . . . 33
2.4 Creating a Table from a File . . . . . . . . . . . . . . . . . . . . . . . . 35
2.4.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 35
2.4.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 36
2.4.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . . . . . 36
2.5 Processing Table Data . . . . . . . . . . . . . . . . . . . . . . . . . . . 37
2.5.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 37
2.5.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 38
2.5.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . . . . . 38
2.6 Using MATLAB Strings . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.1 String Concatenation . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.1.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.1.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.1.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.2 Arrays of Strings . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.2.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.2.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.2.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . 41
2.6.3 Substrings . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 42
2.6.3.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 42
2.6.3.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 42
2.6.3.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . 42
2.7 Summary . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . . 43

**3 MATLAB Graphics** 45

3.1 2D Line Plots . . . 45
3.1.1 Problem . . . 45
3.1.2 Solution . . . 45
3.1.3 How It Works . . . 46
3.2 General 2D Graphics . . . 48
3.2.1 Problem . . . 48
3.2.2 Solution . . . 48
3.2.3 How It Works . . . 48
3.3 Custom Two-Dimensional Diagrams . . . 50
3.3.1 Problem . . . 50
3.3.2 Solution . . . 50
3.3.3 How It Works . . . 50
3.4 Three-Dimensional Box . . . 51
3.4.1 Problem . . . 52
3.4.2 Solution . . . 52
3.4.3 How It Works . . . 52
3.5 Draw a 3D Object with a Texture . . . 54
3.5.1 Problem . . . 54
3.5.2 Solution . . . 54
3.5.3 How It Works . . . 55
3.6 General 3D Graphics . . . 56
3.6.1 Problem . . . 56
3.6.2 Solution . . . 56
3.6.3 How It Works . . . 56
3.7 Building a GUI . . . 58
3.7.1 Problem . . . 58
3.7.2 Solution . . . 58
3.7.3 How It Works . . . 58
3.8 Animating a Bar Chart . . . 63
3.8.1 Problem . . . 64
3.8.2 Solution . . . 64
3.8.3 How It Works . . . 64
3.9 Drawing a Robot . . . 67
3.9.1 Problem . . . 67
3.9.2 Solution . . . 67
3.9.3 How It Works . . . 67
3.10 Summary . . . 71

**4 Kalman Filters** **73**
4.1 A State Estimator Using a Linear Kalman Filter . . . . . 74
4.1.1 Problem . . . . . 74
4.1.2 Solution . . . . . 75
4.1.3 How It Works . . . . . 75
4.2 Using the Extended Kalman Filter for State Estimation . . . . . 92
4.2.1 Problem . . . . . 92
4.2.2 Solution . . . . . 93
4.2.3 How It Works . . . . . 93
4.3 Using the Unscented Kalman Filter for State Estimation . . . . . 97
4.3.1 Problem . . . . . 97
4.3.2 Solution . . . . . 97
4.3.3 How It Works . . . . . 99
4.4 Using the UKF for Parameter Estimation . . . . . 104
4.4.1 Problem . . . . . 104
4.4.2 Solution . . . . . 104
4.4.3 How It Works . . . . . 104
4.5 Summary . . . . . 108

**5 Adaptive Control** **109**
5.1 Self Tuning: Modeling an Oscillator . . . . . 110
5.2 Self Tuning: Tuning an Oscillator . . . . . 112
5.2.1 Problem . . . . . 112
5.2.2 Solution . . . . . 112
5.2.3 How It Works . . . . . 112
5.3 Implement Model Reference Adaptive Control . . . . . 117
5.3.1 Problem . . . . . 117
5.3.2 Solution . . . . . 117
5.3.3 How It Works . . . . . 117
5.4 Generating a Square Wave Input . . . . . 121
5.4.1 Problem . . . . . 121
5.4.2 Solution . . . . . 121
5.4.3 How It Works . . . . . 121
5.5 Demonstrate MRAC for a Rotor . . . . . 123
5.5.1 Problem . . . . . 123
5.5.2 Solution . . . . . 123
5.5.3 How It Works . . . . . 123
5.6 Ship Steering: Implement Gain Scheduling for Steering Control of a Ship . . . . . 126
5.6.1 Problem . . . . . 126
5.6.2 Solution . . . . . 126
5.6.3 How It Works . . . . . 126

5.7 Spacecraft Pointing . . . . . . . . . . . . . . . . . . . . . . . . . . 130
5.7.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 130
5.7.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 130
5.7.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 131
5.8 Summary . . . . . . . . . . . . . . . . . . . . . . . . . . . . 133

**6 Fuzzy Logic 135**
6.1 Building Fuzzy Logic Systems . . . . . . . . . . . . . . . . . . . . 136
6.1.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 136
6.1.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 136
6.1.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 136
6.2 Implement Fuzzy Logic . . . . . . . . . . . . . . . . . . . . . . . 139
6.2.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 139
6.2.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 139
6.2.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 139
6.3 Demonstrate Fuzzy Logic . . . . . . . . . . . . . . . . . . . . . . 142
6.3.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 142
6.3.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 142
6.3.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 143
6.4 Summary . . . . . . . . . . . . . . . . . . . . . . . . . . . . 146

**7 Data Classification with Decision Trees 147**
7.1 Generate Test Data . . . . . . . . . . . . . . . . . . . . . . . . 148
7.1.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 148
7.1.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 148
7.1.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 148
7.2 Drawing Decision Trees . . . . . . . . . . . . . . . . . . . . . . 151
7.2.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 151
7.2.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 151
7.2.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 151
7.3 Implementation . . . . . . . . . . . . . . . . . . . . . . . . . 155
7.3.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 155
7.3.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 155
7.3.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 155
7.4 Creating a Decision tree . . . . . . . . . . . . . . . . . . . . . . 158
7.4.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 158
7.4.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 158
7.4.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 159
7.5 Creating a Handmade Tree . . . . . . . . . . . . . . . . . . . . . 162
7.5.1 Problem . . . . . . . . . . . . . . . . . . . . . . . . . . 162
7.5.2 Solution . . . . . . . . . . . . . . . . . . . . . . . . . . 162
7.5.3 How It Works . . . . . . . . . . . . . . . . . . . . . . . . 163

7.6 Training and Testing . . . . . 165
7.6.1 Problem . . . . . 165
7.6.2 Solution . . . . . 165
7.6.3 How It Works . . . . . 166
7.7 Summary . . . . . 169

**8 Introduction to Neural Nets** **171**
8.1 Daylight Detector . . . . . 171
8.1.1 Problem . . . . . 171
8.1.2 Solution . . . . . 172
8.1.3 How It Works . . . . . 172
8.2 Modeling a Pendulum . . . . . 173
8.2.1 Problem . . . . . 173
8.2.2 Solution . . . . . 174
8.2.3 How It Works . . . . . 174
8.3 Single Neuron Angle Estimator . . . . . 177
8.3.1 Problem . . . . . 177
8.3.2 Solution . . . . . 177
8.3.3 How It Works . . . . . 178
8.4 Designing a Neural Net for the Pendulum . . . . . 182
8.4.1 Problem . . . . . 182
8.4.2 Solution . . . . . 182
8.4.3 How It Works . . . . . 182
8.5 Summary . . . . . 186

**9 Classification of Numbers Using Neural Networks** **187**
9.1 Generate Test Images with Defects . . . . . 188
9.1.1 Problem . . . . . 188
9.1.2 Solution . . . . . 188
9.1.3 How It Works . . . . . 188
9.2 Create the Neural Net Functions . . . . . 192
9.2.1 Problem . . . . . 192
9.2.2 Solution . . . . . 193
9.2.3 How It Works . . . . . 193
9.3 Train a Network with One Output Node . . . . . 197
9.3.1 Problem . . . . . 197
9.3.2 Solution . . . . . 197
9.3.3 How It Works . . . . . 198
9.4 Testing the Neural Network . . . . . 202
9.4.1 Problem . . . . . 202
9.4.2 Solution . . . . . 202
9.4.3 How It Works . . . . . 203

9.5 Train a Network with Many Outputs . . . 203
9.5.1 Problem . . . 203
9.5.2 Solution . . . 203
9.5.3 How It Works . . . 204
9.6 Summary . . . 207

**10 Pattern Recognition with Deep Learning 209**
10.1 Obtain Data Online for Training a Neural Net . . . 211
10.1.1 Problem . . . 211
10.1.2 Solution . . . 211
10.1.3 How It Works . . . 211
10.2 Generating Training Images of Cats . . . 211
10.2.1 Problem . . . 211
10.2.2 Solution . . . 211
10.2.3 How It Works . . . 212
10.3 Matrix Convolution . . . 215
10.3.1 Problem . . . 215
10.3.2 Solution . . . 215
10.3.3 How It Works . . . 215
10.4 Convolution Layer . . . 217
10.4.1 Problem . . . 217
10.4.2 Solution . . . 217
10.4.3 How It Works . . . 217
10.5 Pooling to Outputs of a Layer . . . 218
10.5.1 Problem . . . 218
10.5.2 Solution . . . 218
10.5.3 How It Works . . . 219
10.6 Fully Connected Layer . . . 220
10.6.1 Problem . . . 220
10.6.2 Solution . . . 220
10.6.3 How It Works . . . 220
10.7 Determining the Probability . . . 222
10.7.1 Problem . . . 222
10.7.2 Solution . . . 222
10.7.3 How It Works . . . 222
10.8 Test the Neural Network . . . 223
10.8.1 Problem . . . 223
10.8.2 Solution . . . 223
10.8.3 How It Works . . . 223

10.9 Recognizing a Number . . . 225
10.9.1 Problem . . . 225
10.9.2 Solution . . . 225
10.9.3 How It Works . . . 226
10.10 Recognizing an Image . . . 228
10.10.1 Problem . . . 228
10.10.2 Solution . . . 228
10.10.3 How It Works . . . 228
10.11 Summary . . . 230

**11 Neural Aircraft Control 231**
11.1 Longitudinal Motion . . . 232
11.1.1 Problem . . . 233
11.1.2 Solution . . . 233
11.1.3 How It Works . . . 233
11.2 Numerically Finding Equilibrium . . . 238
11.2.1 Problem . . . 238
11.2.2 Solution . . . 238
11.2.3 How It Works . . . 239
11.3 Numerical Simulation of the Aircraft . . . 240
11.3.1 Problem . . . 240
11.3.2 Solution . . . 240
11.3.3 How It Works . . . 240
11.4 Activation Function . . . 242
11.4.1 Problem . . . 242
11.4.2 Solution . . . 242
11.4.3 How It Works . . . 242
11.5 Neural Net for Learning Control . . . 243
11.5.1 Problem . . . 243
11.5.2 Solution . . . 243
11.5.3 How It Works . . . 243
11.6 Enumeration of All Sets of Inputs . . . 248
11.6.1 Problem . . . 248
11.6.2 Solution . . . 248
11.6.3 How It Works . . . 248
11.7 Write a Sigma-Pi Neural Net Function . . . 249
11.7.1 Problem . . . 249
11.7.2 Solution . . . 249
11.7.3 How It Works . . . 250
11.8 Implement PID Control . . . 251
11.8.1 Problem . . . 251
11.8.2 Solution . . . 251
11.8.3 How It Works . . . 252

11.9 PID Control of Pitch . . . 256
11.9.1 Problem . . . 256
11.9.2 Solution . . . 256
11.9.3 How It Works . . . 256
11.10 Neural Net for Pitch Dynamics . . . 258
11.10.1 Problem . . . 258
11.10.2 Solution . . . 258
11.10.3 How It Works . . . 258
11.11 Nonlinear Simulation . . . 261
11.11.1 Problem . . . 261
11.11.2 Solution . . . 262
11.11.3 How It Works . . . 262
11.12 Summary . . . 264

**12 Multiple Hypothesis Testing** **265**
12.1 Overview . . . 265
12.2 Theory . . . 267
12.2.1 Introduction . . . 267
12.2.2 Example . . . 269
12.2.3 Algorithm . . . 269
12.2.4 Measurement Assignment and Tracks . . . 270
12.2.5 Hypothesis Formation . . . 271
12.2.6 Track Pruning . . . 272
12.3 Billiard Ball Kalman Filter . . . 274
12.3.1 Problem . . . 274
12.3.2 Solution . . . 274
12.3.3 How It Works . . . 274
12.4 Billiard Ball MHT . . . 280
12.4.1 Problem . . . 280
12.4.2 Solution . . . 280
12.4.3 How It Works . . . 280
12.5 One-Dimensional Motion . . . 285
12.5.1 Problem . . . 285
12.5.2 Solution . . . 285
12.5.3 How It Works . . . 285
12.6 One-Dimensional Motion with Track Association . . . 287
12.6.1 Problem . . . 287
12.6.2 Solution . . . 287
12.6.3 How It Works . . . 287
12.7 Summary . . . 289

**13 Autonomous Driving with Multiple Hypothesis Testing** **291**
13.1 Automobile Dynamics . . . 292
13.1.1 Problem . . . 292
13.1.2 Solution . . . 292
13.1.3 How It Works . . . 292
13.2 Modeling the Automobile Radar . . . 295
13.2.1 Problem . . . 295
13.2.2 Solution . . . 295
13.2.3 How It Works . . . 295
13.3 Automobile Autonomous Passing Control . . . 297
13.3.1 Problem . . . 297
13.3.2 Solution . . . 297
13.3.3 How It Works . . . 298
13.4 Automobile Animation . . . 299
13.4.1 Problem . . . 299
13.4.2 How It Works . . . 299
13.4.3 Solution . . . 299
13.5 Automobile Simulation and the Kalman Filter . . . 303
13.5.1 Problem . . . 303
13.5.2 Solution . . . 303
13.5.3 How It Works . . . 303
13.6 Automobile Target Tracking . . . 306
13.6.1 Problem . . . 306
13.6.2 Solution . . . 306
13.6.3 How It Works . . . 306
13.7 Summary . . . 309

**14 Case-Based Expert Systems** **311**
14.1 Building Expert Systems . . . 312
14.1.1 Problem . . . 312
14.1.2 Solution . . . 312
14.1.3 How It Works . . . 313
14.2 Running an Expert System . . . 313
14.2.1 Problem . . . 313
14.2.2 Solution . . . 313
14.2.3 How It Works . . . 313
14.3 Summary . . . 316

**A A Brief History of Autonomous Learning** **317**
A.1 Introduction . . . 317
A.2 Artificial Intelligence . . . 317

A.3 Learning Control . . . 320
A.4 Machine Learning . . . 322
A.5 The Future . . . 323

**B Software for Machine Learning** **325**
B.1 Autonomous Learning Software . . . 325
B.2 Commercial MATLAB Software . . . 326
B.2.1 MathWorks Products . . . 326
B.2.1.1 Statistics and Machine Learning Toolbox . . . 326
B.2.1.2 Neural Network Toolbox . . . 327
B.2.1.3 Computer Vision System Toolbox . . . 327
B.2.1.4 System Identification Toolbox . . . 327
B.2.1.5 MATLAB for Deep Learning . . . 328
B.2.2 Princeton Satellite Systems Products . . . 328
B.2.2.1 Core Control Toolbox . . . 328
B.2.2.2 Target Tracking . . . 328
B.3 MATLAB Open Source Resources . . . 329
B.3.1 DeepLearnToolbox . . . 329
B.3.2 Deep Neural Network . . . 329
B.3.3 MatConvNet . . . 329
B.4 Non- MATLAB Products for Machine Learning . . . 329
B.4.1 R . . . 330
B.4.2 scikit-learn . . . 330
B.4.3 LIBSVM . . . 330
B.5 Products for Optimization . . . 330
B.5.1 LOQO . . . 331
B.5.2 SNOPT . . . 331
B.5.3 GLPK . . . 331
B.5.4 CVX . . . 332
B.5.5 SeDuMi . . . 332
B.5.6 YALMIP . . . 332
B.6 Products for Expert Systems . . . 332
B.7 MATLAB MEX files . . . 333
B.7.1 Problem . . . 333
B.7.2 Solution . . . 333
B.7.3 How It Works . . . 333

**Bibliography** **337**

**Index** **341**

# About the Authors

**Michael Paluszek** is President of Princeton Satellite Systems, Inc. (PSS) in Plainsboro, New Jersey. Mr. Paluszek founded PSS in 1992 to provide aerospace consulting services. He used MATLAB to develop the control system and simulations for the Indostar-1 geosynchronous communications satellite. This led to the launch of Princeton Satellite Systems first commercial MATLAB toolbox, the Spacecraft Control Toolbox, in 1995. Since then he has developed toolboxes and software packages for aircraft, submarines, robotics, and nuclear fusion propulsion, resulting in Princeton Satellite Systems current extensive product line. He is working with the Princeton Plasma Physics Laboratory on a compact nuclear fusion reactor for energy generation and space propulsion.

Prior to founding PSS, Mr. Paluszek was an engineer at GE Astro Space in East Windsor, NJ. At GE he designed the Global Geospace Science Polar despun platform control system and led the design of the GPS IIR attitude control system, the Inmarsat-3 attitude control systems and the Mars Observer delta-V control system, leveraging MATLAB for control design. Mr. Paluszek also worked on the attitude determination system for the DMSP meteorological satellites. Mr. Paluszek flew communication satellites on over twelve satellite launches, including the GSTAR III recovery, the first transfer of a satellite to an operational orbit using electric thrusters. At Draper Laboratory Mr. Paluszek worked on the Space Shuttle, Space Station and submarine navigation. His Space Station work included designing of Control Moment Gyro based control systems for attitude control.

Mr. Paluszek received his bachelors degree in Electrical Engineering, and master's and engineers degrees in Aeronautics and Astronautics from the Massachusetts Institute of Technology. He is author of numerous papers and has over a dozen U.S. Patents. Mr. Paluszek is the author of "MATLAB Recipes" and "MATLAB Machine Learning" both published by Apress.

**Stephanie Thomas** is Vice President of Princeton Satellite Systems, Inc. in Plainsboro, New Jersey. She received her bachelors and masters degrees in Aeronautics and Astronautics from the Massachusetts Institute of Technology in 1999 and 2001. Ms. Thomas was introduced to the PSS Spacecraft Control Toolbox for MATLAB during a summer internship in 1996 and has been using MATLAB for aerospace analysis ever since. In her nearly 20 years of MATLAB experience, she has developed many software tools including the Solar Sail Module for the Spacecraft Control Toolbox; a proximity satellite operations toolbox for the Air Force; collision monitoring Simulink blocks for the Prisma satellite mission; and launch vehicle analysis tools in MATLAB and Java,. She has developed novel methods for space situation assessment such as a numeric approach to assessing the general rendezvous problem between any two satellites implemented in both MATLAB and C++. Ms. Thomas has contributed to PSS Attitude and Orbit Control textbook, featuring examples using the Spacecraft Control Toolbox, and written many software Users Guides. She has conducted SCT training for engineers from diverse locales such as Australia, Canada, Brazil, and Thailand and has performed MATLAB consulting for NASA, the Air Force, and the European Space Agency. Ms. Thomas is the author of “MATLAB Recipes” and “MATLAB Machine Learning” both published by Apress. In 2016, Ms. Thomas was named a NASA NIAC Fellow for the project “Fusion-Enabled Pluto Orbiter and Lander”.

# Introduction

Machine learning is becoming important in every engineering discipline. For example:

1. Autonomous cars. Machine learning is used in almost every aspect of car control systems.
2. Plasma physicists use machine learning to help guide experiments on fusion reactors. TAE Systems has used it with great success in guiding fusion experiments. The Princeton Plasma Physics Laboratory has used it for the National Spherical Torus Experiment to study a promising candidate for a nuclear fusion power plant.
3. It is used in finance for predicting the stock market.
4. Medical professionals use it for diagnoses.
5. Law enforcement, and others, use it for facial recognition. Several crimes have been solved using facial recognition!
6. An expert system was used on NASA's Deep Space 1 spacecraft.
7. Adaptive control systems steer oil tankers.

There are many, many other examples.

Although many excellent packages are available from commercial sources and open-source repositories, it is valuable to understand how these algorithms work. Writing your own algorithms is valuable both because it gives you an insight into the commercial and open-source packages and because it gives you the background to write your own custom machine learning software specialized for your application.

MATLAB® had its origins for that very reason. Scientists who needed to do operations on matrices used numerical software written in FORTRAN. At the time, using computer languages required the user to go through the write-compile-link-execute process, which was time-consuming and error-prone. MATLAB presented the user with a scripting language that allowed the user to solve many problems with a few lines of a script that executed instantaneously. MATLAB has built-in visualization tools that helped the user to better understand the results. Writing MATLAB was a lot more productive and fun than writing FORTRAN.

The goal of *MATLAB Machine Learning Recipes: A Problem–Solution Approach* is to help all users to harness the power of MATLAB to solve a wide range of learning problems. The book has something for everyone interested in machine learning. It also has material that will allow people with an interest in other technology areas to see how machine learning, and MATLAB, can help them to solve problems in their areas of expertise.

## Using the Included Software

This textbook includes a MATLAB toolbox, which implements the examples. The toolbox consists of:

1. MATLAB functions
2. MATLAB scripts
3. html help

The MATLAB scripts implement all of the examples in this book. The functions encapsulate the algorithms. Many functions have built-in demos. Just type the function name in the command window and it will execute the demo. The demo is usually encapsulated in a sub-function. You can copy out this code for your own demos and paste it into a script. For example, type the function name `PlotSet` into the command window and the plot in Figure 1 will appear.

```
>> PlotSet
```

***Figure 1:*** Example plot from the function PlotSet.m.

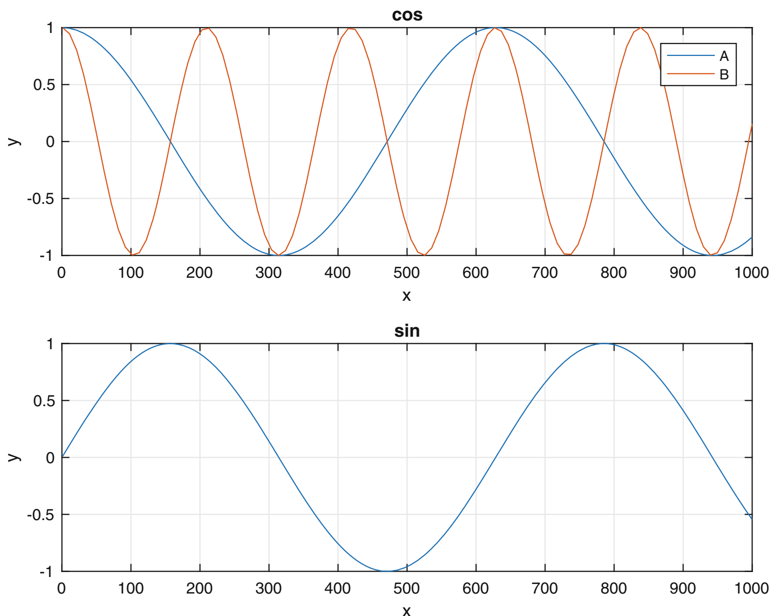

If you open the function you will see the demo:

```
%%% PlotSet>Demo
function Demo

x = linspace(1,1000);
y = [sin(0.01*x);cos(0.01*x);cos(0.03*x)];
disp('PlotSet:␣One␣x␣and␣two␣y␣rows')
PlotSet( x, y, 'figure␣title', 'PlotSet␣Demo',...
    'plot␣set',{[2 3], 1},'legend',{{'A' 'B'},{}},'plot␣title',
    {'cos','sin'});
```

You can use these demos to start your own scripts. Some functions, such as right-hand side functions for numerical integration, don't have demos. If you type:

```
>> RHSAutomobileXY
Error using RHSAutomobileXY (line 17)
a built-in demo is not available.
```

The toolbox is organized according to the chapters in this book. The folder names are Chapter_01, Chapter_02, etc. In addition, there is a general folder with functions that support the rest of the toolbox. You will also need the open-source package GLPK (GNU Linear Programming Kit) to run some of the code. Nicolo Giorgetti has written a MATLAB MEX interface to GLPK that is available on SourceForge and included with this toolbox. The interface consists of:

1. glpk.m
2. glpkcc.mexmaci64, or glpkcc.mexw64, etc.
3. GLPKTest.m

which are available from https://sourceforge.net/projects/glpkmex/. The second item is the MEX file of glpkcc.cpp compiled for your machine, such as Mac or Windows. Go to https://www.gnu.org/software/glpk/ to get the GLPK library and install it on your system. If needed, download the GLPKMEX source code as well and compile it for your machine, or else try another of the available compiled builds.

# CHAPTER 1

# An Overview of Machine Learning

## 1.1 Introduction

Machine learning is a field in computer science where data are used to predict, or respond to, future data. It is closely related to the fields of pattern recognition, computational statistics, and artificial intelligence. The data may be historical or updated in real-time. Machine learning is important in areas such as facial recognition, spam filtering, and other areas where it is not feasible, or even possible, to write algorithms to perform a task.

For example, early attempts at filtering junk emails had the user write rules to determine what was junk or spam. Your success depended on your ability to correctly identify the attributes of the message that would categorize an email as junk, such as a sender address or words in the subject, and the time you were willing to spend on tweaking your rules. This was only moderately successful as junk mail generators had little difficulty anticipating people's hand-made rules. Modern systems use machine-learning techniques with much greater success. Most of us are now familiar with the concept of simply marking a given message as "junk" or "not junk," and take for granted that the email system can quickly learn which features of these emails identify them as junk and prevent them from appearing in our inbox. This could now be any combination of IP or email addresses and words and phrases in the subject or body of the email, with a variety of matching criteria. Note how the machine learning in this example is data-driven, autonomous, and continuously updating itself as you receive email and flag it. However, even today, these systems are not completely successful since they do yet not understand the "meaning" of the text that they are processing.

In a more general sense, what does machine learning mean? Machine learning can mean using machines (computers and software) to gain meaning from data. It can also mean giving machines the ability to learn from their environment. Machines have been used to assist humans for thousands of years. Consider a simple lever, which can be fashioned using a rock and a length of wood, or the inclined plane. Both of these machines perform useful work and assist people but neither has the ability to learn. Both are limited by how they are built. Once built, they cannot adapt to changing needs without human interaction. Figure 1.1 shows early machines that do not learn.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_1

***Figure 1.1:*** Simple machines that do not have the capability to learn.

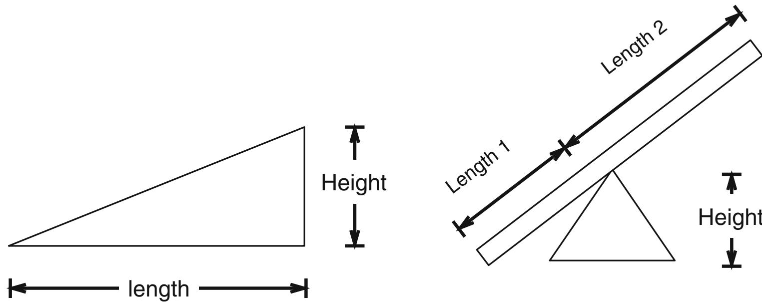


Both of these machines do useful work and amplify the capabilities of people. The knowledge is inherent in their parameters, which are just the dimensions. The function of the inclined plane is determined by its length and height. The function of the lever is determined by the two lengths and the height. The dimensions are chosen by the designer, essentially building in the designer's knowledge of the application and physics.

Machine learning involves memory that can be changed while the machine operates. In the case of the two simple machines described above, knowledge is implanted in them by their design. In a sense, they embody the ideas of the builder, and are thus a form of fixed memory. Learning versions of these machines would automatically change the dimensions after evaluating how well the machines were working. As the loads moved or changed the machines would adapt. A modern crane is an example of a machine that adapts to changing loads, albeit at the direction of a human being. The length of the crane can be changed depending on the needs of the operator.

In the context of the software we will be writing in this book, *machine learning* refers to the process by which an algorithm converts the input data into parameters it can use when interpreting future data. Many of the processes used to mechanize this learning derive from optimization techniques, and in turn are related to the classic field of automatic control. In the remainder of this chapter, we will introduce the nomenclature and taxonomy of machine learning systems.

## 1.2 Elements of Machine Learning

This section introduces key nomenclature for the field of machine learning.

### 1.2.1 Data

All learning methods are data driven. Sets of data are used to train the system. These sets may be collected and edited by humans or gathered autonomously by other software tools. Control systems may collect data from sensors as the systems operate and use that data to identify parameters, or train, the system. The data sets may be very large, and it is the explosion of data storage infrastructure and available databases that is largely driving the growth in machine learning software today. It is still true that a machine learning tool is only as good as the data used to create it, and the selection of training data is practically a field unto itself.

■ **Note** When collecting data from training, one must be careful to ensure that the time variation of the system is understood. If the structure of a system changes with time it may be necessary to discard old data before training the system. In automatic control, this is sometimes called a forgetting factor in an estimator.

### 1.2.2 Models

Models are often used in learning systems. A model provides a mathematical framework for learning. A model is human-derived and based on human observations and experiences. For example, a model of a car, seen from above, might show that it is of rectangular shape with dimensions that fit within a standard parking spot. Models are usually thought of as human-derived and providing a framework for machine learning. However, some forms of machine learning develop their own models without a human-derived structure.

### 1.2.3 Training

A system, which maps an input to an output, needs training to do this in a useful way. Just as people need to be trained to perform tasks, machine learning systems need to be trained. Training is accomplished by giving the system and input and the corresponding output and modifying the structure (models or data) in the learning machine so that mapping is learned. In some ways, this is like curve fitting or regression. If we have enough training pairs, then the system should be able to produce correct outputs when new inputs are introduced. For example, if we give a face recognition system thousands of cat images and tell it that those are cats we hope that when it is given new cat images it will also recognize them as cats. Problems can arise when you don't give it enough training sets or the training data are not sufficiently diverse, for instance, identifying a long-haired cat or hairless cat when the training data only consist of shorthaired cats. Diversity of training data is required for a functioning neural net.

#### 1.2.3.1 Supervised Learning

Supervised learning means that specific training sets of data are applied to the system. The learning is supervised in that the "training sets" are human-derived. It does not necessarily mean that humans are actively validating the results. The process of classifying the system's outputs for a given set of inputs is called "labeling," that is, you explicitly say which results are correct or which outputs are expected for each set of inputs.

The process of generating training sets can be time consuming. Great care must be taken to ensure that the training sets will provide sufficient training so that when real-world data are collected, the system will produce the correct results. They must cover the full range of expected inputs and desired outputs. The training is followed by test sets to validate the results. If the results aren't good then the test sets are cycled into the training sets and the process repeated.

A human example would be a ballet dancer trained exclusively in classical ballet technique. If she were then asked to dance a modern dance, the results might not be as good as required

because the dancer did not have the appropriate training sets; her training sets were not sufficiently diverse.

#### 1.2.3.2 Unsupervised Learning

Unsupervised learning does not utilize training sets. It is often used to discover patterns in data for which there is no "right" answer. For example, if you used unsupervised learning to train a face identification system the system might cluster the data in sets, some of which might be faces. Clustering algorithms are generally examples of unsupervised learning. The advantage of unsupervised learning is that you can learn things about the data that you might not know in advance. It is a way of finding hidden structures in data.

#### 1.2.3.3 Semi-Supervised Learning

With this approach, some of the data are in the form of labeled training sets and other data are not [11]. In fact, typically only a small amount of the input data is labeled while most are not, as the labeling may be an intensive process requiring a skilled human. The small set of labeled data is leveraged to interpret the unlabeled data.

#### 1.2.3.4 Online Learning

The system is continually updated with new data [11]. This is called "online" because many of the learning systems use data collected online. It could also be called recursive learning. It can be beneficial to periodically "batch" process data used up to a given time and then return to the online learning mode. The spam filtering systems from the introduction utilize online learning.

## 1.3 The Learning Machine

Figure 1.2 shows the concept of a learning machine. The machine absorbs information from the environment and adapts. The inputs may be separated into those that produce an immediate response and those that lead to learning. In some cases they are completely separate. For example, in an aircraft a measurement of altitude is not usually used directly for control. Instead, it is used to help select parameters for the actual control laws. The data required for learning and regular operation may be the same, but in some cases separate measurements or data are needed for learning to take place. Measurements do not necessarily mean data collected by a sensor such as radar or a camera. It could be data collected by polls, stock market prices, data in accounting ledgers or any other means. The machine learning is then the process by which the measurements are transformed into parameters for future operation.

Note that the machine produces output in the form of actions. A copy of the actions may be passed to the learning system so that it can separate the effects of the machine actions from those of the environment. This is akin to a feedforward control system, which can result in improved performance.

A few examples will clarify the diagram. We will discuss a medical example, a security system, and spacecraft maneuvering.

A doctor may want to diagnose diseases more quickly. She would collect data on tests on patients and then collate the results. Patient data may include age, height, weight, historical data such as blood pressure readings and medications prescribed, and exhibited symptoms. The

***Figure 1.2:*** A learning machine that senses the environment and stores data in memory.

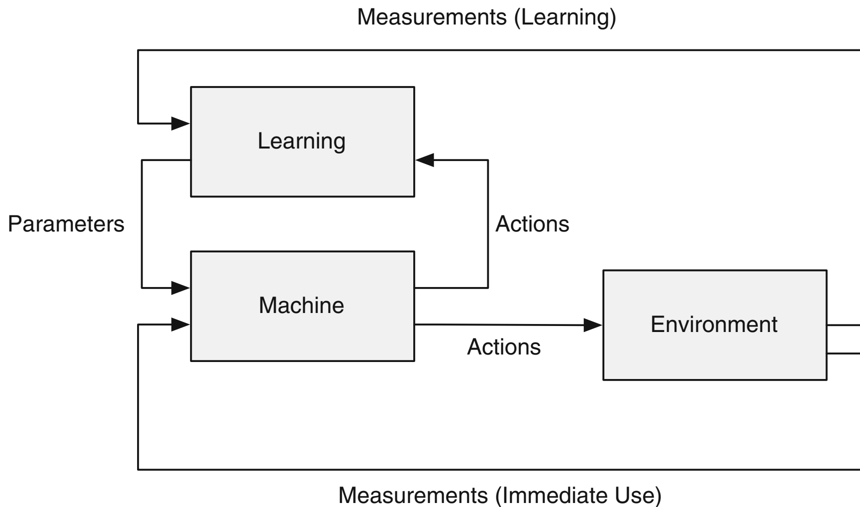


machine learning algorithm would detect patterns so that when new tests were performed on a patient, the machine learning algorithm would be able to suggest diagnoses, or additional tests to narrow down the possibilities. As the machine-learning algorithm were used it would, hopefully, get better with each success or failure. Of course, the definition of success or failure is fuzzy. In this case, the environment would be the patients themselves. The machine would use the data to generate actions, which would be new diagnoses. This system could be built in two ways. In the supervised learning process, test data and known correct diagnoses are used to train the machine. In an unsupervised learning process, the data would be used to generate patterns that may not have been known before and these could lead to diagnosing conditions that would normally not be associated with those symptoms.

A security system may be put into place to identify faces. The measurements are camera images of people. The system would be trained with a wide range of face images taken from multiple angles. The system would then be tested with these known persons and its success rate validated. Those that are in the database memory should be readily identified and those that are not should be flagged as unknown. If the success rate were not acceptable, more training might be needed or the algorithm itself might need to be tuned. This type of face recognition is now common, used in Mac OS X's "Faces" feature in Photos, face identification on the new iPhone X, and Facebook when "tagging" friends in photos.

For precision maneuvering of a spacecraft, the inertia of the spacecraft needs to be known. If the spacecraft has an inertial measurement unit that can measure angular rates, the inertia matrix can be identified. This is where machine learning is tricky. The torque applied to the spacecraft, whether by thrusters or momentum exchange devices, is only known to a certain degree of accuracy. Thus, the system identification must sort out, if it can, the torque scaling factor from the inertia. The inertia can only be identified if torques are applied. This leads to the issue of stimulation. A learning system cannot learn if the system to be studied does not

have known inputs and those inputs must be sufficiently diverse to stimulate the system so that the learning can be accomplished. Training a face recognition system with one picture will not work.

## 1.4 Taxonomy of Machine Learning

In this book, we take a bigger view of machine learning than is typical. Machine learning as described above is the collecting of data, finding patterns, and doing useful things based on those patterns. We expand machine learning to include adaptive and learning control. These fields started off independently, but are now adapting technology and methods from machine learning. Figure 1.3 shows how we organize the technology of machine learning into a consistent taxonomy. You will notice that we created a title that encompasses three branches of learning; we call the whole subject area "Autonomous Learning." That means, learning without human intervention during the learning process. This book is not solely about "traditional" machine learning. There are other, more specialized books that focus on any one of the machine-learning topics. Optimization is part of the taxonomy because the results of optimization can be new discoveries, such as a new type of spacecraft or aircraft trajectory. Optimization is also often a part of learning systems.

***Figure 1.3:*** Taxonomy of machine learning.

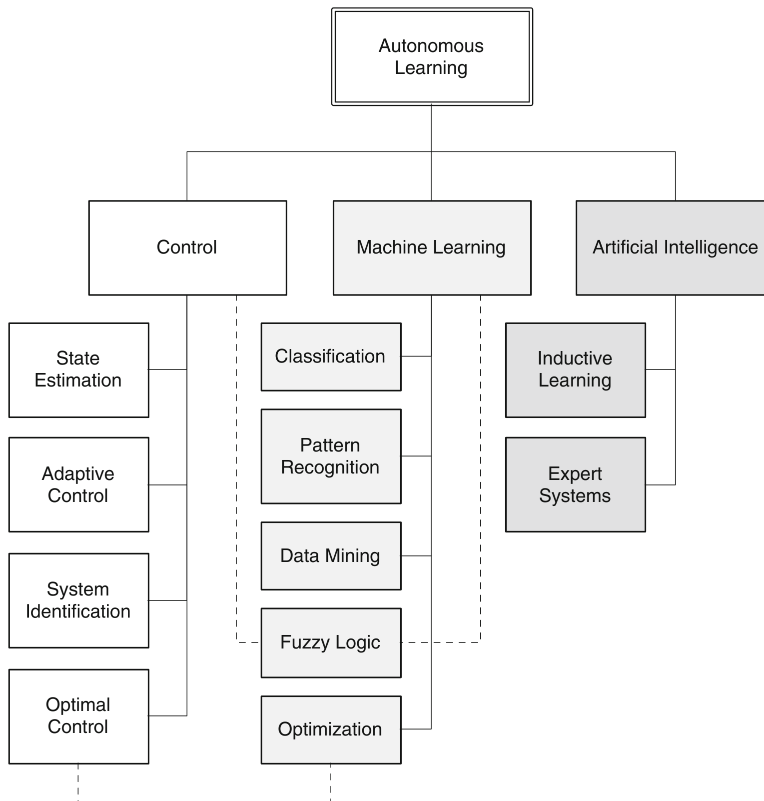

There are three categories under Autonomous Learning. The first is *Control*. Feedback control is used to compensate for uncertainty in a system or to make a system behave differently than it would normally behave. If there were no uncertainty you wouldn't need feedback. For example, if you are a quarterback throwing a football at a running player, assume for a moment and you know everything about the upcoming play. You know exactly where the player should be at a given time, so you can close your eyes, count, and just throw the ball to that spot. Assuming that the player has good hands, you would have a 100% reception rate! More realistically, you watch the player, estimate the player's speed and throw the ball. You are applying feedback to the problem. As stated, this is not a learning system. However, if now you practice the same play repeatedly, look at your success rate and modify the mechanics and timing of your throw using that information, you would have an adaptive control system, the box second from the top of the control list. Learning in control takes place in adaptive control systems and also in the general area of system identification.

System identification is learning about a system. By system we mean the data that represent the system and the relationships between elements of those data. For example, a particle moving in a straight line is a system defined by its mass, the force on that mass, its velocity and position. The position is related to the velocity times time and the velocity is related determined by the acceleration, which is the force divided by the mass.

Optimal control may not involve any learning. For example, what is known as full state feedback produces an optimal control signal, but does not involve learning. In full state feedback, the combination of model and data tells us everything we need to know about the system. However, in more complex systems we can't measure all the states and don't know the parameters perfectly so some form of learning is needed to produce "optimal" or the best possible results.

The second category is what many people consider true *Machine Learning*. This is making use of data to produce behavior that solves problems. Much of its background comes from statistics and optimization. The learning process may be done once in a batch process or continually in a recursive process. For example, in a stock-buying package, a developer may have processed stock data for several years, say prior to 2008, and used that to decide which stocks to buy. That software may not have worked well during the financial crash. A recursive program would continuously incorporate new data. Pattern recognition and data mining fall into this category. Pattern recognition is looking for patterns in images. For example, the early AI Blocks World software could identify a block in its field of view. It could find one block in a pile of blocks. Data mining is taking large amounts of data and looking for patterns, for example, taking stock market data and identifying companies that have strong growth potential. Classification techniques and fuzzy logic are also in this category.

The third category of autonomous learning is *Artificial Intelligence*. Machine learning traces some of its origins to artificial intelligence. Artificial Intelligence is the area of study whose goal is to make machines reason. Although many would say the goal is to "think like people," this is not necessarily the case. There may be ways of reasoning that are not similar to human reasoning, but are just as valid. In the classic Turing test, Turing proposes that the computer only needs to imitate a human in its output to be a "thinking machine," regardless of how

those outputs are generated. In any case, intelligence generally involves learning, so learning is inherent in many Artificial Intelligence technologies such as inductive learning and expert systems. Our diagram includes the two techniques of inductive learning and expert systems.

The recipe chapters of this book are grouped according to this taxonomy. The first chapters cover state estimation using the Kalman Filter and adaptive control. Fuzzy logic is then introduced, which is a control methodology that uses classification. Additional machine-learning recipes follow with chapters on data classification with binary trees, neural nets including deep learning, and multiple hypothesis testing. We then have a chapter on aircraft control that incorporates neural nets, showing the synergy between the different technologies. Finally, we conclude with a chapter on an artificial intelligence technique, case-based expert systems.

## 1.5 Control

Feedback control algorithms inherently learn about the environment through measurements used for control. These chapters show how control algorithms can be extended to effectively design themselves using measurements. The measurements may be the same as used for control, but the adaptation, or learning, happens more slowly than the control response time. An important aspect of control design is stability. A stable controller will produce bounded outputs for bounded inputs. It will also produce smooth, predictable behavior of the system that is controlled. An unstable controller will typically experience growing oscillations in the quantities (such as speed or position) that is controlled. In these chapters, we explore both the performance of learning control and the stability of such controllers. We often break control into two parts, control and estimation. The latter may be done independent of feedback control.

### 1.5.1 Kalman Filters

Chapter 4 shows how Kalman filters allow you to learn about dynamical systems for which we already have a model. This chapter provides an example of a variable gain Kalman Filter for a spring system, that is, a system with a mass connected to its base via a spring and a damper. This is a linear system. We write the system in discrete time. This provides an introduction to Kalman Filtering. We show how Kalman Filters can be derived from Bayesian Statistics. This ties it into many machine-learning algorithms. Originally, the Kalman Filter, developed by R. E. Kalman, C. Bucy, and R. Battin, was not derived in this fashion.

The second recipe adds a nonlinear measurement. A linear measurement is a measurement proportional to the state (in this case position) it measures. Our nonlinear measurement will be the angle of a tracking device that points at the mass from a distance from the line of movement. One way is to use an Unscented Kalman Filter (UKF) for state estimation. The UKF lets us use a nonlinear measurement model easily.

The last part of the chapter describes the Unscented Kalman Filter configured for parameter estimation. This system learns the model, albeit one that has an existing mathematical model. As such, it is an example of model-based learning. In this example, the filter estimates the oscillation frequency of the spring-mass system. It will demonstrate how the system needs to be stimulated to identify the parameters.

### 1.5.2 Adaptive Control

Adaptive control is a branch of control systems in which the gains of the control system change based on measurements of the system. A gain is a number that multiplies a measurement from a sensor to produce a control action such as driving a motor or other actuator. In a nonlearning control system, the gains are computed prior to operation and remain fixed. This works very well most of the time since we can usually pick gains so that the control system is tolerant of parameter changes in the system. Our gain "margins" tell us how tolerant we are to uncertainties in the system. If we are tolerant to big changes in parameters, we say that our system is robust.

Adaptive control systems change the gain based on measurements during operation. This can help a control system perform even better. The better we know a system's model, the tighter we can control the system. This is much like driving a new car. At first, you have to be cautious driving a new car, because you don't know how sensitive the steering is to turning the wheel or how fast it accelerates when you depress the gas pedal. As you learn about the car you can maneuver it with more confidence. If you didn't learn about the car you would need to drive every car in the same fashion.

Chapter 5 starts with a simple example of adding damping to a spring using a control system. Our goal is to get a specific damping time constant. For this, we need to know the spring constant. Our learning system uses a Fast Fourier Transform to measure the spring constant. We'll compare it with a system that does know the spring constant. This is an example of tuning a control system. The second example is model reference adaptive control of a first-order system. This system automatically adapts so that the system behaves like the desired model. This is a very powerful method and applicable to many situations. An additional example will be ship steering control. Ships use adaptive control because it is more efficient than conventional control. This example demonstrates how the control system adapts and how it performs better than its non-adaptive equivalent. This is an example of gain scheduling. We then give a spacecraft example.

The last example is longitudinal control of an aircraft, extensive enough that it is given its own chapter. We can control pitch angle using the elevators. We have five nonlinear equations for the pitch rotational dynamics, velocity in the x direction, velocity in the z direction, and change in altitude. The system adapts to changes in velocity and altitude. Both change the drag and lift forces and the moments on the aircraft and also change the response to the elevators. We use a neural net as the learning element of our control system. This is a practical problem applicable to all types of aircraft ranging from drones to high-performance commercial aircraft.

## 1.6 Autonomous Learning Methods

This section introduces you to popular machine-learning techniques. Some will be used in the examples in this book. Others are available in MATLAB products and open-source products.

### 1.6.1 Regression

Regression is a way of fitting data to a model. A model can be a curve in multiple dimensions. The regression process fits the data to the curve producing a model that can be used to predict future data. Some methods, such as linear regression or least squares, are parametric in that the number of parameters to be fit is known. An example of linear regression is shown in the listing below and Figure 1.4. This model was created by starting with the line $y = x$ and adding noise to $y$. The line was recreated using a least squares fit via MATLAB's `pinv` Pseudo-inverse function.

The first part of the script generates the data.

***Listing 1.1:*** Linear Regression to Data Generation

```
x      = linspace(0,1,500)';
n      = length(x);

% Model a polynomial, y = ax2 + mx + b
a      = 1.0;      % quadratic - make nonzero for larger errors
m      = 1.0;      % slope
b      = 1.0;      % intercept
sigma = 0.1; % standard deviation of the noise
y0     = a*x.^2 + m*x + b;
y      = y0 + sigma*randn(n,1);
```

The actual regression code is just three lines.

***Listing 1.2:*** Linear Regression

```
a      = [x ones(n,1)];
c      = pinv(a)*y;
yR     = c(1)*x + c(2); % the fitted line
```

The last part plots the results using standard MATLAB plotting functions. We use `grid on` rather than `grid`. The latter toggles the grid mode and is usually ok, but sometimes MATLAB gets confused. `grid on` is more reliable.

***Listing 1.3:*** Linear Regression to Plots

```
h = figure;
h.Name = 'Linear Regression';
plot(x,y); hold on;
plot(x,yR,'linewidth',2);
grid on
xlabel('x');
ylabel('y');
title('Linear Regression');
legend('Data','Fit')

figure('Name','Regression Error')
plot(x,yR-y0);
```

```
grid on
xlabel('x');
ylabel('\Delta␣y');
title('Error␣between␣Model␣and␣Regression')
```

This code uses `pinv`. We can solve the problem

$$Ax = b \tag{1.1}$$

by taking the inverse of $A$ if the length of $x$ and $b$ are the same.

$$x = A^{-1}b \tag{1.2}$$

This works because $A$ is a square matrix, but only works if $A$ is not singular, that is, it has a valid inverse. If the length of $x$ and $b$ are the same, we can still find an approximation to $x$ where $x$ = pinv($A$)$b$. For example, in the first case $A$ is 2-by-2. In the second case, it is 3-by-2, meaning there are three elements of $x$ and two of $b$.

```
>> inv(rand(2,2))

ans =

    1.4518   -0.2018
   -1.4398    1.2950

>> pinv(rand(2,3))

ans =

    1.5520   -1.3459
   -0.6390    1.0277
    0.2053    0.5899
```

The system learns the parameters, slope, and $y$ intercept, from the data. The more data, the better the fit. As it happens, our model:

$$y = mx + b \tag{1.3}$$

is correct. However, if it were wrong, the fit would be poor. This is an issue with model-based learning. The quality of the results is highly dependent on the model. If you are sure of your model then it should be used. If not, other methods, such as unsupervised learning, may produce better results. For example, if we add the quadratic term $x^2$, we get the fit in Figure 1.5. Notice how the fit is not as good as we might like.

In these examples, we start with a pattern that we assume fits the data. This is our model. We fit the data to the model. In the first case, we assume that our system is linear; in the second quadratic. If our model is good, the data will fit well. If we choose the wrong model, then the fit

***Figure 1.4:*** Learning with linear regression.

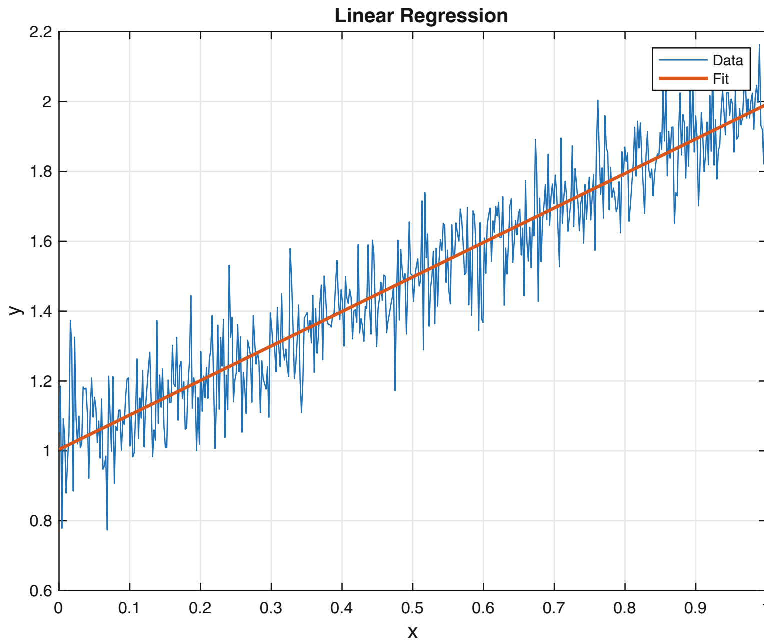


***Figure 1.5:*** Learning with linear regression for a quadratic.

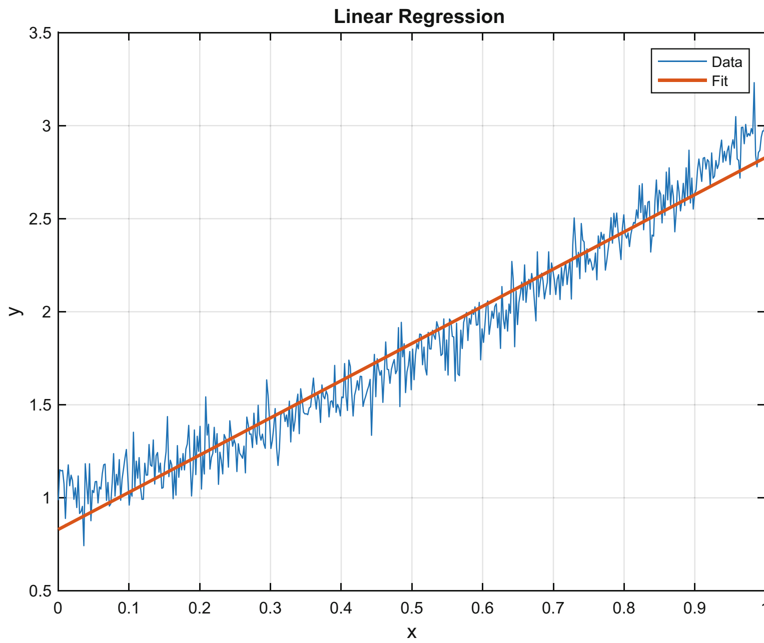

will be poor. If that is the case, we will need to try a different model. For example, our system could be

$$y = \cos(x) \tag{1.4}$$

with the span of $x$ over several cycles. Neither a linear nor a quadratic fit would be good in this case. Limitations in this approach have led to other techniques, including neural networks.

### 1.6.2 Decision Trees

A decision tree is a tree-like graph used to make decisions. It has three kinds of nodes:

1. Decision nodes
2. Chance nodes
3. End nodes

You follow the path from the beginning to the end node. Decision trees are easy to understand and interpret. The decision process is entirely transparent, although very large decision trees may be hard to follow visually. The difficulty is finding an optimal decision tree for a set of training data.

Two types of decision trees are classification trees that produce categorical outputs and regression trees that produce numeric outputs. An example of a classification tree is shown in Figure 1.6. This helps an employee decide where to go for lunch. This tree only has decision nodes.

***Figure 1.6:*** A classification tree.

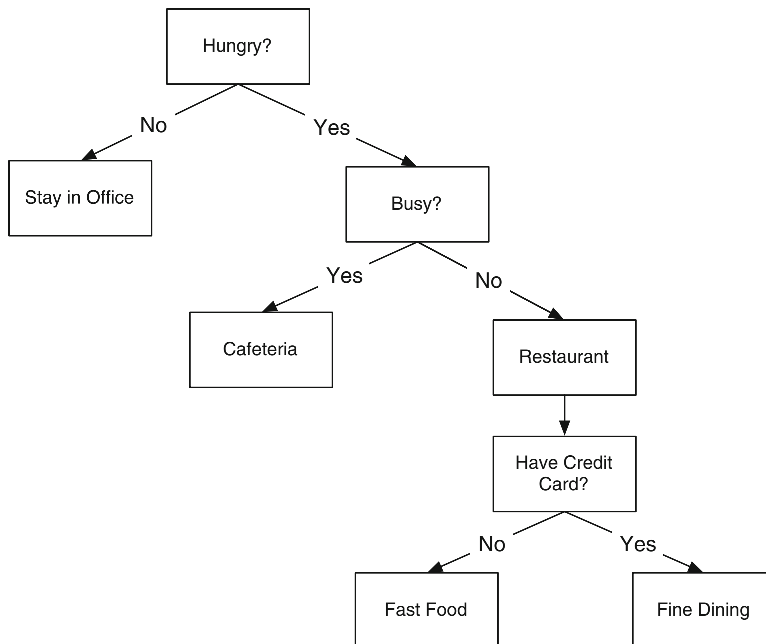

This may be used by management to predict where they could find an employee at lunch time. The decisions are Hungry, Busy, and Have a Credit Card. From that, the tree could be synthesized. However, if there were other factors in the decision of employees, for example, is it someone's birthday, this would result in the employee going to a restaurant, then the tree would not be accurate.

Chapter 7 uses a decision tree to classify data. Classifying data is one of the most widely used areas of machine learning. In this example, we assume that two data points are sufficient to classify a sample and determine to which group it belongs. We have a training set of known data points with membership in one of three groups. We then use a decision tree to classify the data. We'll introduce a graphical display to make understanding the process easier.

With any learning algorithm it is important to know why the algorithm made its decision. Graphics can help you explore large data sets when columns of numbers aren't terribly helpful.

### 1.6.3 Neural Networks

A neural net is a network designed to emulate the neurons in a human brain. Each "neuron" has a mathematical model for determining its output from its input; for example, if the output is a step function with a value of 0 or 1, the neuron can be said to be "firing" if the input stimulus results in a 1 output. Networks are then formed with multiple layers of interconnected neurons. Neural networks are a form of pattern recognition. The network must be trained using sample data, but no a priori model is required. However, usually, the structure of the neural network is specified by giving the number of layers, neurons per layer, and activation functions for each neuron. Networks can be trained to estimate the output of nonlinear processes and the network then becomes the model.

Figure 1.7 displays a simple neural network that flows from left to right, with two input nodes and one output node. There is one "hidden" layer of neurons in the middle. Each node has a set of numeric weights that is tuned during training. This network has two inputs and one output, possibly indicative of a network that solves a categorization problem. Training such a network is called deep learning.

A "deep" neural network is a neural network with multiple intermediate layers between the input and output.

This book presents neural nets in several chapters. Chapter 8 provides an introduction to the fundamentals of neural networks focusing on the neuron and how it can be trained. Chapter 9 provides an introduction to neural networks using a multi-layer feed-forward (MLFF) neural network to classify digits. In this type of network, each neuron depends only on the inputs it receives from the previous layer. The example uses a neural network to classify digits. We will start with a set of six digits and create a training set by adding noise to the digit images. We then see how well our learning network performs at identifying a single digit, and then add more nodes and outputs to identify multiple digits with one network. Classifying digits is one of the oldest uses of machine learning. The U.S. Post Office introduced zip code reading years before machine learning started hitting the front pages of all the newspapers! Earlier digit readers required block letters written in well-defined spots on a form. Reading digits off any envelope is an example of learning in an unstructured environment.

***Figure 1.7:*** A neural net with one intermediate layer between the inputs on the left and the output on the right. The intermediate layer is also known as a hidden layer.

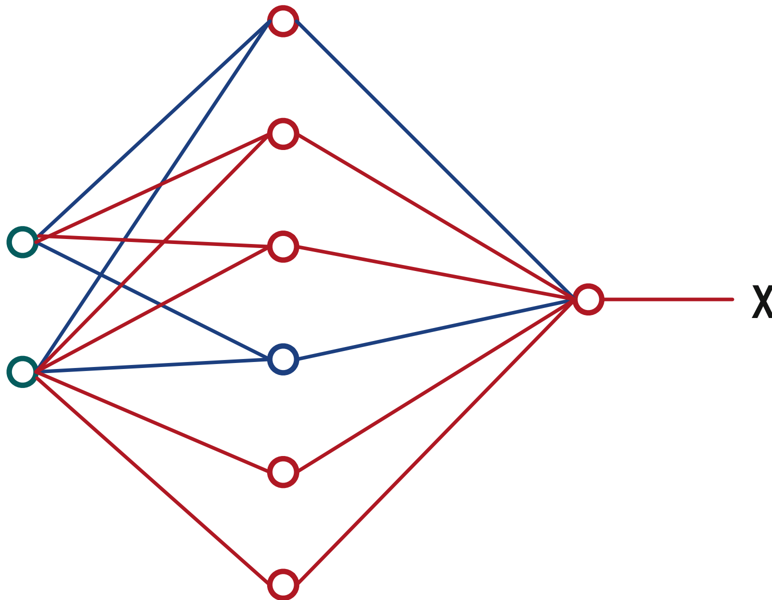


Chapter 10 presents deep learning with distinctive layers. Several different types of elements are in the deep learning chain. This is applied to face recognition. Face recognition is available in almost every photo application. Many social media sites, such as Facebook and Google Plus, also use face recognition. Cameras have built-in face recognition, though not identification, to help with focusing when taking portraits. Our goal is to get the algorithm to match faces, not classify them. Data classification is covered in Chapter 8.

Chapter 11 introduces a neural network as part of an adaptive control system. This ties together learning, via neural networks, and control.

### 1.6.4 Support Vector Machines

Support vector machines (SVMs) are supervised learning models with associated learning algorithms that analyze data used for classification and regression analysis. An SVM training algorithm builds a model that assigns examples into categories. The goal of SVMs is to produce a model, based on the training data, that predicts the target values.

In SVMs, nonlinear mapping of input data in a higher dimensional feature space is done with kernel functions. In this feature space, a separation hyperplane is generated that is the solution to the classification problem. The kernel functions can be polynomials, sigmoidal functions, and radial basis functions. Only a subset of the training data is needed, these are known as the support vectors [8]. The training is done by solving a quadratic program, which can be done with many numerical software.

## 1.7 Artificial Intelligence

### 1.7.1 What is Artificial Intelligence?

A test of artificial intelligence is the Turing test. The idea is that if you have a conversation with a machine and you can't tell it is a machine, then it should be considered intelligent. By this definition, many robo-calling systems might be considered intelligent. Another example, chess programs, can beat all but the best players, but a chess program can't do anything but play chess. Is a chess program intelligent? What we have now is machines that can do things pretty well in a particular context.

### 1.7.2 Intelligent Cars

Our "artificial intelligence" example is really a blending of Bayesian estimation and controls. It still reflects a machine doing what we would consider as intelligent behavior. This, of course, gets back to the question of defining intelligence.

Autonomous driving is an area of great interest to automobile manufacturers and to the general public. Autonomous cars are driving the streets today, but are not yet ready for general use by the public. There are many technologies involved in autonomous driving. These include:

1. Machine vision – turning camera data into information useful for the autonomous control system
2. Sensing – using many technologies including vision, radar, and sound to sense the environment around the car
3. Control – using algorithms to make the car go where it is supposed to go as determined by the navigation system
4. Machine learning – using massive data from test cars to create databases of responses to situations
5. GPS navigation – blending GPS measurements with sensing and vision to figure out where to go
6. Communications/ad hoc networks – talking with other cars to help determine where they are and what they are doing

All of the areas overlap. Communications and ad hoc networks are used with GPS navigation to determine both absolute location (what street and address corresponds to your location) and relative navigation (where you are with respect to other cars). In this context, the Turing test would be if you couldn't tell if a car was driven by a person or the computer. Now, since many drivers are bad, one could argue that a computer that drove really well would fail the Turing test! This gets back to the question of what intelligence is.

This example explores the problem of a car being passed by multiple cars and needing to compute tracks for each one. We are really addressing just the control and collision avoidance

problem. A single sensor version of Track Oriented Multiple Hypothesis Testing is demonstrated for a single car on a two-lane road. The example includes MATLAB graphics that make it easier to understand the thinking of the algorithm. The demo assumes that the optical or radar pre-processing has been done and that each target is measured by a single "blip" in two dimensions. An automobile simulation is included. It involves cars passing the car that is doing the tracking. The passing cars use a passing control system that is in itself a form of machine intelligence.

Our autonomous driving recipes use an Unscented Kalman Filter for the estimation of the state. This is the underlying algorithm that propagates the state (that is, advances the state in time in a simulation) and adds measurements to the state. A Kalman Filter, or other estimator, is the core of many target-tracking systems.

The recipes will also introduce graphics aids to help you understand the tracking decision process. When you implement a learning system you want to make sure it is working the way you think it should, or understand why it is working the way it does.

### 1.7.3 Expert Systems

A system that uses a knowledge base to reason and present the user with a result and an explanation of how it arrived at that result. Expert systems are also known as knowledge-based systems. The process of building an expert system is called knowledge engineering. This involves a knowledge engineer, someone who knows how to build the expert system, interviewing experts for the knowledge needed to build the system. Some systems can induce rules from data, speeding up the data acquisition process.

An advantage of expert systems, over human experts, is that knowledge from multiple experts can be incorporated into the database. Another advantage is that the system can explain the process in detail so that the user knows exactly how the result was generated. Even an expert in a domain can forget to check certain things. An expert system will always methodically check its full database. It is also not affected by fatigue or emotions.

Knowledge acquisition is a major bottleneck in building expert systems. Another issue is that the system cannot extrapolate beyond what is programmed into the database. Care must be taken with using an expert system because it will generate definitive answers for problems where there is uncertainty. The explanation facility is important, because someone with domain knowledge can judge the results from the explanation. In cases where uncertainty needs to be considered, a probabilistic expert system is recommended. A Bayesian network can be used as an expert system. A Bayesian network is also known as a belief network. It is a probabilistic graphical model that represents a set of random variables and their dependencies. In the simplest cases, a Bayesian network can be constructed by an expert. In more complex cases, it needs to be generated from data from machine learning. Chapter 12 delves into expert systems.

In Chapter 14, we explore a simple case-based reasoning system. An alternative would be a rule-based system.

## 1.8 Summary

All of the technologies in this chapter are in current use today. Any one of them can form the basis for a useful product. Many systems, such as autonomous cars, use several. We hope that our broad view of the field of machine learning and our unique taxonomy, which shows the relationships of machine learning and artificial intelligence to the classical fields of control and optimization, are useful to you. In the remainder of the book we will show you how to build software that implements these technologies. This can form the basis of your own more robust production software, or help you to use the many fine commercial products more effectively. Table 1.1 lists the scripts included in the companion code.

***Table 1.1:*** Chapter Code Listing

| File | Description |
|---|---|
| LinearRegression | A script that demonstrates linear regression and curve fitting. |

CHAPTER 2

# Representation of Data for Machine Learning in MATLAB

## 2.1 Introduction to MATLAB Data Types

### 2.1.1 Matrices

By default, all variables in MATLAB are double precision matrices. You do not need to declare a type for these variables. Matrices can be multidimensional and are accessed using 1-based indices via parentheses. You can address elements of a matrix using a single index, taken column-wise, or one index per dimension. To create a matrix variable, simply assign a value to it, such as this 2x2 matrix `a`:

```
>> a = [1 2; 3 4];
>> a(1,1)
     1

>> a(3)
     2
```

---

■ **TIP** A semicolon terminates an expression so that it does not appear in the command window. If you leave out the semicolon, it will print in the command window. Leaving out semicolons is a convenient way of debugging without using the MATLAB debugger, but it can be hard to find those missing semicolons later!

---

You can simply add, subtract, multiply, and divide matrices with no special syntax. The matrices must be the correct size for the linear algebra operation requested. A transpose is indicated using a single quote suffix, `A'`, and the matrix power uses the operator `^`.

```
>> b = a'*a;
>> c = a^2;
>> d = b + c;
```

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_2

By default, every variable is a numerical variable. You can initialize matrices to a given size using the `zeros`, `ones`, `eye`, or `rand` functions, which produce zeros, ones, identity matrices (ones on the diagonal), and random numbers respectively. Use `isnumeric` to identify numeric variables.

***Table 2.1:*** Key Functions for Matrices

| Function | Purpose |
|---|---|
| zeros | Initialize a matrix to zeros |
| ones | Initialize a matrix to ones |
| eye | Initialize an identity matrix |
| rand, randn | Initialize a matrix of random numbers |
| isnumeric | Identify a matrix or scalar numeric value |
| isscalar | Identify a scalar value (a 1 x 1 matrix) |
| size | Return the size of the matrix |

MATLAB can support n-dimensional arrays. A two-dimensional array is like a table. A three-dimensional array can be visualized as a cube where each box inside the cube contains a number. A four-dimensional array is harder to visualize, but we needn't stop there!

### 2.1.2 Cell Arrays

One variable type unique to MATLAB is cell arrays. This is really a list container, and you can store variables of any type in elements of a cell array. Cell arrays can be multi-dimensional, just like matrices, and are useful in many contexts.

Cell arrays are indicated by curly braces, {}. They can be of any dimension and contain any data, including string, structures, and objects. You can initialize them using the `cell` function, recursively display the contents using `celldisp`, and access subsets using parentheses, just like for a matrix. A short example is below.

```
>> c = cell(3,1);
>> c{1} = 'string';
>> c{2} = false;
>> c{3} = [1 2; 3 4];
>> b = c(1:2);
>> celldisp(b)
b{1} =
string

b{2} =
     0
```

Using curly braces for access gives you the element data as the underlying type. When you access elements of a cell array using parentheses, the contents are returned as another cell array, rather than the cell contents. MATLAB help has a special section called *Comma-Separated*

*Lists*, which highlights the use of cell arrays as lists. The code analyzer will also suggest more efficient ways to use cell arrays. For instance,

Replace

```
a = {b{:} c};
```

with

```
a = [b {c}];
```

Cell arrays are especially useful for sets of strings, with many of MATLAB's string search functions optimized for cell arrays, such as `strcmp`.

Use `iscell` to identify cell array variables. Use `deal` to manipulate structure array and cell array contents.

***Table 2.2:*** Key Functions for Cell Arrays

| Function | Purpose |
|---|---|
| cell | Initialize a cell array |
| cellstr | Create a cell array from a character array |
| iscell | Identify a cell array |
| iscellstr | Identify a cell array containing only strings |
| celldisp | Recursively display the contents of a cell array |

### 2.1.3 Data Structures

Data structures in MATLAB are highly flexible, leaving it up to the user to enforce consistency in fields and types. You are not required to initialize a data structure before assigning fields to it, but it is a good idea to do so, especially in scripts, to avoid variable conflicts.

Replace

```
d.fieldName = 0;
```

with

```
d = struct;
d.fieldName = 0;
```

In fact, we have found it generally a good idea to create a special function to initialize larger structures that are used throughout a set of functions. This is similar to creating a class definition. Generating your data structure from a function, instead of typing out the fields in a script, means that you always start with the correct fields. Having an initialization function also allows you to specify the types of variables and provide sample or default data. Remember, since MATLAB does not require you to declare variable types, doing so yourself with default data makes your code that much clearer.

---

■ **TIP** Create an initialization function for data structures.

---

You make a data structure into an array simply by assigning an additional copy. The fields must be identically named (they are case-sensitive) and in the same order, which is yet another

reason to use a function to initialize your structure. You can nest data structures with no limit on depth.

```
d = MyStruct;
d(2) = MyStruct;

function d = MyStruct

d = struct;
d.a = 1.0;
d.b = 'string';
```

MATLAB now allows for *dynamic field names* using variables, i.e., `structName.(dynamicExpression)`. This provides improved performance over `getfield`, where the field name is passed as a string. This allows for all sorts of inventive structure programming. Take our data structure array in the previous code snippet, and let's get the values of field `a` using a dynamic field name; the values are returned in a cell array.

```
>> field = 'a';
>> values = {d.(field)}

values =
    [1]    [1]
```

Use `isstruct` to identify structure variables and `isfield` to check for the existence of fields. Note that `isempty` will return *false* for a struct initialized with `struct`, even if it has no fields.

```
>> d = struct
d =
  struct with no fields.

>> isempty(d)

ans =
  logical
   0
```

***Table 2.3:*** Key Functions for Structs

| Function | Purpose |
|---|---|
| struct | Initialize a structure with or without fields |
| isstruct | Identify a structure |
| isfield | Determine if a field exists in a structure |
| fieldnames | Get the fields of a structure in a cell array |
| rmfield | Remove a field from a structure |
| deal | Set fields in a structure array to a value |

### 2.1.4 Numerics

Although MATLAB defaults to doubles for any data entered at the command line or in a script, you can specify a variety of other numeric types, including `single`, `uint8`, `uint16`, `uint32`, `uint64`, `logical` (i.e., an array of Booleans). Use of the integer types is especially relevant to using large data sets such as images. Use the minimum data type you need, especially when your data sets are large.

### 2.1.5 Images

MATLAB supports a variety of formats including GIF, JPG, TIFF, PNG, HDF, FITS, and BMP. You can read in an image directly using `imread`, which can determine the type automatically from the extension, or `fitsread`. (FITS stands for Flexible Image Transport System and the interface is provided by the CFITSIO library.) `imread` has special syntaxes for some image types, such as handling alpha channels for PNG, so you should review the options for your specific images. `imformats` manages the file format registry and allows you to specify handling of new user-defined types, if you can provide read and write functions.

You can display an image using either `imshow`, `image`, or `imagesc`, which scales the colormap for the range of data in the image.

For example, we use a set of images of cats in Chapter 7, on face recognition. The following is the image information for one of these sample images

```
>> imfinfo('IMG_4901.JPG')
ans =
           Filename: 'MATLAB/Cats/IMG_4901.JPG'
        FileModDate: '28-Sep-2016␣12:48:15'
           FileSize: 1963302
             Format: 'jpg'
      FormatVersion: ''
              Width: 3264
             Height: 2448
           BitDepth: 24
          ColorType: 'truecolor'
    FormatSignature: ''
    NumberOfSamples: 3
       CodingMethod: 'Huffman'
      CodingProcess: 'Sequential'
            Comment: {}
               Make: 'Apple'
              Model: 'iPhone␣6'
        Orientation: 1
        XResolution: 72
        YResolution: 72
     ResolutionUnit: 'Inch'
           Software: '9.3.5'
           DateTime: '2016:09:17␣22:05:08'
   YCbCrPositioning: 'Centered'
```

```
      DigitalCamera: [1x1 struct]
            GPSInfo: [1x1 struct]
      ExifThumbnail: [1x1 struct]
```

These are the metadata that tell camera software, and image databases, where and how the image was generated. This is useful when learning from images as it allows you to correct for resolution (`width` and `height`) bit depth and other factors.

If we view this image using `imshow`, it will publish a warning that the image is too big to fit on the screen and that it is displayed at 33%. If we view it using `image`, there will be a visible set of axes. `image` is useful for displaying other two-dimensional matrix data such as individual elements per pixel. Both functions return a handle to an image object; only the axes' properties are different.

```
>> figure; hI = image(imread('IMG_2398_Zoom.png'))
hI =
  Image with properties:

           CData: [680x680x3 uint8]
     CDataMapping: 'direct'

  Show all properties
```

***Figure 2.1:*** Image display options.

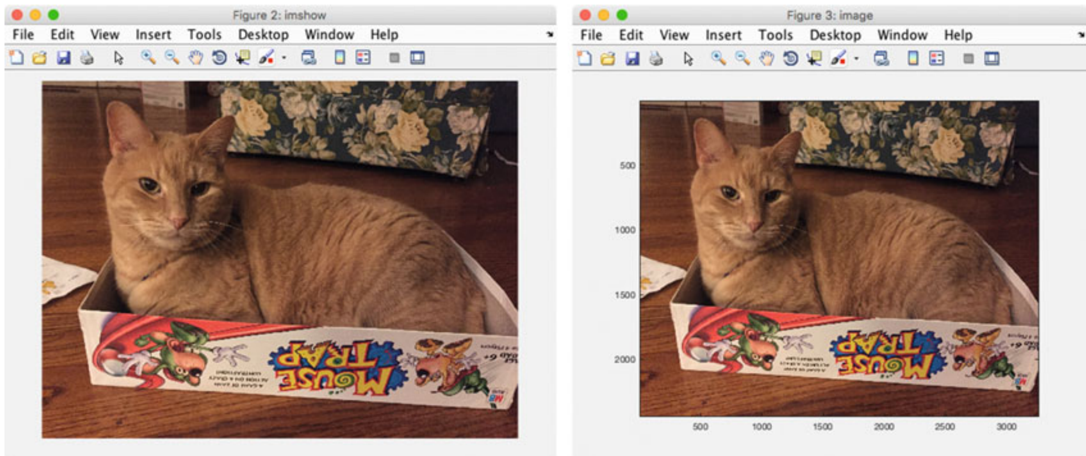

***Table 2.4:*** Key Functions for Images

| Function | Purpose |
|---|---|
| imread | Read an image in a variety of formats |
| imfinfo | Gather information about an image file |
| imformats | Determine if a field exists in a structure |
| imwrite | Write data to an image file |
| image | Display image from array |
| imagesc | Display image data scaled to the current colormap |
| imshow | Display an image, optimizing figure, axes, and image object properties, and taking an array or a filename as an input |
| rgb2gray | Write data to an image file |
| ind2rgb | Convert index data to RGB |
| rgb2ind | Convert RGB data to indexed image data |
| fitsread | Read a FITS file |
| fitswrite | Write data to a FITS file |
| fitsinfo | Information about a FITS file returned in a data structure |
| fitsdisp | Display FITS file metadata for all HDUs in the file |

### 2.1.6 Datastore

Datastores allow you to interact with files containing data that are too large to fit in memory. There are different types of datastores for tabular data, images, spreadsheets, databases, and custom files. Each datastore provides functions to extract smaller amounts of data that fit in the memory for analysis. For example, you can search a collection of images for those with the brightest pixels or maximum saturation values. We will use our directory of cat images as an example.

```
>> location = pwd
location =
/Users/Shared/svn/Manuals/MATLABMachineLearning/MATLAB/Cats
>> ds = datastore(location)
ds =
  ImageDatastore with properties:

      Files: {
             '␣.../Shared/svn/Manuals/MATLABMachineLearning/MATLAB/
                Cats/IMG_0191.png';
             '␣.../Shared/svn/Manuals/MATLABMachineLearning/MATLAB/
                Cats/IMG_1603.png';
             '␣.../Shared/svn/Manuals/MATLABMachineLearning/MATLAB/
                Cats/IMG_1625.png'
              ... and 19 more
             }
     Labels: {}
    ReadFcn: @readDatastoreImage
```

Once the datastore is created, you use the applicable class functions to interact with it. Datastores have standard container-style functions such as `read`, `partition`, and `reset`. Each type of datastore has different properties. The `DatabaseDatastore` requires the Database Toolbox, and allows you to use SQL queries.

MATLAB provides the `MapReduce` framework for working with out-of-memory data in datastores. The input data can be any of the datastore types, and the output is a key-value datastore. The map function processes the datastore input in chunks and the reduce function calculates the output values for each key. `mapreduce` can be sped up by using it with the MATLAB Parallel Computing Toolbox, Distributed Computer Server, or Compiler.

***Table 2.5:*** Key Functions for Datastore

| **Function** | **Purpose** |
| --- | --- |
| datastore | Create a datastore |
| read | Read a subset of data from the datastore |
| readall | Read all of the data in the datastore |
| hasdata | Check to see if there are more data in the datastore |
| reset | Initialize a datastore with the contents of a folder |
| partition | Excerpt a portion of the datastore |
| numpartitions | Estimate a reasonable number of partitions |
| ImageDatastore | Datastore of a list of image files. |
| TabularTextDatastore | A collection of one or more tabular text files. |
| SpreadsheetDatastore | Datastore of spreadsheets. |
| FileDatastore | Datastore for files with a custom format, for which you provide a reader function. |
| KeyValueDatastore | Datastore of key-value pairs. |
| DatabaseDatastore | Database connection, requires the Database Toolbox. |

### 2.1.7 Tall Arrays

Tall arrays were introduced in R2016b Release of MATLAB. They are allowed to have more rows than will fit in the memory. You can use them to work with datastores that might have millions of rows. Tall arrays can use almost any MATLAB type as a column variable, including numeric data, cell arrays, strings, datetimes, and categoricals. The MATLAB documentation provides a list of functions that support tall arrays. Results for operations on the array are only evaluated when they are explicitly requested using the `gather` function. The `histogram` function can be used with tall arrays and will execute immediately.

The MATLAB Statistic and Machine Learning Toolbox™, Database Toolbox, Parallel Computing Toolbox, Distributed Computing Server, and Compiler all provide additional extensions

for working with tall arrays. For more information about this new feature, use the following topics in the documentation:

- Tall Arrays
- Analysis of Big Data with Tall Arrays
- Functions That Support Tall Arrays (A–Z)
- Index and View Tall Array Elements
- Visualization of Tall Arrays
- Extend Tall Arrays with Other Products
- Tall Array Support, Usage Notes, and Limitations

***Table 2.6:*** Key Functions for Tall Arrays

| **Function** | **Purpose** |
|---|---|
| tall | Initialize a tall array |
| gather | Execute the requested operations |
| summary | Display summary information to the command line |
| head | Access the first rows of a tall array |
| tail | Access the last rows of a tall array |
| istall | Check the type of array to determine if it is tall |
| write | Write the tall array to disk |

### 2.1.8 Sparse Matrices

Sparse matrices are a special category of matrix in which most of the elements are zero. They appear commonly in large optimization problems and are used by many such packages. The zeros are "squeezed" out and MATLAB stores only the nonzero elements along with index data such that the full matrix can be recreated. Many regular MATLAB functions, such as `chol` or `diag`, preserve the sparseness of an input matrix.

### 2.1.9 Tables and Categoricals

Tables were introduced in release R2013 of MATLAB and allow tabular data to be stored with metadata in one workspace variable. It is an effective way of storing and interacting with data that one might put in, or import from, a spreadsheet. The table columns can be named, assigned units and descriptions, and accessed as one would fields in a data structure, i.e., `T.DataName`. See `readtable` on creating a table from a file, or try out the Import Data button from the command window.

***Table 2.7:*** Key Functions for Sparse Matrices

| **Function** | **Purpose** |
|---|---|
| sparse | Create a sparse matrix from a full matrix or from a list of indices and values |
| issparse | Determine if a matrix is sparse |
| nnz | Number of nonzero elements in a sparse matrix |
| spalloc | Allocate a nonzero space for a sparse matrix |
| spy | Visualize a sparsity pattern |
| spfun | Selectively apply a function to the nonzero elements of a sparse matrix |
| full | Convert a sparse matrix to full form |

Categorical arrays allow for storage of discrete non-numeric data, and they are often used within a table to define groups of rows. For example, time data may have the day of the week, or geographic data may be organized by state or county. They can be leveraged to rearrange data in a table using `unstack`.

You can also combine multiple data sets into single tables using `join`, `innerjoin`, and `outerjoin`, which will be familiar to you if you have worked with databases.

***Table 2.8:*** Key Functions for Tables

| **Function** | **Purpose** |
|---|---|
| table | Create a table with data in the workspace |
| readtable | Create a table from a file |
| join | Merge tables by matching up variables |
| innerjoin | Join tables A and B retaining only the rows that match |
| outerjoin | Join tables including all rows |
| stack | Stack data from multiple table variables into one variable |
| unstack | Unstack data from a single variable into multiple variables |
| summary | Calculate and display summary data for the table |
| categorical | Arrays of discrete categorical data |
| iscategorical | Create a categorical array |
| categories | List of categories in the array |
| iscategory | Test for a particular category |
| addcats | Add categories to an array |
| removecats | Remove categories from an array |
| mergecats | Merge categories |

### 2.1.10 Large MAT-Files

You can access parts of a large MAT-file without loading the entire file into the memory by using the `matfile` function. This creates an object that is connected to the requested MAT-file without loading it. Data are only loaded when you request a particular variable, or part of a variable. You can also dynamically add new data to the MAT-file.

For example, we can load a MAT-file of neural net weights generated in a later chapter.

```
>> m = matfile('PitchNNWeights','Writable',true)
m =
  matlab.io.MatFile

  Properties:
      Properties.Source: '/Users/Shared/svn/Manuals/
         MATLABMachineLearning/MATLAB/PitchNNWeights.mat'
    Properties.Writable: true
                      w: [1x8 double]
```

We can access a portion of the previously unloaded `w` variable, or add a new variable `name`, all using this object `m`.

```
>> y = m.w(1:4)
y =
     1     1     1     1
>> m.name = 'Pitch␣Weights'
m =
  matlab.io.MatFile

  Properties:
      Properties.Source: '/Users/Shared/svn/Manuals/
         MATLABMachineLearning/MATLAB/PitchNNWeights.mat'
    Properties.Writable: true
                   name: [1x13 char]
                      w: [1x8  double]
>> d = load('PitchNNWeights')
d =
       w: [1 1 1 1 1 1 1 1]
    name: 'Pitch␣Weights'
```

There are some limits to the indexing into unloaded data, such as struct arrays and sparse arrays. Also, `matfile` requires MAT-files using version 7.3, which is not the default for a generic `save` operation as of R2016b Release of MATLAB. You must either create the MAT-file using `matfile` to take advantage of these features or use the `-v7.3'` flag when saving the file.

## 2.2 Initializing a Data Structure Using Parameters

It's always a good idea to use a special function to define a data structure you are using as a type in your codebase, similar to writing a class but with less overhead. Users can then overload individual fields in their code, but there is an alternative way of setting many fields at once: an initialization function that can handle a parameter pair input list. This allows you to do additional processing in your initialization function. Also, your parameter string names can be more descriptive than you would choose to make your field names.

### 2.2.1 Problem

We want to initialize a data structure so that the user clearly knows what he or she is entering.

### 2.2.2 Solution

The simplest way of implementing the parameter pairs is using `varargin` and a switch statement. Alternatively, you could write an `inputParser`, which allows you to specify required and optional inputs as well as named parameters. In that case, you have to write separate or anonymous functions for validation that can be passed to the `inputParser`, rather than just write out the validation in your code.

### 2.2.3 How It Works

We will use the data structure developed for the automobile simulation in Chapter 12 as an example. The header lists the input parameters along with the input dimensions and units, if applicable.

```
%% AUTOMOBILEINITIALIZE Initialize the automobile data structure.
%
%% Form
%  d = AutomobileInitialize( varargin )
%
%% Description
% Initializes the data structure using parameter pairs.
%
%% Inputs
% varargin:  ('parameter',value,...)
%
% 'mass'                                (1,1)  (kg)
% 'steering angle'                      (1,1)  (rad)
% 'position tires'                      (2,4)  (m)
% 'frontal drag coefficient'            (1,1)
% 'side drag coefficient'               (1,1)
% 'tire friction coefficient'           (1,1)
% 'tire radius'                         (1,1)  (m)
% 'engine torque'                       (1,1)  (Nm)
```

```
% 'rotational inertia'                    (1,1)  (kg-m^2)
% 'state'                                 (6,1)  [m;m;m/s;m/s;rad;rad/s]
```

The function first creates the data structure using a set of defaults, then handles the parameter pairs entered by a user. After the parameters have been processed, two areas are calculated using the dimensions and the height.

```
function d = AutomobileInitialize( varargin )

% Defaults
d.mass          = 1513;
d.delta         = 0;
d.r             = [  1.17 1.17 -1.68 -1.68;...
                    -0.77 0.77 -0.77  0.77];
d.cDF           = 0.25;
d.cDS           = 0.5;
d.cF              = 0.01; % Ordinary car tires on concrete
d.radiusTire      = 0.4572; % m
d.torque        = d.radiusTire*200.0; % N
d.inr           = 2443.26;
d.x             = [0;0;0;0;0;0];
d.fRR           = [0.013 6.5e-6];
d.dim           = [1.17+1.68 2*0.77];
d.h             = 2/0.77;
d.errOld        = 0;
d.passState     = 0;
d.model         = 'MyCar.obj';
d.scale         = 4.7981;

for k = 1:2:length(varargin)
  switch lower(varargin{k})
    case 'mass'
      d.mass         = varargin{k+1};
    case 'steering angle'
      d.delta        = varargin{k+1};
    case 'position tires'
      d.r            = varargin{k+1};
    case 'frontal drag coefficient'
      d.cDF          = varargin{k+1};
    case 'side drag coefficient'
      d.cDS          = varargin{k+1};
    case 'tire friction coefficient'
      d.cF           = varargin{k+1};
    case 'tire radius'
      d.radiusTire       = varargin{k+1};
    case 'engine torque'
      d.torque       = varargin{k+1};
    case 'rotational inertia'
```

```
      d.inertia     = varargin{k+1};
    case 'state'
      d.x           = varargin{k+1};
    case 'rolling␣resistance␣coefficients'
      d.fRR         = varargin{k+1};
    case 'height␣automobile'
      d.h           = varargin{k+1};
    case 'side␣and␣frontal␣automobile␣dimensions'
      d.dim         = varargin{k+1};
    case 'car␣model'
      d.model       = varargin{k+1};
    case 'car␣scale'
      d.scale       = varargin{k+1};
  end
end

% Processing
d.areaF = d.dim(2)*d.h;
d.areaS = d.dim(1)*d.h;
d.g     = LoadOBJ(d.model,[],d.scale);
```

To perform the same tasks with `inputParser`, you add a `addRequired`, `addOptional`, or `addParameter` call for every item in the switch statement. The named parameters require default values. You can optionally specify a validation function; in the example below we use `isNumeric` to limit the values to numeric data.

```
>> p = inputParser
p.addParameter('mass',0.25);
p.addParameter('cDF',1513);
p.parse('cDF',2000);
d = p.Results

p =

  inputParser with properties:

       FunctionName: ''
      CaseSensitive: 0
      KeepUnmatched: 0
    PartialMatching: 1
       StructExpand: 1
         Parameters: {1x0 cell}
            Results: [1x1 struct]
          Unmatched: [1x1 struct]
      UsingDefaults: {1x0 cell}
```

```
d =

  struct with fields:

     cDF: 2000
    mass: 0.2500
```

In this case, the results of the parsed parameters are stored in a Results substructure.

## 2.3 Performing MapReduce on an Image Datastore

### 2.3.1 Problem

We discussed the `datastore` class in the introduction to the chapter. Now let's use it to perform analysis on the full set of cat images using `mapreduce`, which is scalable to very large numbers of images. This involves two steps; first a *map* step that operates on the datastore and creates intermediate values, and then a *reduce* step that operates on the intermediate values to produce a final output.

### 2.3.2 Solution

We create the `datastore` by passing in the path to the folder of cat images. We also need to create a map function and a reduce function, to pass into `mapreduce`. If you are using additional toolboxes such as the Parallel Computing Toolbox, you would specify the reduce environment using `mapreducer`.

### 2.3.3 How It Works

First, create the `datastore` using the path to the images.

```
>> imds = imageDatastore('MATLAB/Cats');
imds =
  ImageDatastore with properties:

      Files: {
              '␣.../MATLABMachineLearning/MATLAB/Cats/IMG_0191.png';
              '␣.../MATLABMachineLearning/MATLAB/Cats/IMG_1603.png';
              '␣.../MATLABMachineLearning/MATLAB/Cats/IMG_1625.png'
               ... and 19 more
              }
     Labels: {}
    ReadFcn: @readDatastoreImage
```

Second, we write the map function. This must generate and store a set of intermediate values that will be processed by the reduce function. Each intermediate value must be stored as a key in the intermediate key-value datastore using `add`. In this case, the map function will receive one image each time it is called. We call it `catColorMapper`, since it processes the red, green, and blue values for each image using a simple average.

```
function catColorMapper(data, info, intermediateStore)

% Calculate the average (R,G,B) values
avgRed = mean(mean(data(:,:,1)));
avgGreen = mean(mean(data(:,:,2)));
avgBlue = mean(mean(data(:,:,3)));

% Store the calculated values with text keys
add(intermediateStore, 'Avg␣Red', struct('Filename',info.Filename,'
    Val', avgRed));
add(intermediateStore, 'Avg␣Green', struct('Filename',info.Filename,'
    Val', avgGreen));
add(intermediateStore, 'Avg␣Blue', struct('Filename',info.Filename,'
    Val', avgBlue));
```

The reduce function will then receive the list of image files from the datastore once for each key in the intermediate data. It receives an iterator to the intermediate data store as well as an output data store. Again, each output must be a key-value pair. The `hasnext` and `getnext` functions used are part of the `mapreduce` ValueIterator class. In this case, we find the minimum value for each key across the set of images.

```
function catColorReducer(key, intermediateIter, outputStore)

% Iterate over values for each key
minVal = 255;
minImageFilename = '';
while hasnext(intermediateIter)
  value = getnext(intermediateIter);

  % Compare values to find the minimum
  if value.Val < minVal
      minVal = value.Val;
      minImageFilename = value.Filename;
  end
end

% Add final key-value pair
add(outputStore, ['Minimum␣-␣␣' key], minImageFilename);
```

Finally, we call `mapreduce` using function handles to our two helper functions. Progress updates are printed to the command line, first for the mapping step, and then for the reduce step (once the mapping progress reaches 100%).

```
minRGB = mapreduce(imds, @catColorMapper, @catColorMapper);
```

```
********************************
*      MAPREDUCE PROGRESS      *
********************************
Map   0% Reduce   0%
Map  13% Reduce   0%
Map  27% Reduce   0%
Map  40% Reduce   0%
Map  50% Reduce   0%
Map  63% Reduce   0%
Map  77% Reduce   0%
Map  90% Reduce   0%
Map 100% Reduce   0%
Map 100% Reduce  33%
Map 100% Reduce  67%
Map 100% Reduce 100%
```

The results are stored in a MAT-file, for example, `results_1_28-Sep-2016_16-28-38_347`. The store returned is a key-value store to this MAT-file, which in turn contains the store with the final key-value results.

```
>> output = readall(minRGB)
output =
          Key                              Value
    _______________   ______________________________________________
    ''Minimum - Avg Red'      '/MATLAB/Cats/IMG_1625.png'
    ''Minimum - Avg Blue'     '/MATLAB/Cats/IMG_4866.jpg'
    ''Minimum - Avg Green'    '/MATLAB/Cats/IMG_4866.jpg'
```

You'll notice that the image files are different file types. This is because they came from different sources. MATLAB can handle most image types quite well.

## 2.4 Creating a Table from a File

Often with big data we have complex data in many files. MATLAB provides functions to make it easier to handle massive sets of data. In this section, we will collect data from a set of weather files and perform a Fast Fourier Transform (FFT) on data from two years. First, we will write the FFT function.

### 2.4.1 Problem

We want do FFTs.

### 2.4.2 Solution

Write a function using `fft` and compute the energy from the FFT. The energy is just the real part of the product of the FFT output and its transpose.

### 2.4.3 How It Works

The following functions takes in data `y` with a sample time `tSamp` and performs an FFT

```
function [e, w] = FFTEnergy( y, tSamp )

% Demo
if( nargin < 1 )
  Demo;
  return;
end

[n, m] = size( y );

if( n < m )
  y = y';
end

n = size( y, 1 );

% Check if an odd number and make even
if(2*floor(n/2) ~= n )
  n = n - 1;
  y = y(1:n,:);
end

x  = fft(y);
e  = real(x.*conj(x))/n;

hN = n/2;
e  = e(1:hN,:);
r  = 2*pi/(n*tSamp);
w  = r*(0:(hN-1));

if( nargout == 0 )
  tL = sprintf('FFT Energy Plot: Resolution = %10.2e rad/sec',r);
  PlotSet(w,e','x label','Frequency (rad/sec)','y label', 'Energy',
  'plot title', tL,'plot type', 'xlog', 'figure title', 'FFT');
  clear e
end
```

We get the energy using these two lines

```
x  = fft(y);
```

Taking the real part just accounts for numerical errors. The product of a number and its complex conjugate should be real.

The function computes the resolution. Notice it is a function of the sampling period and number of points.

```
e  = e(1:hN,:);
```

The built-in demo creates a series with a frequency at 1 rad/sec and a second at 2 rad/sec. The higher frequency one, with an amplitude of 2, has more energy as expected.

```
end

function Demo
%% Demo
tSamp    = 0.1;
omega1   = 1;
omega2   = 3;
t        = linspace(0,1000,10000)*tSamp;
y        = sin(omega1*t) + 2*sin(omega2*t);
```

Figure 2.2 shows the data and the FFT. Note the clearly visible frequencies in the FFT plot that match the oscillations in the time plot.

***Figure 2.2:*** The input data for the FFT and the results.

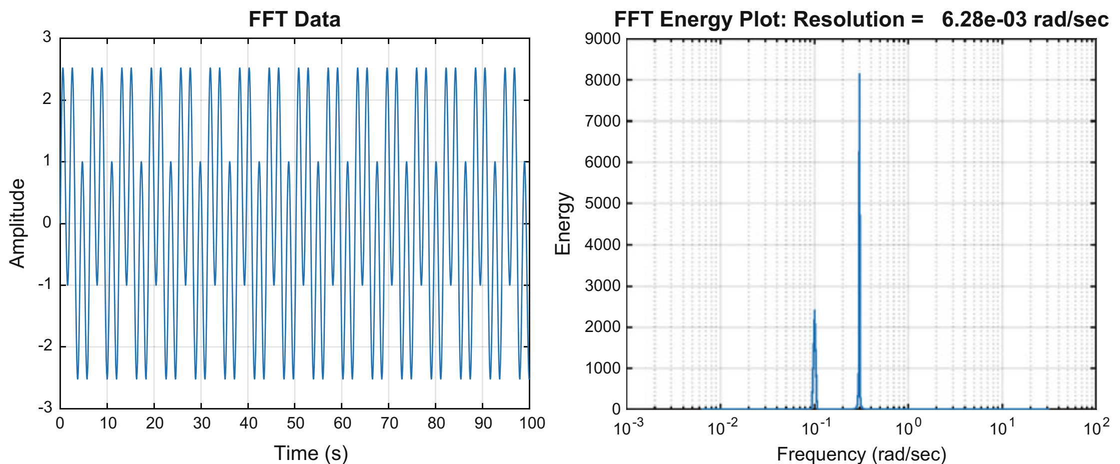


## 2.5 Processing Table Data

### 2.5.1 Problem

We want to compare temperature frequencies in 1999 and 2015 using data from a table.

### 2.5.2 Solution

Use `tabularTextDatastore` to load the data and perform an FFT on the data.

### 2.5.3 How It Works

First, let us look at what happens when we read in the data from the weather files.

```
>> tds = tabularTextDatastore('./Weather')

tds =

  TabularTextDatastore with properties:

                      Files: {
                             '␣.../MATLABMachineLearning2/MATLAB/
                                Chapter_02/Weather/HistKTTN_1990.txt
                                ';
                             '␣.../MATLABMachineLearning2/MATLAB/
                                Chapter_02/Weather/HistKTTN_1993.txt
                                ';
                             '␣.../MATLABMachineLearning2/MATLAB/
                                Chapter_02/Weather/HistKTTN_1999.txt
                                '
                              ... and 5 more
                             }
               FileEncoding: 'UTF-8'
   AlternateFileSystemRoots: {}
          ReadVariableNames: true
              VariableNames: {'EST', 'MaxTemperatureF', '
                 MeanTemperatureF' ... and 20 more}

  Text Format Properties:
             NumHeaderLines: 0
                  Delimiter: ','
               RowDelimiter: '\r\n'
             TreatAsMissing: ''
               MissingValue: NaN

  Advanced Text Format Properties:
            TextscanFormats: {'%{uuuu-MM-dd}D', '%f', '%f' ... and 20
                more}
                   TextType: 'char'
         ExponentCharacters: 'eEdD'
               CommentStyle: ''
                 Whitespace: '␣\b\t'
    MultipleDelimitersAsOne: false
```

```
  Properties that control the table returned by preview, read,
     readall:
      SelectedVariableNames: {'EST', 'MaxTemperatureF', '
         MeanTemperatureF' ... and 20 more}
            SelectedFormats: {'%{uuuu-MM-dd}D', '%f', '%f' ... and 20
                 more}
                   ReadSize: 20000 rows
```

WeatherFFT selects the data to use. It finds all the data in the mess of data in the files. When running the script you need to be in the same folder as WeatherFFT.

```
tDS                        = tabularTextDatastore('./Weather/');
tDS.SelectedVariableNames = {'EST','MaxTemperatureF'};

preview(tDS)

secInDay = 86400;

z = readall(tDS);

% The first column in the cell array is the date. year extracts the
   year
y     = year(z{:,1});
k1993 = find(y == 1993);
k2015 = find(y == 2015);
tSamp = secInDay;
t     = (1:365)*tSamp;
j     = {[1 2]};

%% Plot the FFT

% Get 1993 data
d1993     = z{k1993,2}';
m1993     = mean(d1993);
d1993     = d1993 - m1993;

e1993     = FFTEnergy( d1993, tSamp );

% Get 2015 data
d2015     = z{k2015,2}';
```

If the data do not exist TabularTextDatastore puts NaN in the data points' place. We happened to pick two years without any missing data. We use preview to see what we are getting.

```
>> WeatherFFT
Warning: Variable names were modified to make them valid MATLAB
   identifiers.
```

```
ans =

  8x2 table

        EST          MaxTemperatureF
    __________       _______________

    1990-01-01             39
    1990-01-02             39
    1990-01-03             48
    1990-01-04             51
    1990-01-05             46
    1990-01-06             43
    1990-01-07             42
    1990-01-08             37
```

In this script, we get output from `FFTEnergy` so that we can combine the plots. We chose to put the data on the same axes. Figure 2.3 shows the temperature data and the FFT.

***Figure 2.3:*** 1993 and 2015 data.

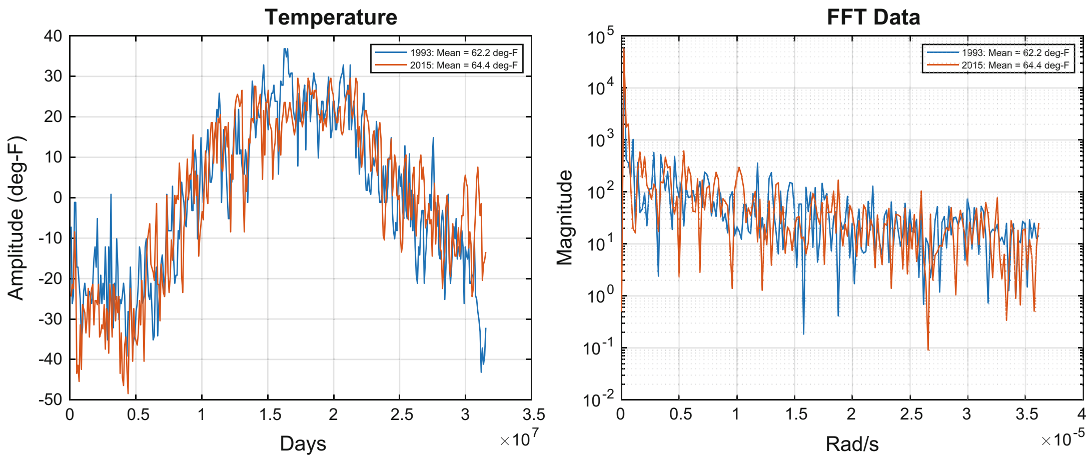


We get a little fancy with `plotset`. Our legend entries are computed to include the mean temperatures.

```
d2015       = d2015 - m2015;
[e2015,f] = FFTEnergy( d2015, tSamp );

lG = {{sprintf('1993:␣Mean␣=␣%4.1f␣deg-F',m1993) sprintf('2015:␣Mean␣
    =␣%4.1f␣deg-F',m2015)}};

PlotSet(t,[d1993;d2015],  'x␣label', 'Days', 'y␣label','Amplitude␣(
    deg-F)',...
```

```
  'plot␣title','Temperature', 'figure␣title', 'Temperature','legend',
     lG,'plot␣set',j);
```

## 2.6 Using MATLAB Strings

Machine learning often requires interaction with humans, which often means processing speech. Also, expert systems and fuzzy logic systems can make use of textual descriptions. MATLAB's string data type makes this easier. Strings are bracketed by double quotes. In this section, we will give examples of operations that work with strings, but not with character arrays.

### 2.6.1 String Concatenation

#### 2.6.1.1 Problem

We want to concatenate two strings.

#### 2.6.1.2 Solution

Create the two strings and use the "+" operator.

#### 2.6.1.3 How It Works

You can use the + operator to concatenate strings. The result is the second string after the first.

```
>> a = "12345";
>> b = "67";
>> c = a + b

c =

    "1234567"
```

### 2.6.2 Arrays of Strings

#### 2.6.2.1 Problem

We want any array of strings.

#### 2.6.2.2 Solution

Create the two strings and put them in a matrix.

#### 2.6.2.3 How It Works

We create the same two strings as above and use the matrix operator. If they were character arrays we would need to pad the shorter with blanks to be the same size as the longer.

```
>> a = "12345";
>> b = "67";
>> c = [a;b]

c =

  2$\times$1 string array
```

```
    "12345"
    "67"

>> c = [a b]

c =

  1$\times$2 string array

    "12345"    "67"
```

You could have used a cell array for this, but strings are often more convenient.

### 2.6.3 Substrings

#### 2.6.3.1 Problem

We want to get strings after a fixed prefix.

#### 2.6.3.2 Solution

Create a string array and use `extractAfter`.

#### 2.6.3.3 How It Works

Create a string array of strings to search and use `extractAfter`.

```
>> a = ["1234";"12456";"12890"];
f = extractAfter(a,"12")

f =

  3$\times$1 string array

    "34"
    "456"
    "890"
```

Most of the string functions work with `char`, but **strings** are a little cleaner. Here is the above example with **cell** arrays.

```
>> a = {'1234';'12456';'12890'};
>> f = extractAfter(a,"12")

f =

  3$\times$1 cell array

    {'34' }
    {'456'}
    {'890'}
```

## 2.7 Summary

There are a variety of data containers in MATLAB to assist you in analyzing your data for machine learning. If you have access to a computer cluster or one of the specialized computing toolboxes you have even more options. Table 2.9 lists the functions and scripts included in the companion code.

***Table 2.9:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| AutomobileInitialize | Data structure initialization example from Chapter 12. |
| catReducer | Image datastore used with `mapreduce`. |
| FFTEnergy | Computes the energy from an FFT. |
| weatherFFT | Does an FFT of weather data. |

CHAPTER 3

# MATLAB Graphics

One of the issues with machine learning is understanding the algorithms and why an algorithm made a particular decision. In addition, you want to be able to easily understand the decision. MATLAB has extensive graphics facilities that can be harnessed for that purpose. Plotting is used extensively in machine learning problems. MATLAB plots can be two- or three-dimensional. MATLAB also has many plot types such as line plots, bar charts, and pie charts. Different types of plots are better at conveying particular types of data. MATLAB also has extensive surface and contour plotting capabilities that can be used to display complex data in an easy-to-grasp fashion. Another facility is 3D modeling. You can draw animated objects, such as robots or automobiles. These are particularly valuable when your machine learning involves simulations.

An important part of MATLAB graphics is Graphical User Interface (GUI) building. MATLAB has extensive facilities for making GUIs. These can be a valuable way of making your design tools or machine learning systems easy for users to operate.

This chapter will provide an introduction to a wide variety of graphics tools in MATLAB. They should allow you to harness MATLAB graphics for your own applications.

## 3.1 2D Line Plots

### 3.1.1 Problem

You want a single function to generate two-dimensional line graphs, avoiding a long list of code for the generation of each graphic.

### 3.1.2 Solution

Write a single function to take the data and *parameter pairs* to encapsulate the functionality of MATLAB's 2D line plotting functions. An example of a plot created with a single line of code is shown in Figure 3.1.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_3

***Figure 3.1:*** `PlotSet`'s built-in demo.

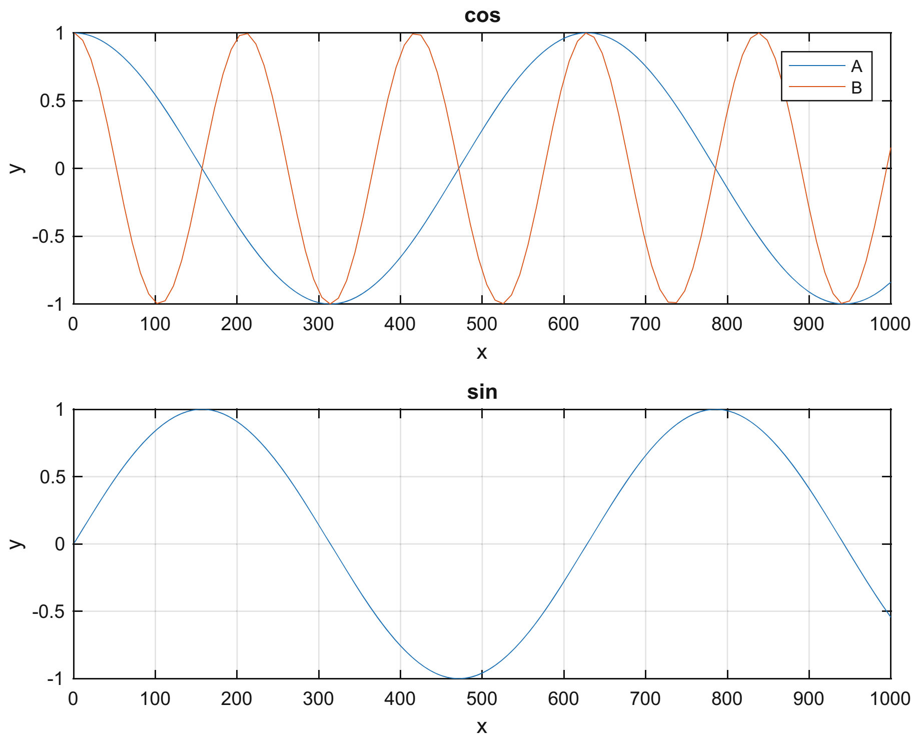


### 3.1.3 How It Works

`PlotSet` generates 2D plots, including multiple plots on a page.

```
function h = PlotSet( x, y, varargin )
```

This code processes `varargin` as parameter pairs to set options. A parameter pair is two inputs. The first is the name of the value and the second is the value. For example, the parameter pair for labeling the x-axis is:

```
'x label','Time (s)'
```

`varargin` makes it easy to expand the plotting options. The core function code is shown below. We supply default values for the $x$ and $y$ axis labels and the figure name. The parameter pairs are handled in a switch statement. The following code is the branch when there is only one x-axis label for all of the plots. It arranges plots by the data in `plotSet` that is a cell array.

```
for k = 1:m
  subplot(m,nCol,k);
  j = plotSet{k};
  for i = 1:length(j)
    plotXY(x,y(j(i),:),plotType);
    hold on
  end
```

```
    hold off
    xlabel(xLabel{1});
    ylabel(yLabel{k});
    if( length(plotTitle) == 1 )
      title(plotTitle{1})
    else
      title(plotTitle{k})
    end
    if( ~isempty(leg{k}) )
      legend(leg{k});
    end
    grid on
  end
```

The plotting is done in a subfunction called `plotXY`. There you see all the familiar MATLAB plotting function calls.

```
function plotXY(x,y,type)

switch type
  case 'plot'
    plot(x,y);
  case {'log' 'loglog' 'log␣log'}
    loglog(x,y);
  case {'xlog' 'semilogx' 'x␣log'}
    semilogx(x,y);
  case {'ylog' 'semilogy' 'y␣log'}
    semilogy(x,y);
  otherwise
    error('%s␣is␣not␣an␣available␣plot␣type',type);
end
```

The example in Figure 3.1 is generated by a dedicated demo function at the end of the `PlotSet` function. This demo shows several of the features of the function. These include:

1. Multiple lines per graph
2. Legends
3. Plot titles
4. Default axes labels

Using a dedicated demo subfunction is a clean way of providing a built-in example of a function, and it is especially important in graphics functions to provide an example of a typical plot. The code is shown below.

```
function Demo

x = linspace(1,1000);
y = [sin(0.01*x);cos(0.01*x);cos(0.03*x)];
disp('PlotSet:␣One␣x␣and␣two␣y␣rows')
```

## 3.2 General 2D Graphics

### 3.2.1 Problem

You want to represent a 2D data set in different ways. Line plots are very useful, but sometimes it is easier to visualize data in different forms. MATLAB has many functions for 2D graphical displays.

### 3.2.2 Solution

Write a script to show MATLAB's different 2D plot types. In our example we use subplots within one figure to help reduce figure proliferation.

### 3.2.3 How It Works

Use the `NewFigure` function to create a new figure window with a suitable name. Then run the following script.

```
>> NewFigure('My figure name')

ans =

  Figure (1: My figure name) with properties:

      Number: 1
        Name: 'My figure name'
       Color: [0.9400 0.9400 0.9400]
    Position: [560 528 560 420]
       Units: 'pixels'

  Show all properties

subplot(4,1,1);
plot(x,y);
subplot(4,1,2);
bar(x,y);
subplot(4,1,3);
barh(x,y);
ax4 = subplot(4,1,4);
pie(y)
colormap(ax4,'gray')
```

Four plot types are shown that are helpful in displaying 2D data. One is the 2D line plot, the same as that used in `PlotSet`. The middle two are bar charts. The final is a pie chart. Each gives you different insight into the data. Figure 3.2 shows the plot types.

There are many MATLAB functions for making these plots more informative. You can:

- Add labels
- Add grids

***Figure 3.2:*** Four different types of MATLAB 2D plots.

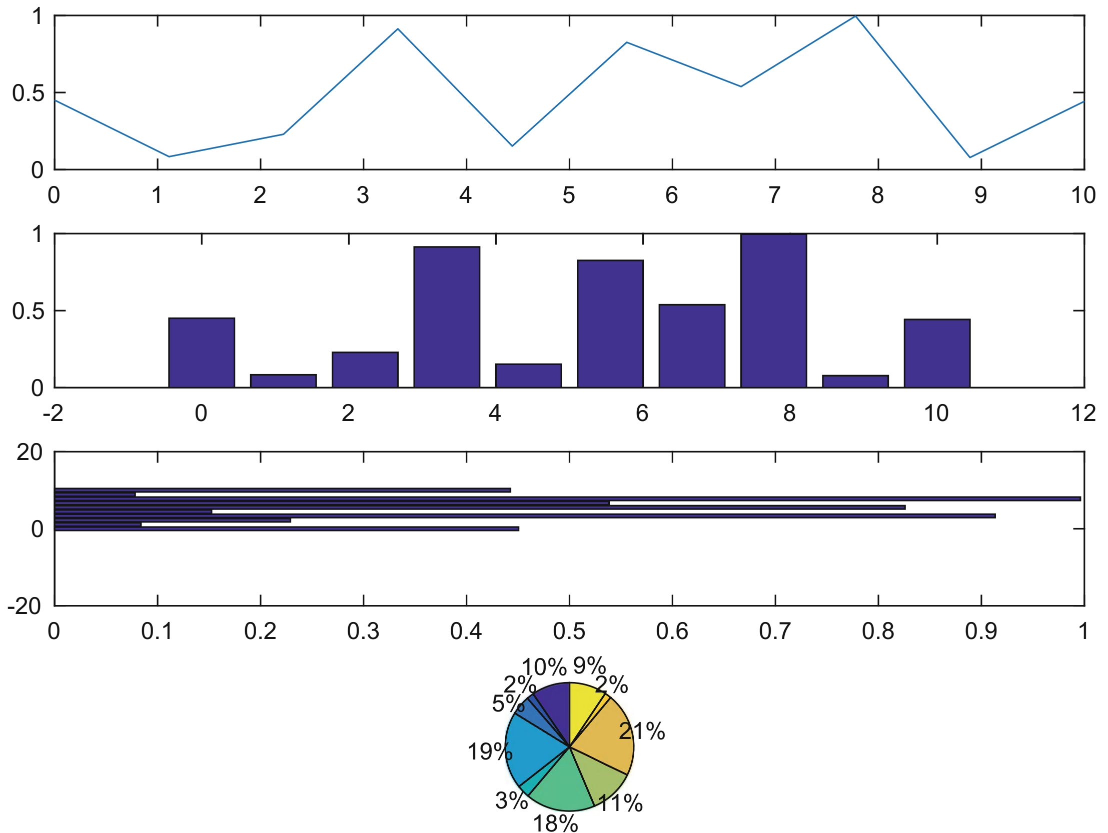


- Change font types and sizes
- Change the thickness of lines
- Add legends
- Change axes limits

The last item requires looking at the axes' properties. Here are the properties for the last plot – the list is very long! `gca` is the handle to the current axes. `get(gca)` returns a huge list, which we will not print here. Every single one of these can be changed by using the `set` function:

```
set(gca,'YMinorGrid','on','YGrid','on')
```

This uses parameter pairs just like `PlotSet`. In this list, `children` are pointers to the children of the axes. You can access those using `get` and change their properties using `set`. Any item that is added to an axis, such as axis labels, titles, lines, or other graphics objects, are all children of that axis.

## 3.3 Custom Two-Dimensional Diagrams

### 3.3.1 Problem

Many machine learning algorithms benefit from two-dimensional diagrams such as tree diagrams, to help the user understand the results and the operation of the software. Such diagrams, automatically generated by the software, are useful in many types of learning systems. This section gives an example of how to write MATLAB code for a tree diagram.

### 3.3.2 Solution

Our solution is to use the MATLAB `patch` function to automatically generate the blocks, and use `line` to generate connecting lines in the function `TreeDiagram`. Figure 3.3 shows the resulting hierarchical tree diagram. The circles are in rows and each row is labeled.

### 3.3.3 How It Works

Tree diagrams are very useful for machine learning. This function generates a hierarchical tree diagram with the nodes as circles with text within each node. The graphics functions used in this function are:

1. `line`
2. `patch`
3. `text`

The data needed to draw the tree are contained in a data structure, which is documented in the header. Each node has a parent field. This information is sufficient to make the connections. The node data are entered as a cell array.

***Figure 3.3:*** A custom tree diagram.

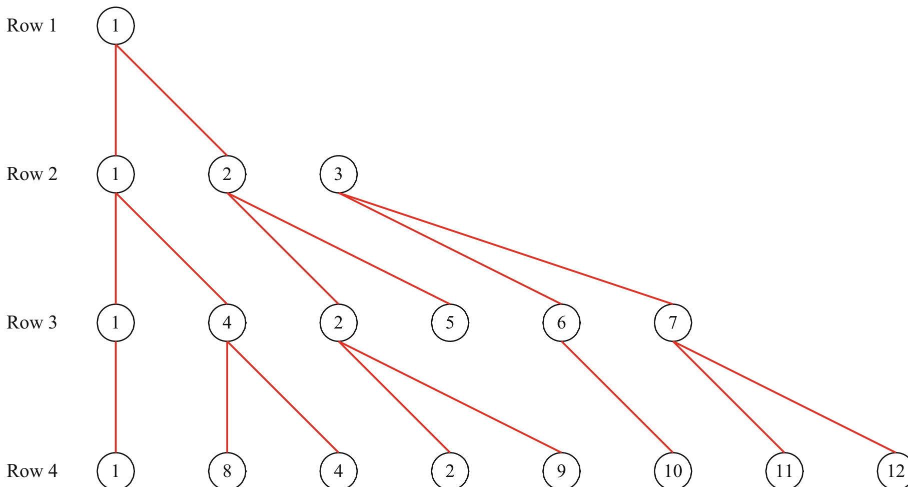

The function uses a figure handle as a persistent variable so that the same figure can be updated with subsequent calls, if desired.

```
if( ~update )
  figHandle = NewFigure(w.name);
else
  clf(figHandle)
end
```

The core drawing code is in `DrawNode`, which draws the boxes and `ConnectNode`, which connects the nodes with lines. Our nodes are circles with 20 segments. The `linspace` code makes sure that both 0 and $2\pi$ are not in the list of angles.

```
function [xC,yCT,yCB] = DrawNode( x0, y0, k, w )

n = 20;
a = linspace(0,2*pi*(1-1/n),n);

x = w.width*cos(a)/2 + x0;
y = w.width*sin(a)/2 + y0;
patch(x,y,'w');
text(x0,y0,sprintf('%d',k),'fontname',w.fontName,'fontsize',w.
   fontSize,'horizontalalignment','center');

xC  = x0;
yCT = y0 + w.width/2;
yCB = y0 - w.width/2;

%% TreeDiagram>ConnectNode
function ConnectNode( n, nP, w )

x = [n.xC nP.xC];
y = [n.yCT nP.yCB];

line(x,y,'linewidth',w.linewidth,'color',w.linecolor);
```

The builtin in demo in `TreeDiagram`.

## 3.4 Three-Dimensional Box

There are two broad classes of three-dimensional graphics. One is to draw an object, like the earth. The other is to draw large data sets. This recipe plus the following one will show you how to do both.

### 3.4.1 Problem

We want to draw a three-dimensional box.

### 3.4.2 Solution

The function Box3 users the `patch` function to draw the object. An example is shown in Figure 3.4.

### 3.4.3 How It Works

Three-dimensional objects are created from vertices and faces. A vertex is a point in space. You create a list of vertices that are the corners of your 3D object. You then create faces that are lists of vertices. A face with two vertices is a line, one with three vertices is a triangle. A polygon can have as many vertices as you would like. However, at the lowest level, graphics processors deal with triangles so you are better off making all patches triangles. You will notice the normal

***Figure 3.4:*** A box drawn with `patch`.

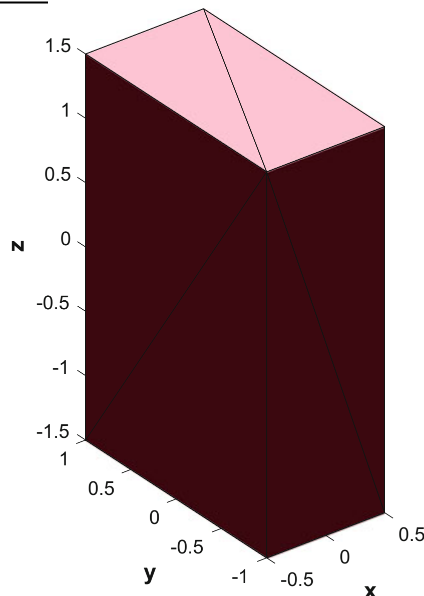

***Figure 3.5:*** A patch. The normal is toward the camera or the "outside" of the object.

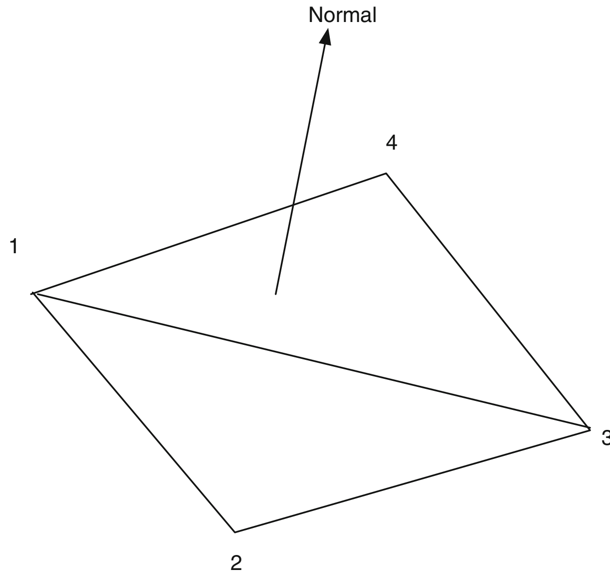


vector. This is the outward vector. Your vertices in your patches should be ordered using the right-hand rule, that is, if the normal is in the direction of your thumb, then the faces are ordered in the direction of your fingers. In this figure, the order for the two triangles would be:

```
[3 2 1]
[1 4 3]
```

MATLAB lighting is not very picky about vertex ordering, but if you export a model, then you will need to follow this convention. Otherwise, you can end up with inside-out objects!

The following code creates a box composed of triangle patches. The face and vertex arrays are created by hand. Vertices are one vertex per row so vertex arrays are n by 3. Face arrays are n by m where m is the largest number of vertices per face. In `Box` we work with triangles only. All graphics processors ultimately draw triangles so, if you can, it is best to create objects only with triangles.

```
function [v, f] = Box( x, y, z )

% Demo
if( nargin < 1 )
  Demo
  return
end

% Faces
f   = [2 3 6;3 7 6;3 4 8;3 8 7;4 5 8;4 1 5;2 6 5;2 5 1;1 3 2;1 4 3;5
    6 7;5 7 8];
```

```
% Vertices
v = [-x  x  x -x -x  x  x -x;...
     -y -y  y  y -y -y  y  y;...
     -z -z -z -z  z  z  z  z]'/2;

% Default outputs
if( nargout == 0 )
  DrawVertices( v, f, 'Box' );
  clear v
end
```

The box is drawn using `patch` in the function `DrawVertices`. There is just one call to `patch`. `patch` accepts parameter pairs to specify face and edge coloring and many other characteristics of the patch. Only one color can be specified for a patch. If you wanted a box with different colors on each side, you would need multiple patches. We turn on `rotate3d` so that we can reorient the object with the mouse. `view3` is a standard MATLAB view with the eye looking down a corner of the grid box.

```
NewFigure(name)
patch('vertices',v,'faces',f,'facecolor',[0.8 0.1 0.2]);
axis image
xlabel('x')
ylabel('y')
zlabel('z')
See SimGUI.m and SimGUI.fig.
view(3)
grid on
rotate3d on
```

We use only the most basic lighting. You can add all sorts of lights in your drawing using `light`. Light can be ambient or from a variety of light sources.

## 3.5 Draw a 3D Object with a Texture

### 3.5.1 Problem

We want to draw a planet with a texture.

### 3.5.2 Solution

Use a surface and overlay a texture onto the surface. Figure 3.6 shows an example with a recent image of Pluto using the function Globe.

```
>> Globe

ans =

  Figure (2: Globe) with properties:
```

***Figure 3.6:*** A three-dimensional globe of Pluto.

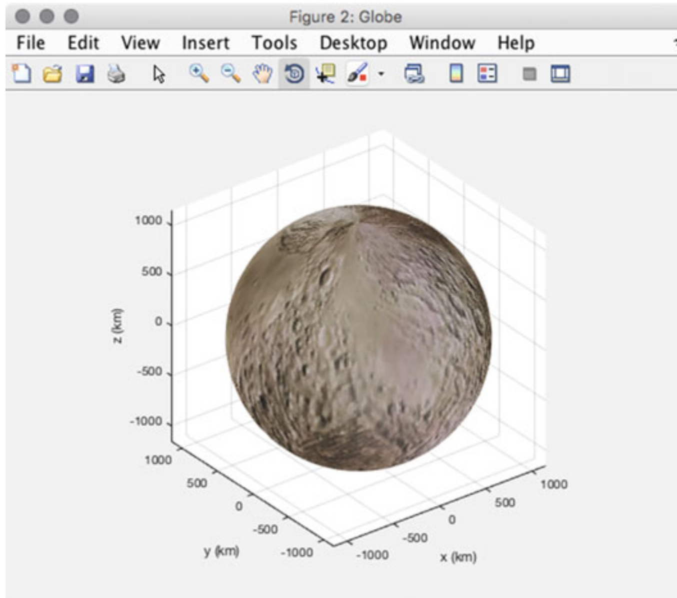


```
     Number: 2
       Name: 'Globe'
      Color: [0.9400 0.9400 0.9400]
   Position: [560 528 560 420]
      Units: 'pixels'

  Show all properties
```

### 3.5.3 How It Works

We generate the picture by first creating $x, y, z$ points on the sphere and then overlaying a texture that is read in from an image file. The texture map can be read from a file using `imread`. If this is color, it will be a three-dimensional matrix. The third element will be an index to the color, red, blue or green. However, if it is a grayscale image, you must create the three-dimensional "color" matrix by replicating the image.

```
p = imread('PlutoGray.png');
p3(:,:,1) = p;
p3(:,:,2) = p;
p3(:,:,3) = p;
```

The starting `p` is a two-dimensional matrix.

You first generate the surface using the coordinates generated from the `sphere` function. This is done with `surface`. You then apply the texture:

```
for i= 1:3
  planetMap(:,:,i)=flipud(planetMap(:,:,i));
end
set(hSurf,'Cdata',planetMap,'Facecolor','texturemap');
set(hSurf,'edgecolor', 'none',...
          'EdgeLighting', 'phong','FaceLighting', 'phong',...
          'specularStrength',0.1,'diffuseStrength',0.9,...
          'SpecularExponent',0.5,'ambientStrength',0.2,...
          'BackFaceLighting','unlit');
```

`flipup` makes the map look "normal." Phong is a type of lighting. It takes the colors at the vertices and interpolates the colors at the pixels on the polygon based on the interpolated normals. Diffuse and specular refer to different types of reflections of light. They aren't too important when you apply a texture to the surface.

# 3.6 General 3D Graphics

## 3.6.1 Problem

We want to use 3D graphics to study a 2D data set. A 2D data set is a matrix or an n by m array.

## 3.6.2 Solution

Use MATLAB surface, mesh, bar and contour functions. TwoDDataDisplay gives an example of a random data set with different visualizations is shown in Figure 3.7.

## 3.6.3 How It Works

We generate a random 2D data set that is 8x8 using `rand`. We display it in several ways in a figure with subplots. In this case, we create two rows and three columns of subplots. Figure 3.7 shows six types of 2D plots. `surf`, `mesh` and `surfl` (3D shaded surface with lighting) are very similar. The surface plots are more interesting when lighting is applied. The two `bar3` plots show different ways of coloring the bars. In the second bar plot, the color varies with length. This requires a bit of code changing the `CData` and `FaceColor`.

```
m = rand(8,8);

h = NewFigure('Two␣Dimensional␣Data');

colormap(h,'gray')

subplot(2,3,1)
surf(m)
title('surf')

subplot(2,3,2)
surfl(m,'light')
```

***Figure 3.7:*** Two-dimensional data shown with six different plot types.

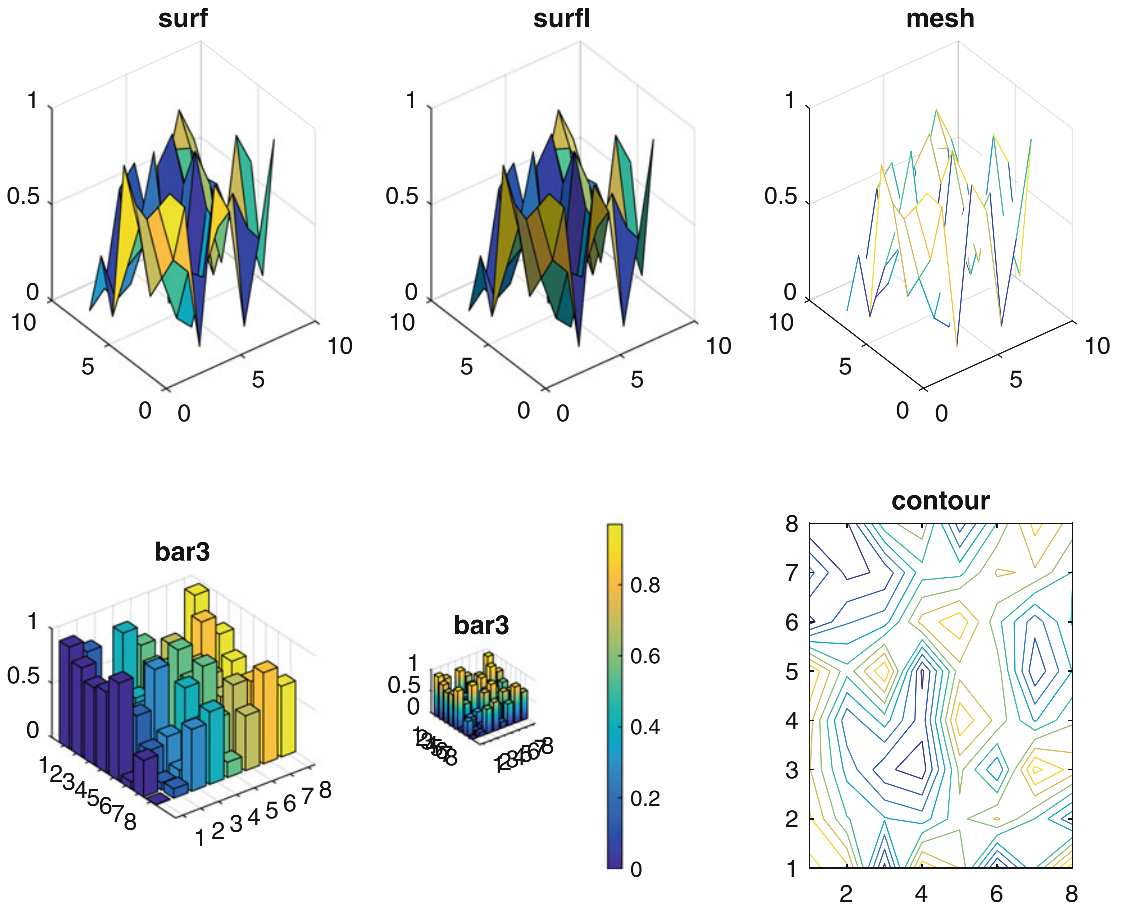


```
title('surfl')

subplot(2,3,3)
mesh(m)
title('mesh')

subplot(2,3,4)
bar3(m)
title('bar3')

subplot(2,3,5)
h = bar3(m);
title('bar3')

colorbar
for k = 1:length(h)
```

```
        zdata = h(k).ZData;
        h(k).CData = zdata;
        h(k).FaceColor = 'interp';
end

subplot(2,3,6)
contour(m);
title('contour')
```

## 3.7 Building a GUI

### 3.7.1 Problem

We want a GUI to provide a graphical interface for a second-order system simulation.

### 3.7.2 Solution

We will use the MATLAB GUIDE to build a GUI that will allow us to:

1. Set the damping constant
2. Set the end time for the simulation
3. Set the type of input (pulse, step or sinusoid)
4. Display the inputs and outputs plot

### 3.7.3 How It Works

We want to build a GUI to interface with `SecondOrderSystemSim` shown below. The first part of SecondOrderSystemSim is the simulation code in a loop.

```
omega    = max([d.omega d.omegaU]); % Maximum frequency for the
   simulation
dT       = 0.1*2*pi/omega; % Get the time step from the frequency
n        = floor(d.tEnd/dT); % Get an integer numbeer of steps
xP       = zeros(2,n); % Size the plotting array
x        = [0;0]; % Initial condition on the [position;velocity]
t        = 0; % Initial time

for k = 1:n
  [~,u]    = RHS(t,x,d);
  xP(:,k)  = [x(1);u];
  x        = RungeKutta( @RHS, t, x, dT, d );
  t        = t + dT;
end
```

***Figure 3.8:*** Second-order system simulation.

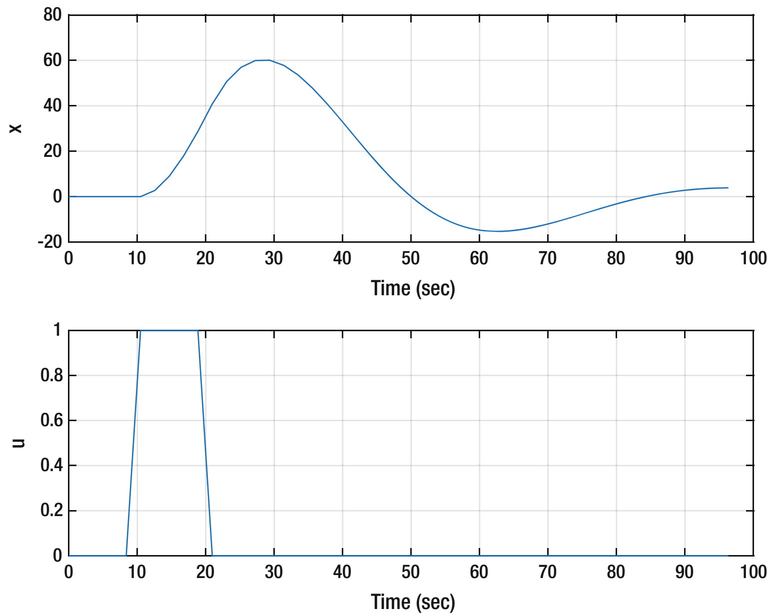


Running it gives the following plot Figure 3.8. The plot code is

```
[t,tL] = TimeLabel((0:n-1)*dT);

if( nargout == 0 )
  PlotSet(t,xP,'x␣label',tL,'y␣label', {'x' 'u'}, 'figure␣title','
      Filter');
end
```

`TimeLabel` makes time units that are reasonable for the length of the simulation. It automatically rescales the time vector. The function has the simulation loop built in.

The MATLAB GUI building system, GUIDE, is invoked by typing `guide` at the command line. We are using MATLAB R2018a. There may be subtle differences in your version.

There are several options for GUI templates, or a blank GUI. We will start from a blank GUI. First, let's make a list of the controls we will need from our desired features list above:

- Edit boxes for:
  - Simulation duration
  - Damping ratio
  - Undamped natural frequency
  - Sinusoid input frequency
  - Pulse start and stop time

***Figure 3.9:*** Blank GUI.

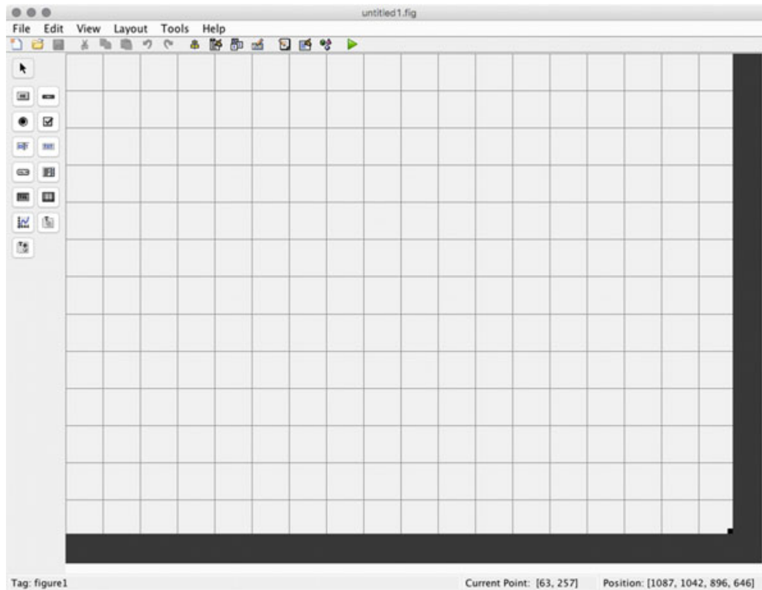


- Radio button for the type of input
- Run button for starting a simulation
- Plot axes

We type "guide" in the command window and it asks us to either pick an existing GUI or create a new one. We choose a blank GUI. Figure 3.9 shows the template GUI in GUIDE before we make any changes to it. You add elements by dragging and dropping from the table at the left.

Figure 3.10 shows the GUI inspector. You edit GUI elements here. You can see that the elements have a lot of properties. We aren't going to try and make this GUI really slick, but with some effort you can make it a work of art. The ones we will change are the tag and text properties. The tag gives the software a name to use internally. The text is just what is shown on the device.

We then add all the desired elements by dragging and dropping. We choose to name our GUI "`GUI`". The resulting initial GUI is shown in Figure 3.11. In the inspector for each element you will see a field for "tag." Change the names from things like `edit1` to names you can easily identify. When you save them and run the GUI from the .fig file the code in `GUI.m` will automatically change.

We create a radio button group and add the radio buttons. This handles disabling all but the selected radio button. When you hit the green arrow in the layout box, it saves all changes to the m-file and also simulates it. It will warn you about bugs.

***Figure 3.10:*** The GUI inspector.

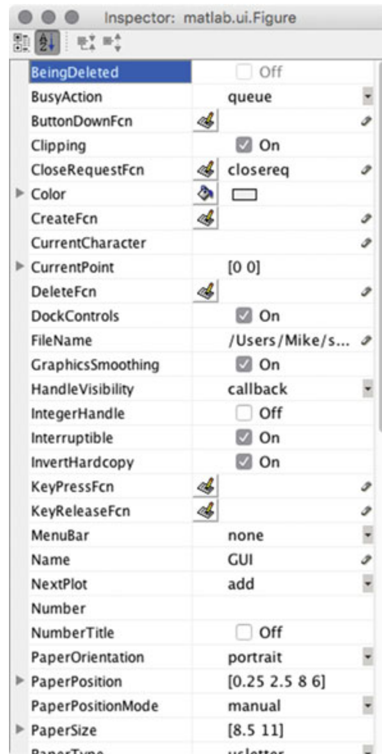


At this point, we can start work on the GUI code itself. The template GUI stores its data, calculated from the data the user types into the edit boxes, in a field called `simdata`. The autogenerated code is in `SimGUI`.

When the GUI loads, we initialize the text fields with the data from the default data structure. Make sure that the initialization corresponds to what is seen in the GUI. You need to be careful about radio buttons and button states.

```
function SimGUI_OpeningFcn(hObject, eventdata, handles, varargin)

% Choose default command line output for SimGUI
handles.output = hObject;

% Get the default data
handles.simData = SecondOrderSystemSim;

% Set the default states
set(handles.editDuration,'string',num2str(handles.simData.tEnd));
set(handles.editUndamped,'string',num2str(handles.simData.omega));
set(handles.editPulseStart,'string',num2str(handles.simData.
    tPulseBegin));
```

***Figure 3.11:*** Snapshot of the GUI in the editing window after adding all the elements.

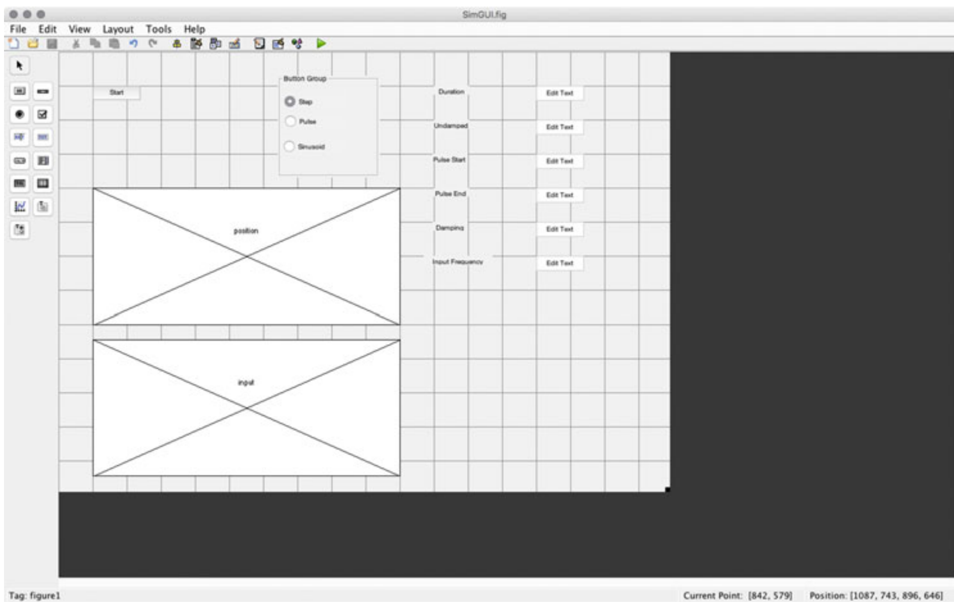


```
set(handles.editPulseEnd,'string',num2str(handles.simData.tPulseEnd))
    ;
set(handles.editDamping,'string',num2str(handles.simData.zeta));
set(handles.editInputFrequency,'string',num2str(handles.simData.
    omegaU));

% Update handles structure
guidata(hObject, handles);
```

When the start button is pushed we run the simulation and plot the results. This essentially is the same as the demo code in the second-order simulation.

```
function start_Callback(hObject, eventdata, handles)

[xP, t, tL] = SecondOrderSystemSim(handles.simData);

axes(handles.position)
plot(t,xP(1,:));
ylabel('Position')
grid

axes(handles.input)
plot(t,xP(2,:));
xlabel(tL);
ylabel('input');
grid
```

***Figure 3.12:*** Snapshot of the GUI in simulation.

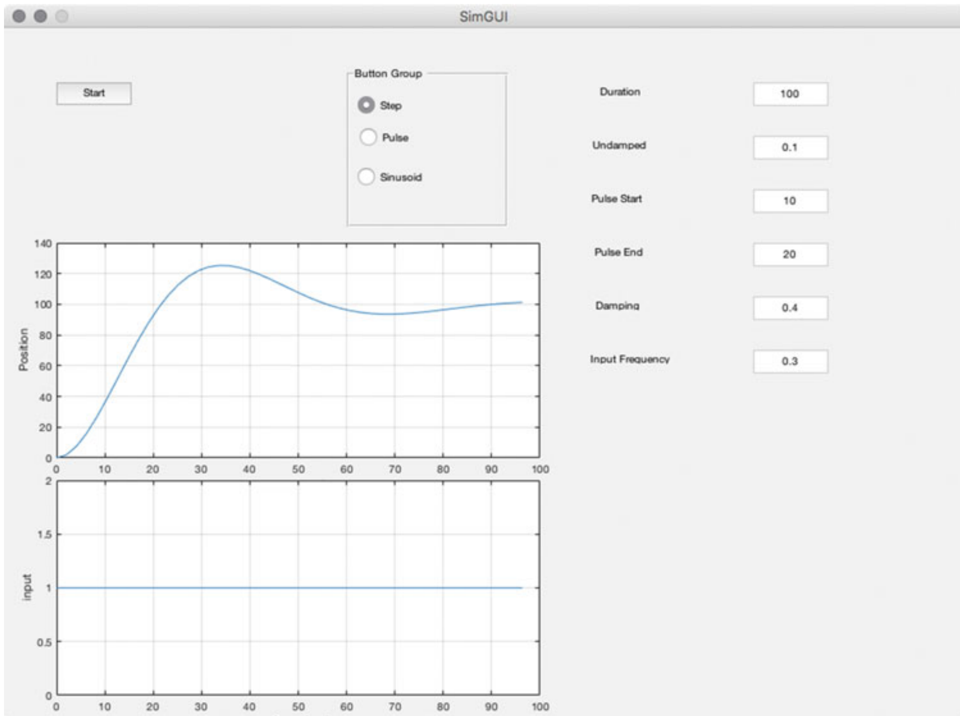


The callbacks for the edit boxes require a little code to set the data in the stored data. All data are stored in the GUI handles. `guidata` must be called to store new data in the handles.

```
function editDuration_Callback(hObject, eventdata, handles)

handles.simData.tEnd = str2double(get(hObject,'String'));
guidata(hObject, handles);
```

One simulation is shown in Figure 3.12. Another simulation in the GUI is shown in Figure 3.13.

## 3.8 Animating a Bar Chart

Two-dimensional arrays are often produced as part of machine-learning algorithms. For situations where they change dynamically we would like to animate a display.

***Figure 3.13:*** Snapshot of the GUI in simulation.

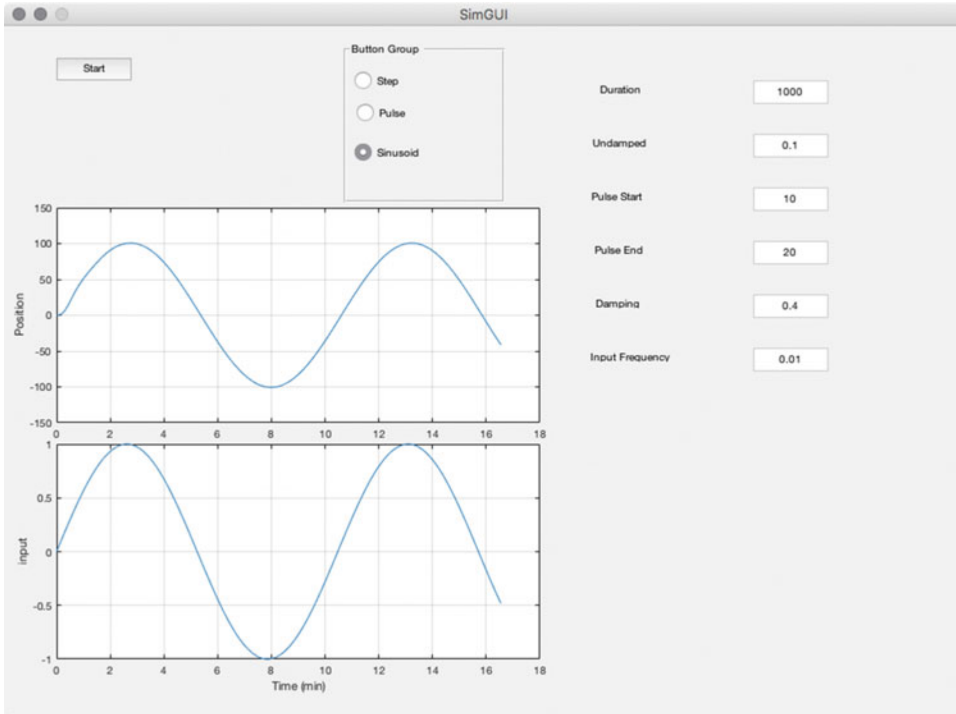


### 3.8.1 Problem

We want to animate a 3D bar chart.

### 3.8.2 Solution

We will write code to animate the MATLAB `bar3` function.

### 3.8.3 How It Works

Our function Bar3D will set up the figure using `bar3` and then replace the values for the length of the bars. This is trickier than it sounds.

The following is an example of `bar3`. We use the handle to get the z data.

```
>> m = [1 2 3;4 5 6];
h = bar3(m);
>> z = get(h(1),'zdata')
z =

   NaN       0       0    NaN
     0       1       1      0
```

```
     0       1       1       0
   NaN       0       0     NaN
   NaN       0       0     NaN
   NaN     NaN     NaN     NaN
   NaN       0       0     NaN
     0       4       4       0
     0       4       4       0
   NaN       0       0     NaN
   NaN       0       0     NaN
   NaN     NaN     NaN     NaN
```

We see each column in the array. We will need to replace all four values for each number in `m`. Look at `h`. It is length 3. Each column in `m` has a `surface` data structure.

```
>> h

h =

  1x3 Surface array:

    Surface     Surface     Surface
```

Figure 3.14 shows the bar graph.

The code is shown below. We have two actions, "initialize," which creates the figure, and "update," which updates the z values. Fortunately, the z-values are always in the same spot so it

***Figure 3.14:*** Two by three bar chart.

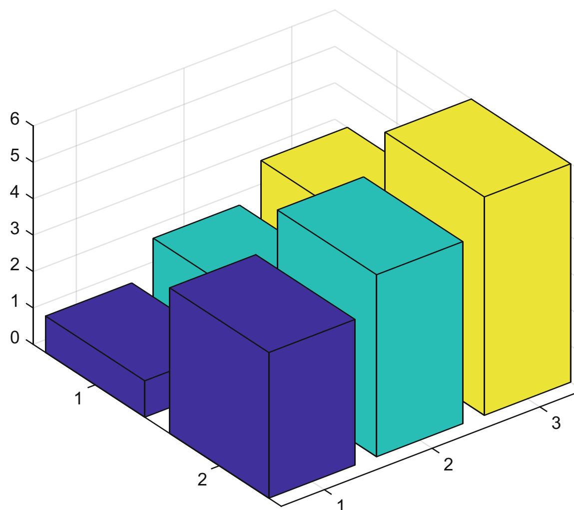

is not too hard to replace them. `colorbar` draws the color bar seen on the right of Figure 3.15. We use `persistent` to store the handle to `bar3`.

```
function Bar3D(action,v,xL,yL,zL,t)

if( nargin < 1 )
  Demo
  return
end

persistent h

switch lower(action)
  case 'initialize'

    NewFigure('3D␣Bar␣Animation');
    h = bar3(v);

    colorbar

    xlabel(xL)
    xlabel(yL)
    xlabel(zL)
    title(t);
    view(3)
    rotate3d on

  case 'update'
    nRows = length(h);
    for i = 1:nRows
      z = get(h(i),'zdata');
      n = size(v,1);
      j = 2;
      for k = 1:n
        z(j,   2) = v(k,i);
        z(j,   3) = v(k,i);
        z(j+1,2) = v(k,i);
        z(j+1,3) = v(k,i);
        j        = j + 6;
      end
      set(h(i),'zdata',z);
    end
end
```

***Figure 3.15:*** Two by three bar chart and the end of the animation.

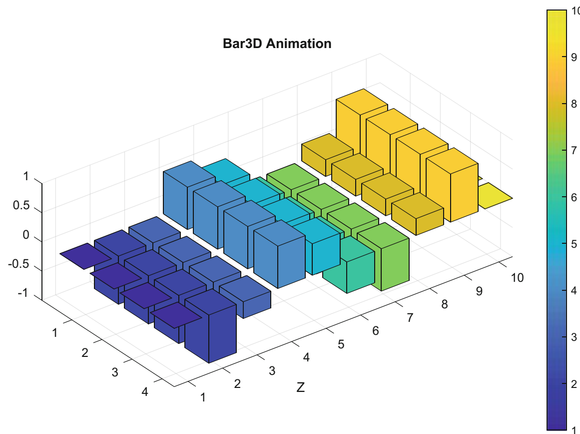


The figure at the end of the animation is shown in Figure 3.15.

## 3.9 Drawing a Robot

This section shows the elements of writing graphics code to draw a robot. If you are doing machine learning involving humans or robots, this is a useful code to have. We'll show how to animate a robot arm.

### 3.9.1 Problem

We want to animate a robot arm.

### 3.9.2 Solution

We write code to create vertices and faces for use in the MATLAB `patch` function.

### 3.9.3 How It Works

`DrawSCARA` draws and animates a robot. The first part of the code really just organizes the operation of the function using a `switch` statement.

```
switch( lower(action) )
    case 'defaults'
        m = Defaults;

    case 'initialize'
        if( nargin < 2 )
            d    = Defaults;
```

```
        else
            d   = x;
        end

        p = Initialize( d );

    case 'update'
        if( nargout == 1 )
            m = Update( p, x );
        else
            Update( p, x );
        end
end
```

Initialize creates the vertex and faces using functions Box, Frustrum and UChannel. These are tedious to write and are geometry-specific. You can apply them to a wide variety of problems, however. You should note that it stores the patches so that we just have to pass in new vertices when animating the arm. The "new" vertices are just the vertices of the arm rotated and translated to match the position of the arm. The arm itself does not deform. We do the computations in the right order so that transformations are passed up/down the chain to get everything moving correctly.

Update updates the arm positions by computing new vertices and passing them to the patches. drawnow draws the arm. We can also save the frames to animate it using MATLAB's movie functions.

```
function m = Update( p, x )

for k = 1:size(x,2)

    % Link 1
    c       = cos(x(1,k));
    s       = sin(x(1,k));

    b1      = [c -s 0;s c 0;0 0 1];
    v       = (b1*p.v1')';

    set(p.link1,'vertices',v);

    % Link 2
    r2      = b1*[p.a1;0;0];

          c       = cos(x(2,k));
    s       = sin(x(2,k));
```

```
    b2       = [c -s 0;s c 0;0 0 1];
    v        = (b2*b1*p.v2')';

    v(:,1)   = v(:,1) + r2(1);
    v(:,2)   = v(:,2) + r2(2);

    set(p.link2,'vertices',v);

    % Link 3
    r3       = b2*b1*[p.r3;0;0] + r2;
    v        = p.v3;

    v(:,1)   = v(:,1) + r3(1);
    v(:,2)   = v(:,2) + r3(2);
    v(:,3)   = v(:,3) + x(3,k);

    set(p.link3,'vertices',v);

    % Link 4
          c        = cos(x(4,k));
    s        = sin(x(4,k));

    b4       = [c -s 0;s c 0;0 0 1];
    v        = (b4*b2*b1*p.v4')';
    r4       = b2*b1*[p.r4;0;0] + r2;

    v(:,1)   = v(:,1) + r4(1);
    v(:,2)   = v(:,2) + r4(2);
    v(:,3)   = v(:,3) + x(3,k);

    set(p.link4,'vertices',v);

    if( nargout > 0 )
        m(k) = getframe;
    else
      drawnow;
    end

end
```

The SCARA robot arm in the demo is shown at the end in Figure 3.16. The demo code could be replaced by a simulation of the arm dynamics. In this case, we pick angular rates and generate an array of angles. Note that this alternate demo function does not need to be a built-in demo function at all. This same block of code can be executed directly from the command line.

***Figure 3.16:*** Robot arm generated by `DrawSCARA`.

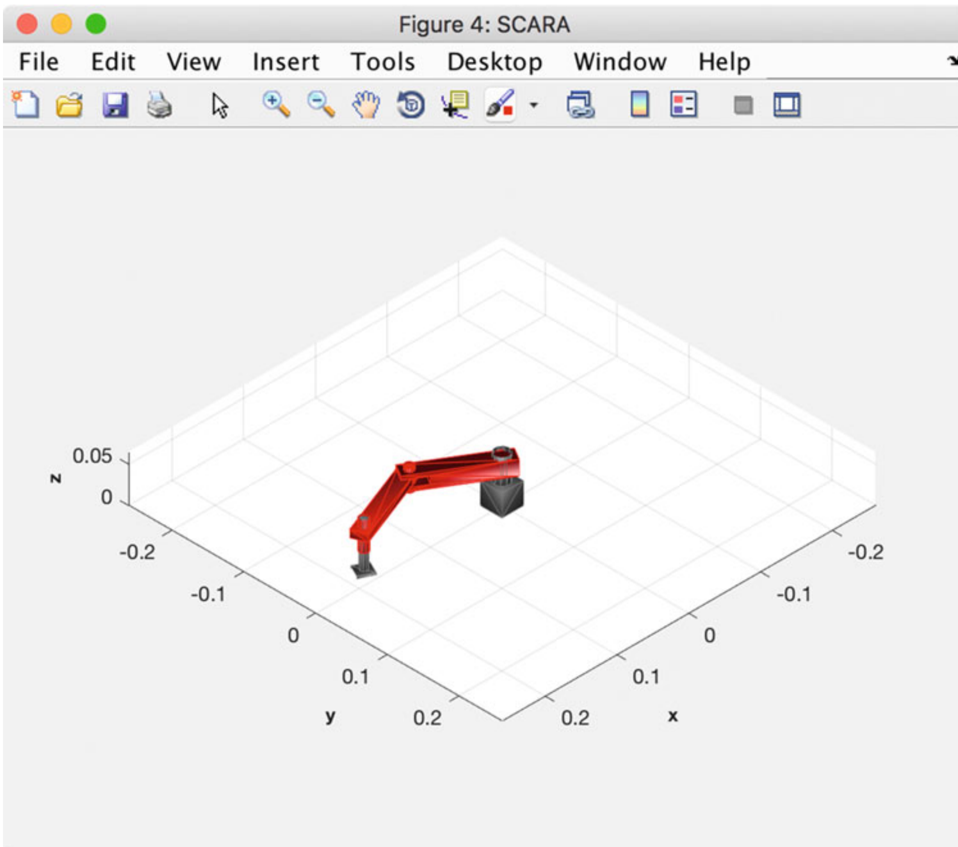


```
function Demo

DrawSCARA( 'initialize' );
t         = linspace(0,100);
omega1   = 0.1;
omega2   = 0.2;
omega3   = 0.3;
omega4   = 0.4;
x        = [sin(omega1*t);sin(omega2*t);0.01*sin(omega3*t);sin(omega4*
    t)];
DrawSCARA( 'update', x );
```

## 3.10 Summary

This chapter has demonstrated graphics that can help you to understand the results of machine learning software. Two- and three-dimensional graphics were demonstrated. The chapter also showed how to build a Graphical User Interface to help you to automate functions. Table 3.1 lists the functions and scripts included in the companion code.

***Table 3.1:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| Bar3D | 3D bar plots |
| Box | Draw a box. |
| DrawSCARA | Draw a robot arm. |
| DrawVertices | Draw a set of vertices and faces. |
| Frustrum | Draw a frustrum (a cone with the top chopped off) |
| Globe | Draw a texture-mapped globe. |
| PlotSet | 2D line plots. |
| SecondOrderSystemSim | Simulates a second-order system. |
| SimGUI | Code for the simulation GUI. |
| SimGUI.fig | The figure. |
| SurfaceOfRevolution | Draw a surface of revolution |
| TreeDiagram | Draw a tree diagram. |
| TwoDDataDisplay | A script to display two-dimensional data in three-dimensional graphics. |
| UChannel | Draw a U shaped channel |

CHAPTER 4

# Kalman Filters

Understanding or controlling a physical system often requires a model of the system, that is, knowledge of the characteristics and structure of the system. A model can be a pre-defined structure or can be determined solely through data. In the case of Kalman Filtering, we create a model and use the model as a framework for learning about the system. This is part of the Control branch of our Autonomous Learning taxonomy from Chapter 1.

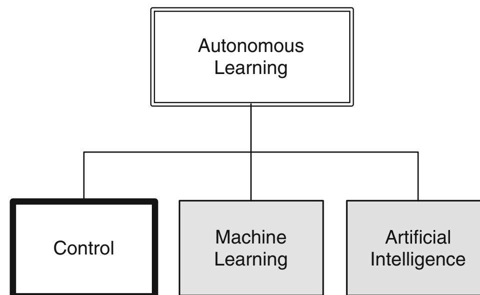


What is important about Kalman Filters is that they rigorously account for uncertainty in a system that you want to know more about. There is uncertainty in the model of the system, if you have a model, and uncertainty (i.e., noise) in measurements of a system.

A system can be defined by its dynamical states and its parameters, which are nominally constant. For example, if you are studying an object sliding on a table, the states would be the position and velocity. The parameters would be the mass of the object and the friction coefficient. There may also be an external force on the object that we may want to estimate. The parameters and states comprise the model. You need to know both to properly understand the system. Sometimes it is hard to decide if something should be a state or a parameter. Mass is usually a parameter, but in a aircraft, car or rocket where the mass changes as fuel is consumed, it is often modeled as a state.

The Kalman Filters, invented by R. E. Kalman and others, are a mathematical framework for estimating or learning the states of a system. An estimator gives you statistically best estimates of the dynamical states of the system, such as the position and velocity of a moving point mass.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_4

Kalman Filters can also be written to identify the parameters of a system. Thus, the Kalman Filter provides a framework for both state and parameter identification.

Another application of Kalman Filters is system identification. System identification is the process of identifying the structure and parameters of any system. For example, with a simple mass on a spring it would be the identification or determination of the mass and spring constant values along with determining the differential equation for modeling the system. It is a form of machine learning that has its origins in control theory. There are many methods of system identification. In this chapter, we will only study the Kalman Filter. The term "learning" is not usually associated with estimation, but it is really the same thing.

An important aspect of the system identification problem is determining what parameters and states can actually be estimated given the measurements that are available. This applies to all learning systems. The question is, can we learn what we need to know about something through our observations? For this, we want to know if a parameter or state is observable and can be independently distinguished. For example, suppose we are using Newton's law:

$$F = ma \tag{4.1}$$

where $F$ is force, $m$ is mass, and $a$ is acceleration as our model, and our measurement is acceleration. Can we estimate both force and mass? The answer is no, because we are measuring the *ratio* of force to mass

$$a = \frac{F}{m} \tag{4.2}$$

We can't separate the two. If we had a force sensor or a mass sensor we could determine each separately. You need to be aware of this issue in all learning systems, including Kalman Filters.

## 4.1 A State Estimator Using a Linear Kalman Filter

### 4.1.1 Problem

You want to estimate the velocity and position of a mass attached through a spring and damper to a structure. The system is shown in Figure 4.1. $m$ is the mass, $k$ is the spring constant, $c$ is the damping constant, and $f$ is an external force. $x$ is the position. The mass moves in only one direction.

Suppose we had a camera that was located near the mass. The camera would be pointed at the mass during its ascent. This would result in a measurement of the angle between the ground and the boresight of the camera. The angle measurement geometry is shown in Figure 4.2. The angle is measured from an offset baseline.

We want to use a conventional linear Kalman Filter to estimate the state of the system. This is suitable for a simple system that can be modeled with linear equations.

***Figure 4.1:*** Spring-mass-damper system. The mass is on the right. The spring is on the top to the left of the mass. The damper is below.

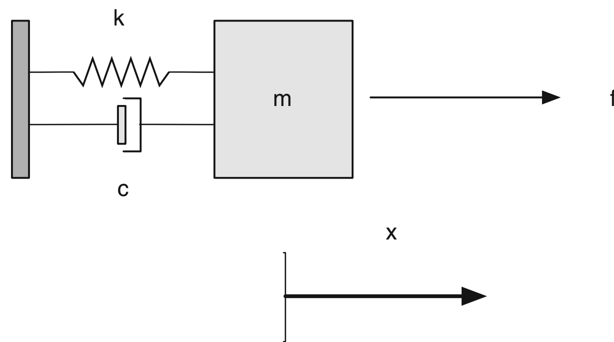


***Figure 4.2:*** The angle measurement geometry.

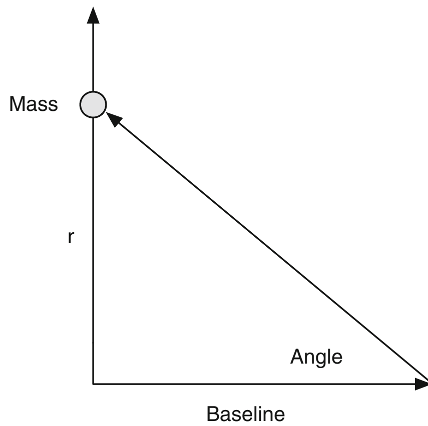


### 4.1.2 Solution

First, we will need to define a mathematical model for the mass system and code it up. Then we will derive the Kalman Filter from first principles, using Bayes theorem. Finally, we present code implementing the Kalman Filter estimator for the spring-mass problem.

### 4.1.3 How It Works

#### Spring-Mass System Model

The continuous time differential equations modeling the system are

$$\frac{dr}{dt} = v \tag{4.3}$$

$$m\frac{dv}{dt} = f - cv - kx \tag{4.4}$$

This says the change in position $r$ with respect to time $t$ is the velocity $v$. The change in velocity with respect to time (times mass) is an external force, minus the damping constant times velocity, minus the spring constant times the position. The second equation is just Newton's law where the total force is $F$ and the total acceleration, $a_T$, is the total force divided by the mass, $\frac{F}{m}$

$$F = f - cv - kx \tag{4.5}$$

$$\frac{dv}{dt} = a_T \tag{4.6}$$

To simplify the problem we divide both sides of the second equation by mass and get:

$$\frac{dr}{dt} = v \tag{4.7}$$

$$\frac{dv}{dt} = a - 2\zeta\omega v - \omega^2 x \tag{4.8}$$

where

$$\frac{c}{m} = 2\zeta\omega \tag{4.9}$$

$$\frac{k}{m} = \omega^2 \tag{4.10}$$

$a$ is the acceleration due to external forces $\frac{f}{m}$, $\zeta$ is the damping ratio, and $\omega$ is the undamped natural frequency. The undamped natural frequency is the frequency at which the mass would oscillate if there were no damping. The damping ratio indicates how fast the system damps and what level of oscillations we observe. With a damping ratio of zero, the system never damps and the mass oscillates forever. With a damping ratio of one you don't see any oscillation. This form makes it easier to understand what damping and oscillation to expect. You immediately know the frequency and the rate at which the oscillation should subside. $m$, $c$, and $k$, although they embody the same information, don't make this as obvious.

The following shows a simulation of the oscillator with damping (`OscillatorDampingRatioSim`). It shows different damping ratios. The loop that runs the simulation with different damping ratios is shown.

```
for j = 1:length(zeta)
  % Initial state [position;velocity]
  x = [0;1];
  % Select damping ratio from array
  d.zeta= zeta(j);

  % Print a string for the legend
  s{j} = sprintf('zeta = %6.4f',zeta(j));
```

```
  for k = 1:nSim
    % Plot storage
    xPlot(j,k)   = x(1);

    % Propagate (numerically integrate) the state equations
    x   = RungeKutta( @RHSOscillator, 0, x, dT, d );
  end
end
```

The results of the damping ratio demo are shown in Figure 4.3. The initial conditions are zero position and a velocity of one. The responses to different levels of damping ratios are seen. When `zeta` is zero it is undamped and oscillates forever. Critical damping, which is desirable from minimizing actuator effort, is 0.7071. A damping ratio of 1 results in no overshoot to a step disturbance. In this case, we have "overshoot," since we are not at a rest initial condition.

***Figure 4.3:*** Spring-mass-damper system simulation with different damping ratios `zeta`.

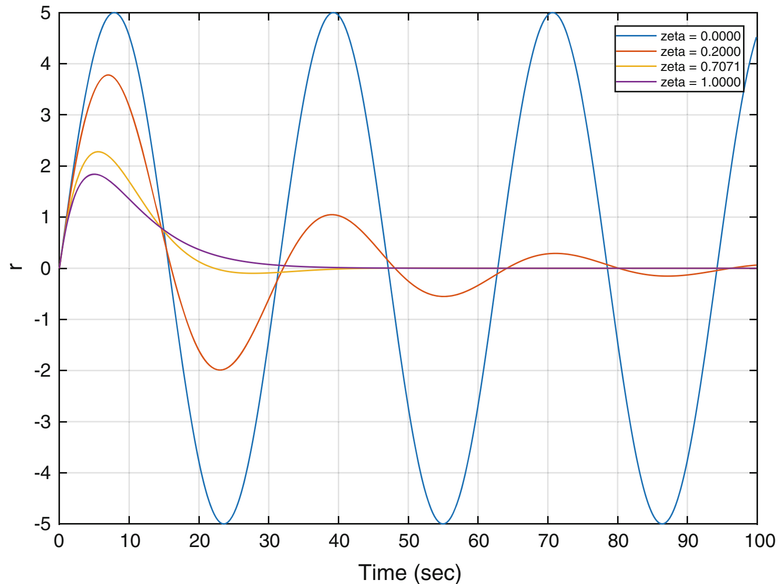


The dynamical equations are in what is called state-space form because the derivative of the state vector:

$$x = \begin{bmatrix} r \\ v \end{bmatrix} \tag{4.11}$$

has nothing multiplying it and there are only first derivatives on the left-hand side. Sometimes you see equations like:

$$Q\dot{x} = Ax + Bu \tag{4.12}$$

If $Q$ is not invertible then you can't do:

$$\dot{x} = Q^{-1}Ax + Q^{-1}Bu \tag{4.13}$$

to make state space equations. Conceptually, if $Q$ is not invertible, that is the same thing as having fewer than $N$ unique equations (where $N$ is the length of $x$, the number of states).

All of our filter derivations work with dynamical equations in state space form. Also, most numerical integration schemes are designed for sets of first-order differential equations.

The right-hand side for the state equations (first-order differential equations), `RHSOscillator`, is shown in the following listing. Notice that if no inputs are requested, it returns the default data structure. The code, `if (nargin < 1)`, tells the function to return the data structure if no inputs are given. This is a convenient way of making your functions self-documenting and keeping your data structures consistent. The actual working code is just one line.

```
xDot = [x(2);d.a-2*d.zeta*d.omega*x(2)-d.omega^2*x(1)];
```

The following listing gives the simulation script `OscillatorSim`. It causes the right-hand side, `RHSOscillator`, to be numerically integrated using the `RungeKutta` function. We start by getting the default data structure from the right-hand side. We fill it in with our desired parameters. Measurements $y$ are created for each step, including random noise. There are two measurements: position and angle.

The following code shows just the simulation loop of `OscillatorSim`. The angle measurement is just trigonometry. The first measurement line computes the angle, which is a non-linear measurement. The second measures the vertical distance, which is linear.

```
for k = 1:nSim
  % Measurements
  yTheta = atan(x(1)/baseline) + yTheta1Sigma*randn(1,1);
  yR     = x(1) + yR1Sigma*randn(1,1);

  % Plot storage
  xPlot(:,k) = [x;yTheta;yR];

  % Propagate (numerically integrate) the state equations
  x = RungeKutta( @RHSOscillator, 0, x, dT, dRHS );
end
```

The results of the simulation are shown in Figure 4.4. The input is a disturbance acceleration that goes from zero to its value at time $t = 0$. It is constant for the duration of the simulation. This is known as a step disturbance. This causes the system to oscillate. The magnitude of the oscillation slowly goes to zero because of the damping. If the damping ratio were 1, we would not see any oscillation, as seen in Figure 4.3.

***Figure 4.4:*** Spring-mass-damper system simulation. The input is a step acceleration. The oscillation slowly damps out, that is, it goes to zero over time. The position $r$ develops an offset due to the constant acceleration.

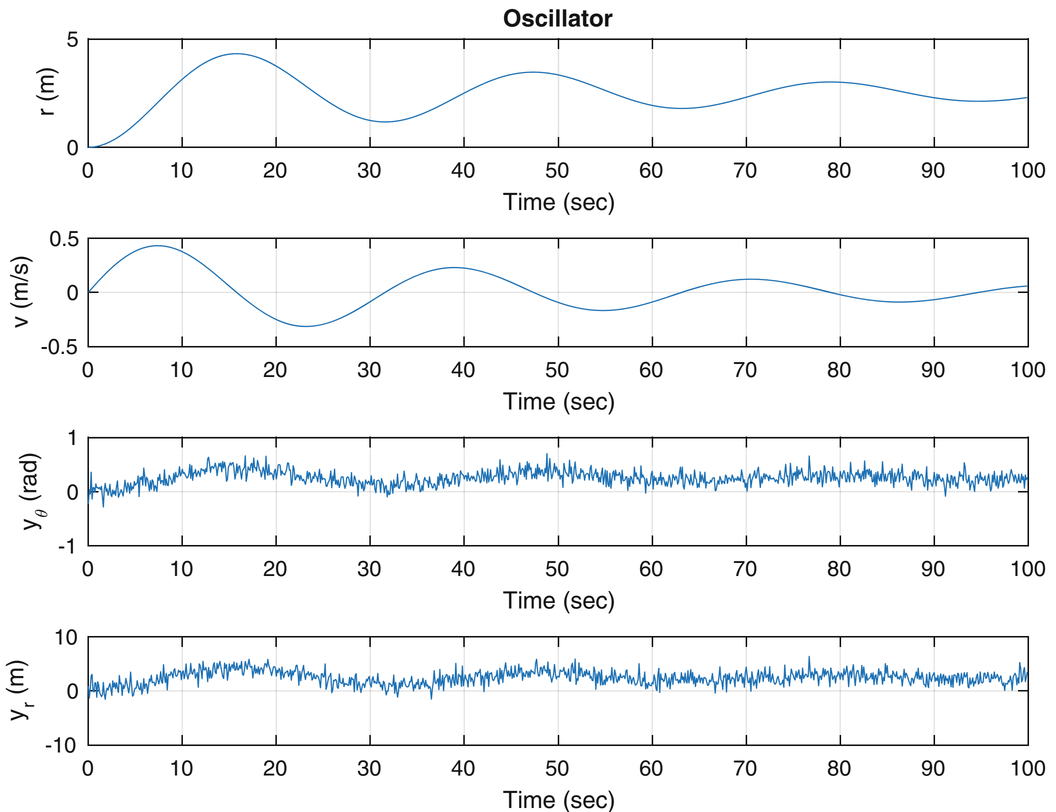


The offset seen in the plot of $r$ can be found analytically by setting $v = 0$. Essentially, the spring force is balancing the external force.

$$0 = \frac{dv}{dt} = a - \omega^2 x \tag{4.14}$$

$$x = \frac{a}{\omega^2} \tag{4.15}$$

We have now completed the derivation of our model and can move on to building the Kalman Filters.

### Kalman Filter Derivation

Kalman filters can be derived from Bayes' Theorem. What is Bayes' Theorem? Bayes' Theorem is:

$$P(A_i|B) = \frac{P(B|A_i)P(A_i)}{\sum P(B|A_i)} \tag{4.16}$$

$$P(A_i|B) = \frac{P(B|A_i)P(A_i)}{P(B)} \tag{4.17}$$

which is just the probability of $A_i$ given $B$. $P$ means "probability." The vertical bar | means "given." This assumes that the probability of $B$ is not zero, that is, $P(B) \neq 0$. In the Bayesian interpretation, the theorem introduces the effect of evidence on belief. This provides a rigorous framework for incorporating any data for which there is a degree of uncertainty. Put simply, given all evidence (or data) to date, Bayes' Theorem allows you to determine how new evidence affects the belief. In the case of state estimation this is the belief in the accuracy of the state estimate.

Figure 4.5 shows the Kalman Filter family and how it relates to the Bayesian Filter. In this book we are covering only the ones in the colored boxes. The complete derivation of the Kalman Filter is given below; this provides a coherent framework for all Kalman filtering implementations. The different filters fall out of the Bayesian models based on assumptions about the model and sensor noise and the linearity or nonlinearity of the measurement and dynamics models. Let's look at the branch that is colored blue. Additive Gaussian noise filters can be linear or nonlinear depending on the type of dynamical and measurement models. In many cases you can take a nonlinear system and linearize it about the normal operating conditions. You can then use a linear Kalman Filter. For example, a spacecraft dynamical model is nonlinear and an Earth sensor that measures the Earth's chord width for roll and pitch information is nonlinear. However, if we are only concerned with Earth pointing, and small deviations from nominal pointing, we can linearize both the dynamical equations and the measurement equations and use a linear Kalman Filter.

If nonlinearities are important, we have to use a nonlinear filter. The Extended Kalman Filter (EKF) uses partial derivatives of the measurement and dynamical equations. These are computed each time step or with each measurement input. In effect, we are linearizing the system each step and using the linear equations. We don't have to do a linear state propagation, that is, propagating the dynamical equations, and could propagate them using numerical integration. If we can get analytical derivatives of the measurement and dynamical equations, this is a reasonable approach. If there are singularities in any of the equations, this may not work.

The Unscented Kalman Filter (UKF) uses the nonlinear equations directly. There are two forms, augmented and non-augmented. In the former, we created an augmented state vector that includes both the states and the state and measurement noise variables. This may result in better results at the expense of more computation.

All of the filters in this chapter are Markov, that is, the current dynamical state is entirely determined from the previous state. Particle filters are not addressed in this book. They are a class of Monte Carlo methods. Monte Carlo (named after the famous casino) methods are computational algorithms that rely on random sampling to obtain results. For example, a Monte

***Figure 4.5:*** The Kalman Filter family tree. All are derived from a Bayesian filter. This chapter covers those in colored boxes.

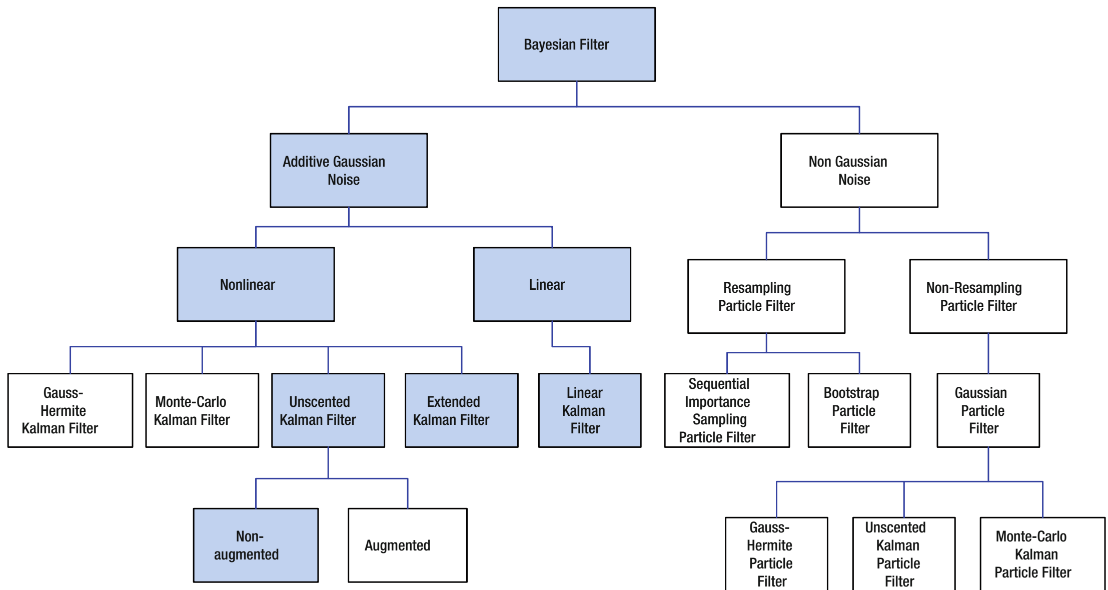


Carlo approach to our oscillator simulation would be to use the MATLAB function `nrandn` to generate the accelerations. We'd run many tests to verify that our mass moves as expected.

Our derivation will use the notation $N(\mu, \sigma^2)$ to represent a normal variable. A normal variable is another word for a Gaussian variable. Gaussian means it is distributed as the normal distribution with mean $\mu$ (average) and variance $\sigma^2$. The following code from `Gaussian` computes a Gaussian or Normal distribution around a mean of 2 for a range of standard deviations. Figure 4.6 shows a plot. The height of the plot indicates how likely a given measurement of the variable is to have that value.

```
%% Initialize
mu              = 2;            % Mean
sigma           = [1 2 3 4]; % Standard deviation
n               = length(sigma);
x               = linspace(-7,10);

%% Simulation
xPlot = zeros(n,length(x));
s     = cell(1,n);

for k = 1:length(sigma)
  s{k}          = sprintf('Sigma = %3.1f',sigma(k));
  f             = -(x-mu).^2/(2*sigma(k)^2);
  xPlot(k,:)  = exp(f)/sqrt(2*pi*sigma(k)^2);
end
```

***Figure 4.6:*** Normal or Gaussian random variable about a mean of 2.

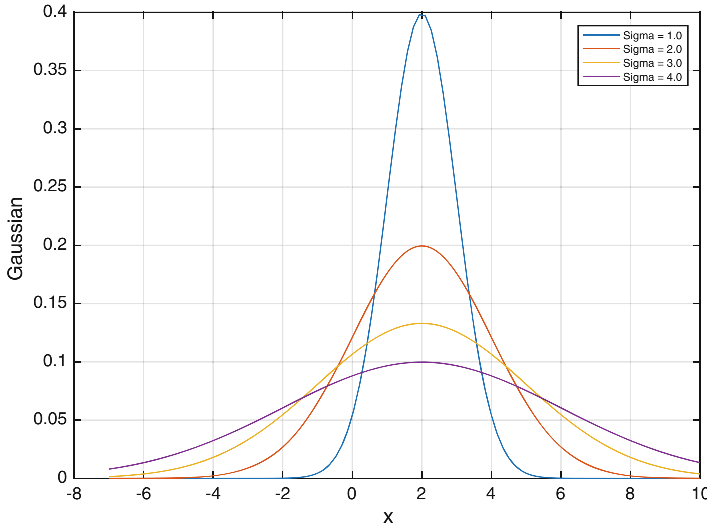


Given the probabilistic state space model in discrete time [23]

$$x_k = f_k(x_{k-1}, w_{k-1}) \tag{4.18}$$

where $x$ is the state vector and $w$ is the noise vector, the measurement equation is:

$$y_k = h_k(x_k, v_n) \tag{4.19}$$

where $v_n$ is the measurement noise. This has the form of a hidden Markov model (HMM) because the state is hidden.

If the process is Markovian, then the future state $x_k$ is dependent only on the current state $x_{k-1}$ and is not dependent on the past states. This can be expressed in the equation:

$$p(x_k|x_{1:k-1}, y_{1:k-1}) = p(x_k|x_{k-1}) \tag{4.20}$$

The $|$ means given. In this case, the first term is read as "the probability of $x_k$ given $x_{1:k-1}$ and $y_{1:k-1}$." This is the probability of the current state given all past states and all measurements up to the $k-1$ measurement. The past, $x_{k-1}$, is independent of the future given the present.

$$p(x_{k-1}|x_{k:T}, y_{k:T}) = p(x_{k-1}|x_k) \tag{4.21}$$

where $T$ is the last sample and the measurements $y_k$ are conditionally independent given $x_k$; that is, they can be determined using only $x_k$ and are not dependent on $x_{1:k}$ or $y_{1:k-1}$. This can be expressed as:

$$p(y_k|x_{1:k}, y_{1:k-1}) = p(y_k|x_k) \tag{4.22}$$

We can define the recursive Bayesian optimal filter that computes the distribution:

$$p(x_k|y_{1:k}) \tag{4.23}$$

given:

- The prior distribution $p(x_0)$, where $x_0$ is the state prior to the first measurement,
- The state space model

$$\begin{aligned} x_k &\sim p(x_k|x_{k-1}) && (4.24)\\ y_k &\sim p(y_k|x_k) && (4.25) \end{aligned}$$

- The measurement sequence $y_{1:k} = y_1, \ldots, y_k$.

Computation is based on the recursion rule

$$p(x_{k-1}|y_{1:k-1}) \rightarrow p(x_k|y_{1:k}) \tag{4.26}$$

This means that we get the current state $x_k$ from the prior state $x_{k-1}$ and all the past measurements $y_{1:k-1}$. Assume that we know the posterior distribution of the previous time step:

$$p(x_{k-1}|y_{1:k-1}) \tag{4.27}$$

The joint distribution of $x_k, x_{k-1}$ given $y_{1:k-1}$ can be computed as:

$$\begin{aligned} p(x_k, x_{k-1}|y_{1:k-1}) &= p(x_k|x_{k-1}, y_{1:k-1})p(x_{k-1}|y_{1:k-1}) && (4.28)\\ &= p(x_k|x_{k-1})p(x_{k-1}|y_{1:k-1}) && (4.29) \end{aligned}$$

because this is a Markov process. Integrating over $x_{k-1}$ gives the prediction step of the optimal filter, which is the Chapman–Kolmogorov equation

$$p(x_k|y_{1:k-1}) = \int p(x_k|x_{k-1}, y_{1:k-1})p(x_{k-1}|y_{1:k-1})dx_{k-1} \tag{4.30}$$

The Chapman–Kolmogorov equation is an identity relating the joint probability distributions of different sets of coordinates on a stochastic process. The measurement update state is found from Bayes' Rule:

$$P(x_k|y_{1:k}) = \frac{1}{C_k}p(y_k|x_k)p(x_k|y_{k-1}) \tag{4.31}$$

$$C_k = p(y_k|y_{1:k-1}) = \int p(y_k|x_k)p(x_k|y_{1:k-1})dx_k \tag{4.32}$$

$C_k$ is the probability of the current measurement, given all past measurements.

If the noise is additive and Gaussian with the state covariance $Q_n$ and the measurement covariance $R_n$, the model and measurement noise have zero mean, and we can write the state equation as:

$$x_k = f_k(x_{k-1}) + w_{k-1} \tag{4.33}$$

where $x$ is the state vector and $w$ is the noise vector. The measurement equation becomes:

$$y_k = h_k(x_k) + v_n \tag{4.34}$$

Given that $Q$ is not time-dependent we can write:

$$p(x_k|x_{k-1}, y_{1:k-1}) = N(x_k; f(x_{k-1}), Q) \tag{4.35}$$

where recall that $N$ is a normal variable, in this case, with mean $x_k; f(x_{k-1})$, which means ($x_k$ given $f(x_{k-1})$ and variance $Q$. We can now write the prediction step Equation 4.30 as:

$$p(x_k|y_{1:k-1}) = \int N(x_k; f(x_{k-1}), Q)p(x_{k-1}|y_{1:k-1})dx_{k-1} \tag{4.36}$$

We need to find the first two moments of $x_k$. A moment is the expected value (or mean) of the variable. The first moment is of the variable, the second is of the variable squared and so forth. They are:

$$E[x_k] = \int x_k p(x_k|y_{1:k-1})dx_k \tag{4.37}$$

$$E[x_k x_k^T] = \int x_k x_k^T p(x_k|y_{1:k-1})dx_k \tag{4.38}$$

$E$ means expected value. $E[x_k]$ is the mean and $E[x_k x_k^T]$ is the covariance. Expanding the first moment and using the identity $E[x] = \int xN(x; f(s), \Sigma)dx = f(s)$ where $s$ is any argument.

$$\begin{aligned}
E[x_k] &= \int x_k \left[\int \mathrm{d}(x_k; f(x_{k-1}), Q)p(x_{k-1}|y_{1:k-1})dx_{k-1}\right] dx_k && (4.39)\\
&= \int x_k \left[\int N(x_k; f(x_{k-1}), Q)dx_k\right] p(x_{k-1}|y_{1:k-1})dx_{k-1} && (4.40)\\
&= \int f(x_{k-1})p(x_{k-1}|y_{1:k-1})dx_{k-1} && (4.41)
\end{aligned}$$

Assuming that $p(x_{k-1}|y_{1:k-1}) = N(x_{k-1}; \hat{x}_{k-1|k-1}, P^{xx}_{k-1|k-1})$ where $P^{xx}$ is the covariance of $x$ and noting that $x_k = f_k(x_{k-1}) + w_{k-1}$ we get:

$$\hat{x}_{k|k-1} = \int f(x_{k-1})N(x_{k-1}; \hat{x}_{k-1|k-1}, P^{xx}_{k-1|k-1})dx_{k-1} \tag{4.42}$$

For the second moment:

$$\begin{aligned}
E[x_k x_k^T] &= \int x_k x_k^T p(x_k|y_{1:k-1})dx_k && (4.43)\\
&= \int \left[\int N(x_k; f(x_{k-1}), Q)x_k x_k^T dx_k\right] p(x_{k-1}|y_{1:k-1})dx_{k-1} && (4.44)
\end{aligned}$$

which results in:

$$P^{xx}_{k|k-1} = Q + \int f(x_{k-1}) f^T(x_{k-1}) N(x_{k-1}; \hat{x}_{k-1|k-1}, P^{xx}_{k-1|k-1}) dx_{k-1} - \hat{x}^T_{k|k-1} \hat{x}_{k|k-1} \tag{4.45}$$

The covariance for the initial state is Gaussian and is $P^{xx}_0$. The Kalman Filter can be written without further approximations as

$$\hat{x}_{k|k} = \hat{x}_{k|k-1} + K_n \left[ y_k - \hat{y}_{k|k-1} \right] \tag{4.46}$$

$$P^{xx}_{k|k} = P^{xx}_{k|k-1} - K_n P^{yy}_{k|k-1} K^T_n \tag{4.47}$$

$$K_n = P^{xy}_{k|k-1} \left[ P^{yy}_{k|k-1} \right]^{-1} \tag{4.48}$$

where $K_n$ is the Kalman gain and $P^{yy}$ is the measurement covariance. The solution of these equations requires the solution of five integrals of the form:

$$I = \int g(x) N(x; \hat{x}, P^{xx}) dx \tag{4.49}$$

The three integrals needed by the filter are:

$$P^{yy}_{k|k-1} = R + \int h(x_n) h^T(x_n) N(x_n; \hat{x}_{k|k-1}, P^{xx}_{k|k-1}) dx_k - \hat{x}^T_{k|k-1} \hat{y}_{k|k-1} \tag{4.50}$$

$$P^{xy}_{k|k-1} = \int x_n h^T(x_n) N(x_n; \hat{x}_{k|k-1}, P^{xx}_{k|k-1}) dx \tag{4.51}$$

$$\hat{y}_{k|k-1} = \int h(x_k) N(x_k; \hat{x}_{k|k-1}, P^{xx}_{k|k-1}) dx_k \tag{4.52}$$

Assume that we have a model of the form:

$$x_k = A_{k-1} x_{k-1} + B_{k-1} u_{k-1} + q_{k-1} \tag{4.53}$$

$$y_k = H_k x_k + r_k \tag{4.54}$$

where

- $x_k \in \Re^n$ is the state of system at time $k$
- $m_k$ is the mean state at time $k$
- $A_{k-1}$ is the state transition matrix at time $k-1$
- $Bk-1$ is the input matrix at time $k-1$
- $uk-1$ is the input at time $k-1$
- $q_{k-1}, N(0, Q_k)$, is the Gaussian process noise at time $k-1$

- $y_k \in \Re^m$ is the measurement at time $k$
- $H_k$ is the measurement matrix at time $k$. This is found from the Jacobian (derivatives) of $h(x)$.
- $r_k = N(0, R_k)$, is the Gaussian measurement noise at time $k$
- The prior distribution of the state is $x_0 = N(m_0, P_0)$ where parameters $m_0$ and $P_0$ contain all prior knowledge about the system. $m_0$ is the mean at time zero and $P_0$ is the covariance. Since our state is Gaussian, this completely describes the state.
- $\hat{x}_{k|k-1}$ is the mean of $x$ at $k$ given $\hat{x}$ at $k-1$
- $\hat{y}_{k|k-1}$ is the mean of $y$ at $k$ given $\hat{x}$ at $k-1$

$\Re^n$ means real numbers in a vector of order $n$, that is, the state has $n$ quantities. In probabilistic terms the model is:

$$p(x_k|x_{k-1}) = N(x_k; A_{k-1}x_{k-1}, Q_k) \quad (4.55)$$

$$p(y_k|x_k) = N(y_k; H_k x_k, R_k) \quad (4.56)$$

The integrals become simple matrix equations. In the following equations, $P_k^-$ means the covariance prior to the measurement update.

$$P_{k|k-1}^{yy} = H_k P_k^- H_k^T + R_k \quad (4.57)$$

$$P_{k|k-1}^{xy} = P_k^- H_k^T \quad (4.58)$$

$$P_{k|k-1}^{xx} = A_{k-1} P_{k-1} A_{k-1}^T + Q_{k-1} \quad (4.59)$$

$$\hat{x}_{k|k-1} = m_k^- \quad (4.60)$$

$$\hat{y}_{k|k-1} = H_k m_k^- \quad (4.61)$$

The prediction step becomes:

$$m_k^- = A_{k-1} m_{k-1} \quad (4.62)$$

$$P_k^- = A_{k-1} P_{k-1} A_{k-1}^T + Q_{k-1} \quad (4.63)$$

The first term in the above covariance equation propagates the covariance based on the state transition matrix, $A$. $Q_{k+1}$ adds to this to form the next covariance. Process noise $Q_{k+1}$ is a measure of the accuracy of the mathematical model, $A$, in representing the system. For example, suppose $A$ was a mathematical model that damped all states to zero. Without $Q$, $P$ would go to zero. But if we really weren't that certain about the model, the covariance would never be less than $Q$. Picking $Q$ can be difficult. In a dynamical system with uncertain disturbances you can compute the standard deviation of the disturbances to compute $Q$. If the model, $A$ were uncertain, then you might do a statistical analysis of the range of models. Or you can try different $Q$ in simulation and see which ones work the best!

The update step is:

$$
\begin{aligned}
v_k &= y_k - H_k m_k^- &(4.64)\\
S_k &= H_k P_k^- H_k^T + R_k &(4.65)\\
K_k &= P_k^- H_k^T S_k^{-1} &(4.66)\\
m_k &= m_k^- + K_k v_k &(4.67)\\
P_k &= P_k^- - K_k S_k K_k^T &(4.68)
\end{aligned}
$$

$S_k$ is an intermediate quantity. $v_k$ is the residual. The residual is the difference between the measurement and your estimate of the measurement given the estimated states. $R$ is just the covariance matrix of the measurements. If the noise is not white, a different filter should be used. White noise has equal energy at all frequencies. Many types of noise, such as the noise from an imager, is not really white noise, but is band limited, that is, it has noise in a limited range of frequencies. You can sometimes add additional states to $A$ to model the noise better, for example, adding a low-pass filter to band limit the noise. This makes $A$ bigger, but is generally not an issue.

### Kalman Filter Implementation

Now we will implement a Kalman Filter estimator for the mass-spring oscillator. First, we need a method of converting the continuous time problem to discrete time. We only need to know the states at discrete times or at fixed intervals, $T$. We use the continuous to discrete transform, which uses the MATLAB `expm` function, which performs the matrix exponential. This transform is coded in `CToDZOH`, the body of which is shown in the following listing. `T` is the sampling period.

```
[n,m] = size(b);
q     = expm([a*T b*T;zeros(m,n+m)]);
f     = q(1:n,1:n);
g     = q(1:n,n+1:n+m);
```

`CToDZOH` includes a demo for a double integrator. A double integrator is a system in which the second derivative of the state is directly dependent upon an external input. In this example, $x$ is the state, representing a position, and $a$ is an external input of acceleration.

$$\frac{d^2r}{dt^2} = a \tag{4.69}$$

Written in state space form it is:

$$
\begin{aligned}
\frac{dr}{dt} &= v &(4.70)\\
\frac{dv}{dt} &= a &(4.71)
\end{aligned}
$$

or in matrix form

$$\dot{x} = Ax + Bu \tag{4.72}$$

where

$$x = \begin{bmatrix} r \\ v \end{bmatrix} \quad (4.73)$$

$$u = \begin{bmatrix} 0 \\ a \end{bmatrix} \quad (4.74)$$

$$A = \begin{bmatrix} 0 & 1 \\ 0 & 0 \end{bmatrix} \quad (4.75)$$

$$B = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad (4.76)$$

To run the demo, simply run `CToDZOH` from the command line without any inputs.

```
>> CToDZOH
Double integrator with a  0.5 second time step.
a =
     0     1
     0     0
b =
     0
     1
f =
    1.0000    0.5000
         0    1.0000
g =
    0.1250
    0.5000
```

The discrete plant matrix $f$ is easy to understand. The position state at step $k+1$ is the state at $k$ plus the velocity at step $k$ multiplied by the time step $T$ of 0.5 s. The velocity at step $k+1$ is the velocity at $k$ plus the time step times the acceleration at step $k$. The acceleration at the time $k$ multiplies $\frac{1}{2}T^2$ to get the contribution to the position. This is just the standard solution to a particle under constant acceleration.

$$r_{k+1} = r_k + Tv_k + \frac{1}{2}T^2 a_k \quad (4.77)$$

$$v_{k+1} = v_k + Ta_k \quad (4.78)$$

In matrix form this is:

$$x_{k+1} = fx_k + bu_k \quad (4.79)$$

With the discrete time approximation, we can change the acceleration every step $k$ to get the time history. This assumes that the acceleration is constant over the period $T$. We need to pick $T$ to be sufficiently small so that this is approximately true if we are to get good results.

The script for testing the Kalman Filter is `KFSim.m`. `KFInitialize` is used to initialize the filter (a Kalman Filter, `'kf'`, in this case). This function has been written to handle multiple

types of Kalman Filters and we will use it again in the recipes for EKF and UKF (`'ekf'` and `'ukf'` respectively). We show it below. This function uses dynamic field names to assign the input values to each field.

The simulation starts by assigning values to all of the variables used in the simulation. We get the data structure from the function `RHSOscillator` and then modify its values. We write the continuous time model in matrix form and then convert it to discrete time. Note that the measurement equation matrix that multiples the state, `h`, is [1 0], indicating that we are measuring the position of the mass. MATLAB's `randn` random number function is used to add Gaussian noise to the simulation. The rest of the script is the simulation loop with plotting afterward.

The first part of the script creates continuous time state space matrices and converts them to discrete time using `CToDZOH`. You then use `KFInitialize` to initialize the Kalman Filter.

```
%% Initialize
tEnd          = 100.0;             % Simulation end time (sec)
dT            = 0.1;               % Time step (sec)
d             = RHSOscillator();   % Get the default data structure
d.a           = 0.1;               % Disturbance acceleration
d.omega       = 0.2;               % Oscillator frequency
d.zeta        = 0.1;               % Damping ratio
x             = [0;0];             % Initial state [position;velocity]
y1Sigma       = 1;                 % 1 sigma position measurement
   noise

% xdot = a*x + b*u
a = [0 1;-2*d.zeta*d.omega -d.omega^2]; % Continuous time model
b = [0;1];                              % Continuous time input
   matrix

% x[k+1] = f*x[k] + g*u[k]
[f,g] = CToDZOH(a,b,dT);   % Discrete time model
xE    = [0.3; 0.1];        % Estimated initial state
q     = [1e-6 1e-6];       % Model noise covariance ;
                           % [1e-6 1e-6] is for low model noise test
                           % [1e-4 1e-4] is for high model noise test
dKF   = KFInitialize('kf','m',xE,'a',f,'b',g,'h',[1 0],...
                     'r',y1Sigma^2,'q',diag(q),'p',diag(xE.^2));
```

The simulation loop cycles through measurements of the state and the Kalman Filter update and prediction state with the code `KFPredict` and `KFUpdate`. The integrator is between the two to get the phasing of the update and prediction correct. You have to be careful to put the predict and update steps in the right places in the script so that the estimator is synchronized with simulation time.

```
%% Simulation
nSim  = floor(tEnd/dT) + 1;
xPlot = zeros(5,nSim);
```

```
for k = 1:nSim
  % Position measurement with random noise
  y = x(1) + y1Sigma*randn(1,1);

  % Update the Kalman Filter
  dKF.y = y;
  dKF   = KFUpdate(dKF);

  % Plot storage
  xPlot(:,k) = [x;y;dKF.m-x];

  % Propagate (numerically integrate) the state equations
  x = RungeKutta( @RHSOscillator, 0, x, dT, d );

  % Propagate the Kalman Filter
  dKF.u = d.a;
  dKF   = KFPredict(dKF);
end
```

The prediction Kalman Filter step, `KFPredict`, is shown in the following listing with an abbreviated header. The prediction propagates the state one time step and propagates the covariance matrix with it. It is saying that when we propagate the state, there is uncertainty, so we must add that to the covariance matrix.

```
%% KFPREDICT Linear Kalman Filter prediction step.

function d = KFPredict( d )

% The first path is if there is no input matrix b
if( isempty(d.b) )
  d.m = d.a*d.m;
else
  d.m = d.a*d.m + d.b*d.u;
end

d.p = d.a*d.p*d.a' + d.q;
```

The update Kalman Filter step, `KFUpdate`, is shown in the following listing. This adds the measurements to the estimate and accounts for the uncertainty (noise) in the measurements.

```
%% KFUPDATE Linear Kalman Filter measurement update step.

function d = KFUpdate( d )

s   = d.h*d.p*d.h' + d.r;        % Intermediate value
k   = d.p*d.h'/s;          % Kalman gain
v   = d.y - d.h*d.m;       % Residual
d.m = d.m + k*v;           % Mean update
d.p = d.p - k*s*k';        % Covariance update
```

You will note that the "memory" of the filter is stored in the data structure `d`. No persistent data storage is used. This makes it easier to use these functions in multiple places in your code. Note also that you don't have to call `KFUpdate` every time step. You need only call it when you have new data. However, the filter does assume uniform time steps.

The script gives two examples for the model noise covariance matrix. Figure 4.7 shows results when high numbers, `[1e-4 1e-4]`, for the model covariance are used. Figure 4.8 shows results when lower numbers, `[1e-6 1e-6]`, are used. We don't change the measurement covariance because only the ratio between noise covariance and model covariance is important.

When the higher numbers are used, the errors are Gaussian but noisy. When the low numbers are used, the result is very smooth, with little noise seen. However, the errors are large in the low model covariance case. This is because the filter is essentially ignoring the measurements, since it thinks the model is very accurate. You should try different options in the script and see how it performs. As you can see, the parameters make a huge difference in how well the filter learns about the states of the system.

***Figure 4.7:*** The Kalman Filter results with the higher model noise matrix, `[1e-4 1e-4]`.

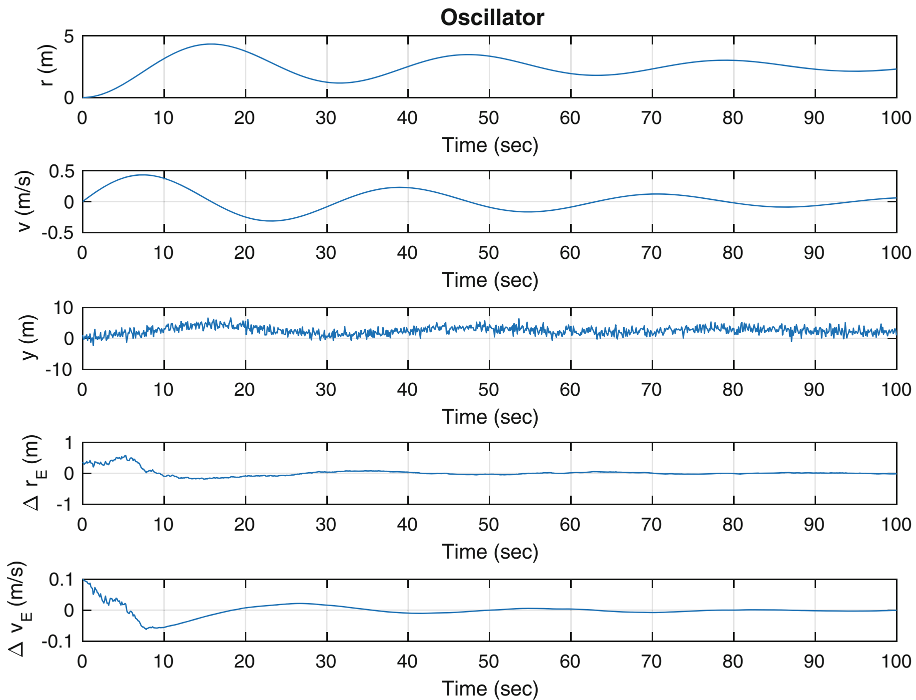

***Figure 4.8:*** The Kalman Filter results with the lower model noise matrix, `[1e-6 1e-6]`. Less noise is seen but the errors are large.

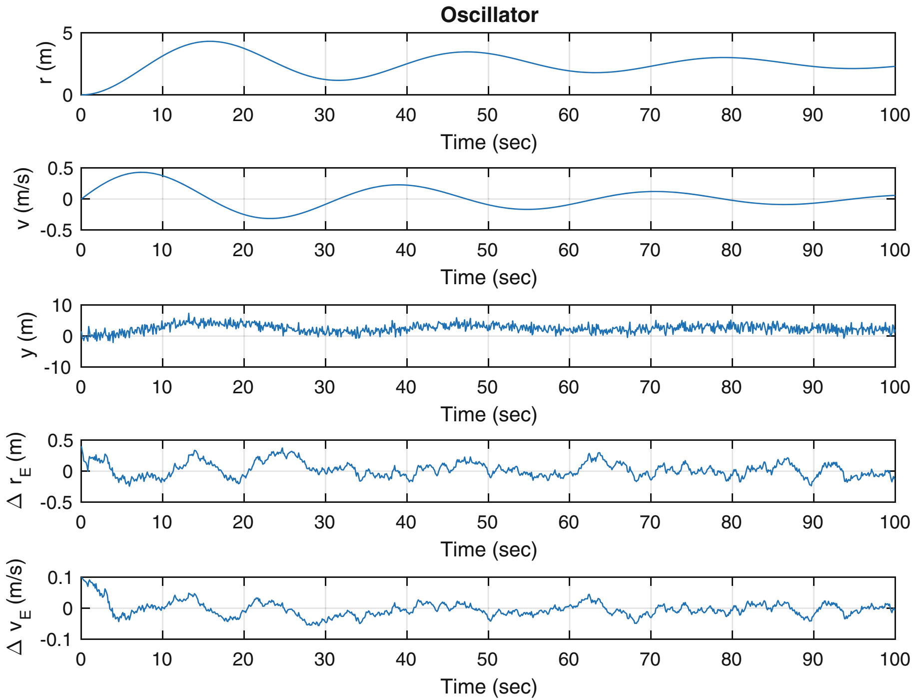


## 4.2 Using the Extended Kalman Filter for State Estimation

### 4.2.1 Problem

We want to track the damped oscillator using an EKF with the nonlinear angle measurement. The EKF was developed to handle models with nonlinear dynamical models and/or nonlinear measurement models. The conventional, or linear, filter requires linear dynamical equations and linear measurements models, that is, the measurement is a linear function of the state. If the model is not linear, linear filters will not track the states very well.

Given a nonlinear model of the form:

$$x_k = f(x_{k-1}, k-1) + q_{k-1} \tag{4.80}$$
$$y_k = h(x_k, k) + r_k \tag{4.81}$$

The prediction step is:

$$m_k^- = f(m_{k-1}, k-1) \tag{4.82}$$
$$P_k^- = F_x(m_{k-1}, k-1) P_{k-1} F_x(m_{k-1}, k-1)^T + Q_{k-1} \tag{4.83}$$

$F$ is the Jacobian of $f$. The update step is:

$$
\begin{aligned}
v_k &= y_k - h(m_k^-, k) & (4.84)\\
S_k &= H_x(m_k^-, k)P_k^- H_x(m_k^-, k)^T + R_k & (4.85)\\
K_k &= P_k^- H_x(m_k^-, k)^T S_k^{-1} & (4.86)\\
m_k &= m_k^- + K_k v_k & (4.87)\\
P_k &= P_k^- - K_k S_k K_k^T & (4.88)
\end{aligned}
$$

$F_x(m, k-1)$ and $H_x(m, k)$ are the Jacobians of the nonlinear functions $f$ and $h$. The Jacobians are just a matrix of partial derivatives of $F$ and $H$. This results in matrices from the vectors $F$ and $H$. For example, assume we have $f(x, y)$, which is:

$$
f = \begin{bmatrix} f_x(x, y) \\ f_y(x, y) \end{bmatrix} \quad (4.89)
$$

The Jacobian is

$$
F_k = \begin{bmatrix} \frac{\partial f_x(x_k, y_k)}{\partial x} & \frac{\partial f_x(x_k, y_k)}{\partial y} \\ \frac{\partial f_y(x_k, y_k)}{\partial x} & \frac{\partial f_y(x_k, y_k)}{\partial y} \end{bmatrix} \quad (4.90)
$$

The matrix is computed at $x_k, y_k$.

The Jacobians can be found analytically or numerically. If done numerically, the Jacobian needs to be computed about the current value of $m_k$. In the Iterated EKF, the update step is done in a loop using updated values of $m_k$ after the first iteration. $H_x(m, k)$ needs to be updated at each step.

### 4.2.2 Solution

We will use the same `KFInitialize` function as created in the previous recipe, but now using the `'ekf'` input. We will need functions for the derivative of the model dynamics, the measurement, and the measurement derivatives. These are implemented in `RHSOscillatorPartial`, `AngleMeasurement`, and `AngleMeasurementPartial`.

We will also need custom versions of the filter to predict and update steps.

### 4.2.3 How It Works

The EKF requires a measurement function, a measurement derivative function, and a state derivative function. The state derivative function computes the $a$ matrix:

$$
x_{k+1} = a_k x_k \quad (4.91)
$$

You would only use the EKF if $a_k$ changed with time. In this problem, it does not. The function to compute $a$ is `RHSOscillatorPartial`. It uses `CToDZOH`. We could have computed $a$ once, but using `CToDZOH` makes the function more general.

```
function a = RHSOscillatorPartial( ~, ~, dT, d )

if( nargin < 1 )
  a = struct('zeta',0.7071,'omega',0.1);
  return
end

b = [0;1];
a = [0 1;d.omega^2 -2*d.zeta*d.omega];
a = CToDZOH( a, b, dT );
```

Our measurement is nonlinear (being an arctangent) and needs to be linearized about each value of position. `AngleMeasurement` computes the measurement, which is nonlinear but smooth.

```
y = atan(x(1)/d.baseline);
```

`AngleMeasurementPartial` computes the derivative. The following function computes the $c$ matrix

$$y_k = c_k x_k \tag{4.92}$$

The partial measurement if found by taking the derivative of the arc-tangent of the angle from the baseline. The comment reminds you of this fact.

```
% y = atan(x(1)/d.baseline);

u   = x(1)/d.baseline;
dH  = 1/(1+u^2);
h   = [dH 0]/d.baseline;
```

It is convenient that the measurement function is smooth. If there were discontinuities, the measurement partials would be difficult to compute. The EKF implementation can handle either functions for the derivatives or matrices. In the case of the functions, we use `feval` to call them. This can be seen in the `EKFPredict` and `EKFUpdate` functions.

`EKFPredict` is the state propagation step for an EKF. It numerically integrates the right-hand side using `RungeKutta`. `RungeKutta` may be overkill in some problems and a simple Euler integration may be appropriate. Euler integration is just:

$$x_{k+1} = x_k + \Delta T f(x, u, t) \tag{4.93}$$

where $f(x, u, t)$ is the right-hand side that can be a function of the state, $x$, time $t$, and the inputs $u$.

```
function d = EKFPredict( d )

% Get the state transition matrix
if( isempty(d.a) )
  a = feval( d.fX, d.m, d.t, d.dT, d.fData );
else
```

```
  a = d.a;
end

% Propagate the mean
d.m = RungeKutta( d.f, d.t, d.m, d.dT, d.fData );

% Propagate the covariance
d.p = a*d.p*a' + d.q;

%% EKFUPDATE Extended Kalman Filter measurement update step.
%% Form
%  d = EKFUpdate( d )
%
%% Description
% All inputs are after the predict state (see EKFPredict). The h
% data field may contain either a function name for computing
% the estimated measurements or an m by n matrix. If h is a function
% name you must include hX which is a function to compute the m by n
% matrix as a linearized version of the function h.
%
%% Inputs
%   d   (.)  EKF data structure
%              .m       (n,1) Mean
%              .p       (n,n) Covariance
%              .h       (m,n) Either a matrix or name/handle of
   function
%              .hX      (*)   Name or handle of Jacobian function for
    h
%              .y       (m,1) Measurement vector
%              .r       (m,m) Measurement covariance vector
%              .hData   (.)   Data structure for the h and hX
   functions
%
%% Outputs
%   d   (.)  Updated EKF data structure
%              .m       (n,1)   Mean
%              .p       (n,n)   Covariance
%              .v       (m,1)   Residuals

function d = EKFUpdate( d )

% Residual
if( isnumeric( d.h ) )
  h   = d.h;
  yE  = h*d.m;
```

```
else
  h    = feval( d.hX, d.m, d.hData );
  yE   = feval( d.h,  d.m, d.hData );
end

% Residual
d.v      = d.y - yE;

% Update step
s    = h*d.p*h' + d.r;
k    = d.p*h'/s;
d.m = d.m + k*d.v;
d.p = d.p - k*s*k';
```

The EKFSim script implements the EKF with all of the above functions as shown in the following listing. The functions are passed to the EKF in the data structure produced by KFInitialize. Note the use of function handles using @, i.e., @RHSOscillator. Notice that KFInitialize requires hX and fX for computing partial derivatives of the dynamical equations and measurement equations.

```
%% Simulation
xPlot = zeros(5,nSim);

for k = 1:nSim
  % Angle measurement with random noise
  y = AngleMeasurement( x, dMeas ) + y1Sigma*randn;

  % Update the Kalman Filter
  dKF.y = y;
  dKF    = EKFUpdate(dKF);

  % Plot storage
  xPlot(:,k)  = [x;y;dKF.m-x];

  % Propagate (numerically integrate) the state equations
  x = RungeKutta( @RHSOscillator, 0, x, dT, d );

  % Propagate the Kalman Filter
  dKF = EKFPredict(dKF);
end
```

Figure 4.9 shows the results. The errors are small. Since the problem dynamics are linear, we don't expect any differences from a conventional Kalman Filter.

***Figure 4.9:*** The Extended Kalman filter tracks the oscillator using the angle measurement.

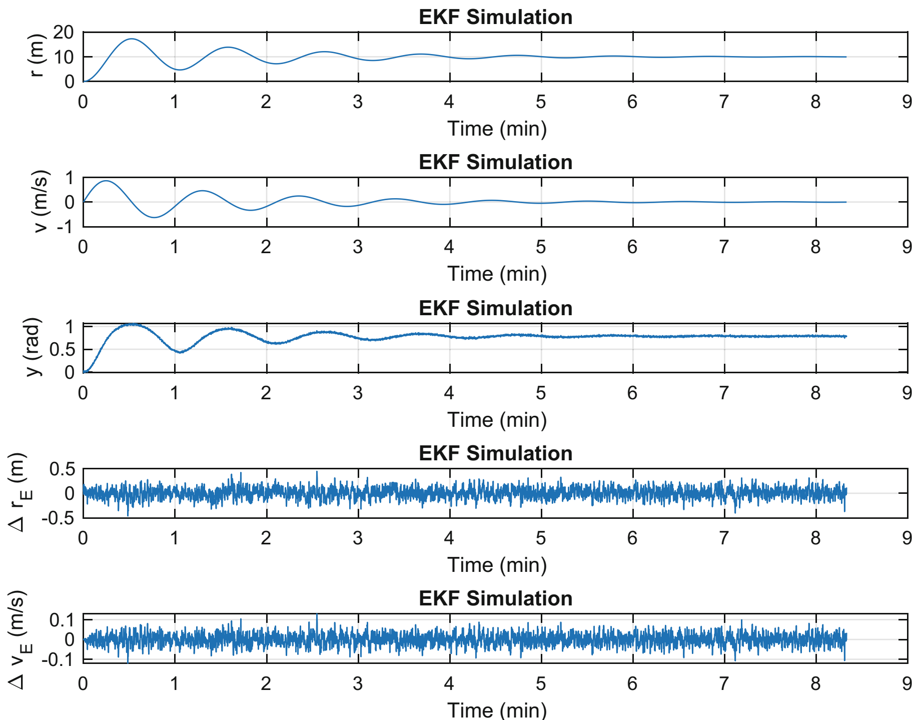


## 4.3 Using the Unscented Kalman Filter for State Estimation

### 4.3.1 Problem

You want to learn the states of the spring, damper, mass system given a nonlinear angle measurement. This time we'll use an UKF. With the UKF, we work with the nonlinear dynamical and measurement equations directly. We don't have to linearize them as we did for the EKF with `RHSOscillatorPartial` and `AngleMeasurementPartial`. The UKF is also known as a sigma $\sigma$ point filter because it simultaneously maintains models one sigma (standard deviation) from the mean.

### 4.3.2 Solution

We will create an UKF as a state estimator. This will absorb measurements and determine the state. It will autonomously learn about the state of the system based on a pre-existing model.

In the following text we develop the equations for the non-augmented Kalman Filter. This form only allows for additive Gaussian noise. Given a nonlinear model of the form

$$x_k = f(x_{k-1}, k-1) + q_{k-1} \quad (4.94)$$

$$y_k = h(x_k, k) + r_k \quad (4.95)$$

Define weights as

$$W_m^0 = \frac{\lambda}{n+\lambda} \tag{4.96}$$

$$W_c^0 = \frac{\lambda}{n+\lambda} + 1 - \alpha^2 + \beta \tag{4.97}$$

$$W_m^i = \frac{\lambda}{2(n+\lambda)}, i = 1, \ldots, 2n \tag{4.98}$$

$$W_c^i = \frac{\lambda}{2(n+\lambda)}, i = 1, \ldots, 2n \tag{4.99}$$

$m$ are weights on the mean state( $m$ for mean) and $c$ weights on the covariances. Note that $W_m^i = W_c^i$.

$$\lambda = \alpha^2(n+\kappa) - n \tag{4.100}$$

$$c = \lambda + n = \alpha^2(n+\kappa) \tag{4.101}$$

$c$ scales the covariances to compute the sigma points, that is, the distribution of points around the mean for computing the additional states to propagate. $\alpha$, $\beta$, and $\kappa$ are scaling constants. General rules for the scaling constants are:

- $\alpha$ – 0 for state estimation, 3 minus the number of states for parameter estimation.
- $\beta$ – Determines the spread of sigma points. Smaller means more closely spaced sigma points.
- $\kappa$ – Constant for prior knowledge. Set to 2 for Gaussian processes.

$n$ is the order of the system. The weights can be put into matrix form:

$$w_m = \left[W_m^0 \cdots W_m^{2n}\right]^T \tag{4.102}$$

$$W = (I - [w_m \cdots w_m]) \begin{bmatrix} W_c^0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & W_c^{2n} \end{bmatrix} (I - [w_m \cdots w_m])^T \tag{4.103}$$

$I$ is the $2n+1$ by $2n+1$ identity matrix. In the equation vector $w_m$ is replicated $2n+1$ times. $W$ is $2n+1$ by $2n+1$.

The prediction step is:

$$X_{k-1} = \begin{bmatrix} m_{k-1} & \cdots & m_{k-1} \end{bmatrix} + \sqrt{c} \begin{bmatrix} 0 & \sqrt{P_{k-1}} & -\sqrt{P_{k-1}} \end{bmatrix} \tag{4.104}$$

$$\hat{X}_k = f(X_{k-1}, k-1) \tag{4.105}$$

$$m_k^- = \hat{X}_k w_m \tag{4.106}$$

$$P_k^- = \hat{X}_k W \hat{X}_k^T + Q_{k-1} \tag{4.107}$$

where $X$ is a matrix where its column is the state vector possibly with an added sigma point vector. The update step is:

$$X_k^- = \begin{bmatrix} m_k^- & \cdots & m_k^- \end{bmatrix} + \sqrt{c}\begin{bmatrix} 0 & \sqrt{P_k^-} & -\sqrt{P_k^-} \end{bmatrix} \quad (4.108)$$

$$Y_k^- = h(X_k^-, k) \quad (4.109)$$

$$\mu_k = Y_k^- w_m \quad (4.110)$$

$$S_k = Y_k^- W[Y_k^-]^T + R_k \quad (4.111)$$

$$C_k = X_k^- W[Y_k^-]^T \quad (4.112)$$

$$K_k = C_k S_k^{-1} \quad (4.113)$$

$$m_k = m_k^- + K_k(y_k - \mu_k) \quad (4.114)$$

$$P_k = P_k^- - K_k S_k K_k^T \quad (4.115)$$

$\mu_k$ is a matrix of the measurements in which each column is a copy modified by the sigma points. $S_k$ and $C_k$ are intermediate quantities. The brackets around $Y_k^-$ are just for clarity.

### 4.3.3 How It Works

The weights are computed in `UKFWeight`.

```
%% UKFWEIGHT Unscented Kalman Filter weight calculation
%% Form
%   d = UKFWeight( d )
%
%% Description
% Unscented Kalman Filter weights.
%
% The weight matrix is used by the matrix form of the Unscented
% Transform. Both UKFPredict and UKFUpdate use the data structure
% generated by this function.
%
% The constant alpha determines the spread of the sigma points around
%   x and is usually set to between 10e-4 and 1. beta incorporates
%   prior knowledge of the distribution of x and is 2 for a Gaussian
%   distribution. kappa is set to 0 for state estimation and 3 -
%   number of states for parameter estimation.
%
%% Inputs
%   d   (.)       Data structure with constants
%               .kappa  (1,1)   0 for state estimation, 3-#states for
%                               parameter estimation
%               .m      (:,1)   Vector of mean states
%               .alpha  (1,1)   Determines spread of sigma points
%               .beta   (1,1)   Prior knowledge - 2 for Gaussian
```

```
%
%% Outputs
%  d    (.)       Data structure with constants
%                .w       (2*n+1,2*n+1)   Weight matrix
%                .wM      (1,2*n+1)       Weight array
%                .wC      (2*n+1,1)       Weight array
%                .c       (1,1)           Scaling constant
%                .lambda  (1,1)           Scaling constant
%

function d = UKFWeight( d )

% Compute the fundamental constants
n            = length(d.m);
a2           = d.alpha^2;
d.lambda     = a2*(n + d.kappa) - n;
nL           = n + d.lambda;
wMP          = 0.5*ones(1,2*n)/nL;
d.wM         = [d.lambda/nL               wMP]';
d.wC         = [d.lambda/nL+(1-a2+d.beta) wMP];

d.c          = sqrt(nL);

% Build the matrix
f            = eye(2*n+1) - repmat(d.wM,1,2*n+1);
d.w          = f*diag(d.wC)*f';
```

The prediction UKF step is shown in the following excerpt from `UKFPredict`.

```
%% UKFPREDICT Unscented Kalman Filter measurement update step

function d = UKFPredict( d )

pS      = chol(d.p)';
nS      = length(d.m);
nSig    = 2*nS + 1;
mM      = repmat(d.m,1,nSig);
x       = mM + d.c*[zeros(nS,1) pS -pS];

xH      = Propagate( x, d );
d.m     = xH*d.wM;
d.p     = xH*d.w*xH' + d.q;
d.p     = 0.5*(d.p + d.p'); % Force symmetry

%% Propagate each sigma point state vector
function x = Propagate( x, d )
```

```
for j = 1:size(x,2)
  x(:,j) = RungeKutta( d.f, d.t, x(:,j), d.dT, d.fData );
end
```

`UKFPredict` uses `RungeKutta` for prediction, which is done by numerical integration. In effect, we are running a simulation of the model and just correcting the results with the next function, `UKFUpdate`. This gets to the core of the Kalman Filter. It is just a simulation of your model with a measurement correction step. In the case of the conventional linear Kalman Filter, we use a linear discrete time model.

The update UKF step is shown in the following listing. The update propagates the state one time step.

```
%% UKFUPDATE Unscented Kalman Filter measurement update step.

function d = UKFUpdate( d )

% Get the sigma points
pS      = d.c*chol(d.p)';
nS      = length(d.m);
nSig    = 2*nS + 1;
mM      = repmat(d.m,1,nSig);
x       = mM + [zeros(nS,1) pS -pS];
[y, r]  = Measurement( x, d );
mu      = y*d.wM;
s       = y*d.w*y' + r;
c       = x*d.w*y';
k       = c/s;
d.v     = d.y - mu;
d.m     = d.m + k*d.v;
d.p     = d.p - k*s*k';

%% Measurement estimates from the sigma points
function [y, r] = Measurement( x, d )

nSigma = size(x,2);

% Create the arrays
lR  = length(d.r);
y   = zeros(lR,nSigma);
r   = d.r;

for j = 1:nSigma
  f         = feval(d.hFun, x(:,j), d.hData );
  iR        = 1:lR;
  y(iR,j)   = f;
end
```

The sigma points are generated using `chol`. `chol` is Cholesky factorization and generates an approximate square root of a matrix. A true matrix square root is more computationally expensive and the results don't really justify the penalty. The idea is to distribute the sigma points around the mean and `chol` works well. Here is an example that compares the two approaches:

```
>> z = [1 0.2;0.2 2]
z =
    1.0000    0.2000
    0.2000    2.0000
>> b = chol(z)
b =
    1.0000    0.2000
         0    1.4000
>> b*b
ans =
    1.0000    0.4800
         0    1.9600
>> q = sqrtm(z)
q =
    0.9965    0.0830
    0.0830    1.4118
>> q*q
ans =
    1.0000    0.2000
    0.2000    2.0000
```

The square root actually produces a square root! The diagonal of `b*b` is close to `z`, which is all that is important.

The script for testing the UKF, `UKFSim`, is shown below. As noted earlier, we don't need to convert the continuous time model into discrete time as we did for the Kalman Filter and EKF. Instead, we pass the filter the right-hand side of the differential equations. You must also pass it a measurement model, which can be nonlinear. You add `UKFUpdate` and `UKFPredict` function calls to the simulation loop. We start by initializing all parameters. `KFInitialize` takes parameter pairs, after `'ukf'` to initialize the filter. The remainder is the simulation loop and plotting. Initialization requires computation of the weighting matrices after calling `KFInitialize`.

```
%% Initialize
dKF   = KFInitialize( 'ukf','m',xE,'f',@RHSOscillator,'fData',d,...
                      'r',y1Sigma^2,'q',q,'p',p,...
                      'hFun',@AngleMeasurement,'hData',dMeas,'dT',dT);
dKF   = UKFWeight( dKF );
```

We show the simulation loop here:

```
%% Simulation
xPlot = zeros(5,nSim);
```

```
for k = 1:nSim
  % Measurements
  y = AngleMeasurement( x, dMeas ) + y1Sigma*randn;

  % Update the Kalman Filter
  dKF.y = y;
  dKF   = UKFUpdate(dKF);

  % Plot storage
  xPlot(:,k)   = [x;y;dKF.m-x];

  % Propagate (numerically integrate) the state equations
  x = RungeKutta( @RHSOscillator, 0, x, dT, d );

  % Propagate the Kalman Filter
  dKF = UKFPredict(dKF);
end
```

The results are shown in Figure 4.10. The errors $\Delta r_E$ and $\Delta v_E$ are just noise. The measurement goes over a large angle range which would make a linear approximation problematic.

***Figure 4.10:*** The Unscented Kalman filter results for state estimation.

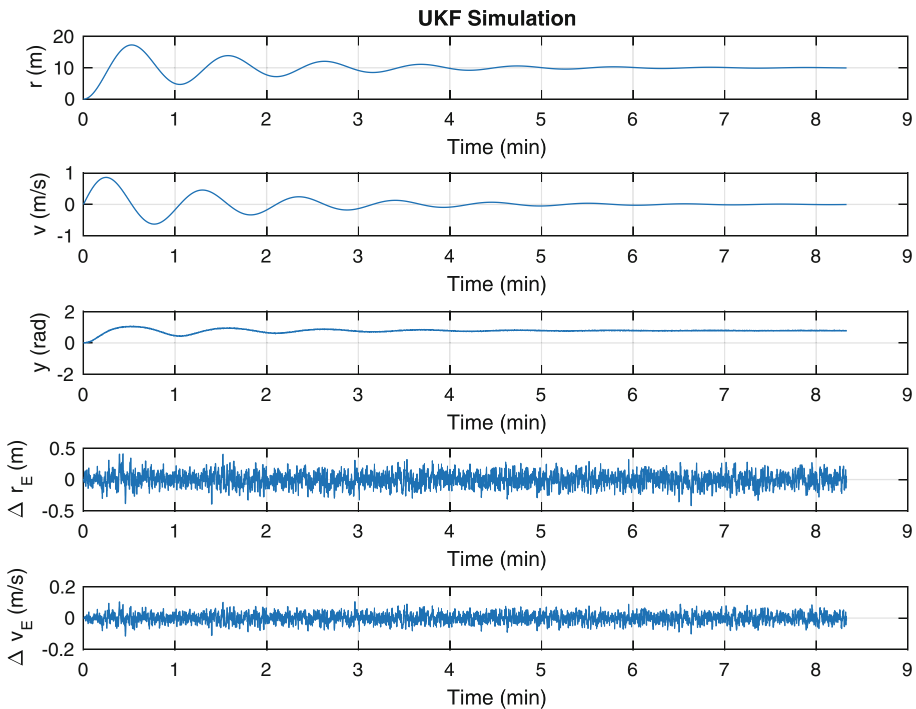

# 4.4 Using the UKF for Parameter Estimation

## 4.4.1 Problem

You want to learn the parameters of the spring-damper-mass system given a nonlinear angle measurement. The UKF can be configured to do this.

## 4.4.2 Solution

The solution is to create a UKF configured as a parameter estimator. This will absorb measurements and determine the undamped natural frequency. It will autonomously learn about the system based on a pre-existing model. We develop the version that requires an estimate of the state that could be generated with a UKF running in parallel, as in the previous recipe.

## 4.4.3 How It Works

Initialize the parameter filter with the expected value of the parameters, $\eta$ [28]

$$\hat{\eta}(t_0) = E\{\hat{\eta}_0\} \tag{4.116}$$

and the covariance for the parameters

$$P_{\eta_o} = E\{(\eta(t_0) - \hat{\eta}_0)(\eta(t_0) - \hat{\eta}_0)^T\} \tag{4.117}$$

The update sequence begins by adding the parameter model uncertainty, $Q$, to the covariance, $P$,

$$P = P + Q \tag{4.118}$$

$Q$ is for the parameters, not the states. The sigma points are then calculated. These are points found by adding the square root of the covariance matrix to the current estimate of the parameters.

$$\eta_\sigma = \left[\begin{array}{ccc} \hat{\eta} & \hat{\eta} + \gamma\sqrt{P} & \hat{\eta} - \gamma\sqrt{P} \end{array}\right] \tag{4.119}$$

$\gamma$ is a factor that determines the spread of the sigma points. We use `chol` for the square root. If there are $L$ parameters, the $P$ matrix is $L \times L$, so this array will be $L \times (2L + 1)$.

The state equations are of the form:

$$\dot{x} = f(x, u, t) \tag{4.120}$$

and the measurement equations are:

$$y = h(x, u, t) \tag{4.121}$$

$x$ is the previous state of the system, as identified by the state estimator or other process. $u$ is a structure with all other inputs to the system that are not being estimated. $\eta$ is a vector of parameters that are being estimated and $t$ is time. $y$ is the vector of measurements. This is the dual estimation approach in that we are not estimating $x$ and $\eta$ simultaneously.

The script UKFPSim for testing the UKF parameter estimation is shown below. We are not doing the UKF state estimation to simplify the script. Normally, you would run the UKF in parallel. We start by initializing all parameters. KFInitialize takes parameter pairs to initialize the filters. The remainder is the simulation loop and plotting. Notice that there is only an update call, since parameters, unlike states, do not propagate.

```
for k = 1:nSim
  % Update the Kalman Filter parameter estimates
  dKF.x = x;

  % Plot storage
  xPlot(:,k) = [y;x;dKF.eta;dKF.p];

  % Propagate (numerically integrate) the state equations
  x = RungeKutta( @RHSOscillator, 0, x, dT, d );

  % Incorporate measurements
  y      = LinearMeasurement( x ) + y1Sigma*randn;
  dKF.y = y;
  dKF    = UKFPUpdate(dKF);
end
```

The UKF parameter update function is shown in the following code. It uses the state estimate generated by the UKF. As noted, we are using the exact value of the state generated by the simulation. This function needs a specialized right-hand side that uses the parameter estimate, d.eta. We modified RHSOscillator for this purpose and wrote RHSOscillatorUKF.

```
function d = UKFPUpdate( d )

d.wA   = zeros(d.L,d.n);
D      = zeros(d.lY,d.n);
yD     = zeros(d.lY,1);

% Update the covariance
d.p = d.p + d.q;

% Compute the sigma points
d = SigmaPoints( d );

% We are computing the states, then the measurements
% for the parameters +/- 1 sigma
for k = 1:d.n
  d.fData.eta = d.wA(:,k);
  x           = RungeKutta( d.f, d.t, d.x, d.dT, d.fData );
  D(:,k)      = feval( d.hFun, x, d.hData );
  yD          = yD + d.wM(k)*D(:,k);
end
```

```
pWD = zeros(d.L,d.lY);
pDD = d.r;
for k = 1:d.n
  wD  = D(:,k) - yD;
  pDD = pDD + d.wC(k)*(wD*wD');
  pWD = pWD + d.wC(k)*(d.wA(:,k) - d.eta)*wD';
end

pDD = 0.5*(pDD + pDD');

% Incorporate the measurements
K       = pWD/pDD;
dY      = d.y - yD;
d.eta   = d.eta + K*dY;
d.p     = d.p - K*pDD*K';
d.p     = 0.5*(d.p + d.p'); % Force symmetry

%% Create the sigma points for the parameters
function d = SigmaPoints( d )

n          = 2:(d.L+1);
m          = (d.L+2):(2*d.L + 1);
etaM       = repmat(d.eta,length(d.eta));
sqrtP      = chol(d.p);
d.wA(:,1)  = d.eta;
d.wA(:,n)  = etaM + d.gamma*sqrtP;
d.wA(:,m)  = etaM - d.gamma*sqrtP;
```

It also has its own weight initialization function `UKFPWeight.m`. The weight matrix is used by the matrix form of the Unscented Transform. The constant alpha determines the spread of the sigma points around the parameter vector and is usually set to between 10e-4 and 1. `beta` incorporates prior knowledge of the distribution of the parameter vector and is 2 for a Gaussian distribution. `kappa` is set to 0 for state estimation and 3 the number of states for parameter estimation.

```
function d = UKFPWeight( d )

d.L           = length(d.eta);
d.lambda      = d.alpha^2*(d.L + d.kappa) - d.L;
d.gamma       = sqrt(d.L + d.lambda);
d.wC(1)       = d.lambda/(d.L + d.lambda) + (1 - d.alpha^2 + d.beta);
d.wM(1)       = d.lambda/(d.L + d.lambda);
d.n           = 2*d.L + 1;
for k = 2:d.n
  d.wC(k) = 1/(2*(d.L + d.lambda));
  d.wM(k) = d.wC(k);
end
```

```
d.wA          = zeros(d.L,d.n);
y             = feval( d.hFun, d.x, d.hData );
d.lY          = length(y);
d.D           = zeros(d.lY,d.n);
```

`RHSOscillatorUKF` is the oscillator model used by the UKF. It has a different input format than `RHSOscillator`. There is only one line of code.

```
xDot = [x(2);d.a-2*d.zeta*d.eta*x(2)-d.eta^2*x(1)];
```

`LinearMeasurement` is a simple measurement function for demonstration purposes. The UKF can use arbitrarily complex measurement functions.

The results of a simulation of an undamped oscillator are shown in Figure 4.11. The filter rapidly estimates the undamped natural frequency. The result is noisy, however. You can explore this script by varying the numbers in the script.

***Figure 4.11:*** The Unscented Kalman parameter estimation results. $p$ is the covariance. It shows that our parameter estimate has converged.

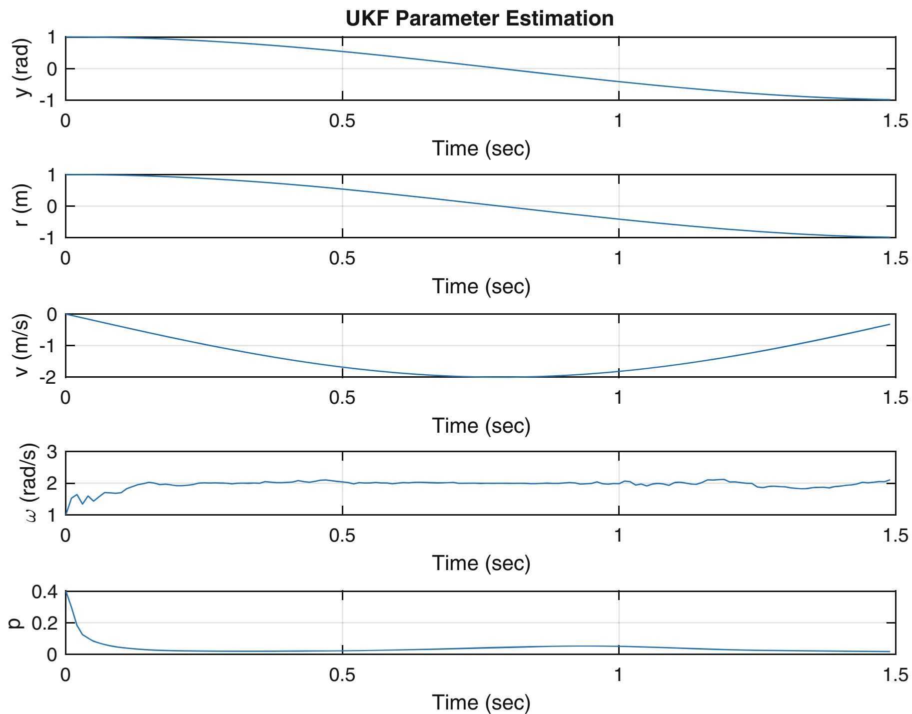

## 4.5 Summary

This chapter has demonstrated learning using Kalman Filters. In this case, learning is the estimation of states and parameters for a damped oscillator. We looked at conventional Kalman Filters and Unscented Kalman Filters. We looked at the parameter learning version of the latter. All examples were done using a damped oscillator. Table 4.1 lists the functions and scripts included in the companion code.

***Table 4.1:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| AngleMeasurement | Angle measurement of the mass. |
| AngleMeasurementPartial | Angle measurement derivative. |
| LinearMeasurement | Position measurement of the mass. |
| OscillatorSim | Simulation of the damped oscillator. |
| OscillatorDampingRatioSim | Simulation of the damped oscillator with different damping ratios. |
| RHSOscillator | Dynamical model for the damped oscillator. |
| RHSOscillatorPartial | Derivative model for the damped oscillator. |
| RungeKutta | Fourth-order Runge–Kutta integrator. |
| PlotSet | Create two-dimensional plots from a data set. |
| TimeLabel | Produce time labels and scaled time vectors. |
| Gaussian | Plot a Gaussian distribution. |
| KFInitialize | Initialize Kalman Filters. |
| KFSim | Demonstration of a conventional Kalman Filter. |
| KFPredict | Prediction step for a conventional Kalman Filter. |
| KFUpdate | Update step for a conventional Kalman Filter. |
| EKFPredict | Prediction step for an Extended Kalman Filter. |
| EKFUpdate | Update step for an Extended Kalman Filter. |
| UKFPredict | Prediction step for an Unscented Kalman Filter. |
| UKFUpdate | Update step for an Unscented Kalman Filter. |
| UKFPUpdate | Update step for an Unscented Kalman Filter parameter update. |
| UKFSim | Demonstration of an Unscented Kalman Filter. |
| UKFPSim | Demonstration of parameter estimation for the Unscented Kalman Filter. |
| UKFWeights | Generates weights for the Unscented Kalman Filter. |
| UKFPWeights | Generates weights for the Unscented Kalman Filter parameter estimator. |
| RHSOscillatorUKF | Dynamical model for the damped oscillator for use in Unscented Kalman Filter parameter estimation. |

# CHAPTER 5

# Adaptive Control

Control systems need to react to the environment in a predictable and repeatable fashion. Control systems take measurements and use them to control the process. For example, a ship measures its heading and changes its rudder angle to attain a desired heading.

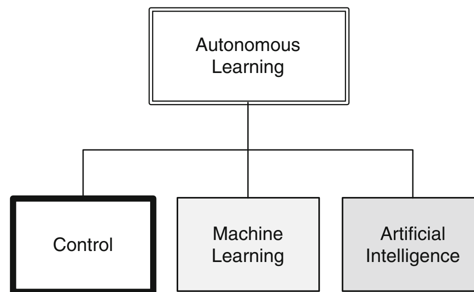


Typically control systems are designed and implemented with all of the parameters hard coded into the software. This works very well in most circumstances, particularly when the system is well known during the design process. When the system is not well defined, or is expected to change significantly during operation, it may be necessary to implement learning control. For example, the batteries in an electric car degrade over time. This leads to less range. An autonomous driving system would need to learn that that range was decreasing. This would be done by comparing the distance traveled with the battery state of charge. More drastic, and sudden, changes can alter a system. For example, in an aircraft the air data system may fail owing to a sensor malfunction. If GPS were still operating, the plane would want to switch to a GPS-only system. In a multi-input-multi-output control system a branch may fail, because of a failed actuator or sensor. The system may have to modify to operating branches in that case.

Learning and adaptive control are often used interchangeably. In this chapter, you will learn a variety of techniques for adaptive control for different systems. Each technique is applied to a different system, but all are generally applicable to any control system.

Figure 5.1 provides a taxonomy of adaptive and learning control. The paths depend on the nature of the dynamical system. The right-most branch is tuning. This is something a designer would do during testing, but it could also be done automatically, as will be described in the

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_5

***Figure 5.1:*** Taxonomy of adaptive or learning control.

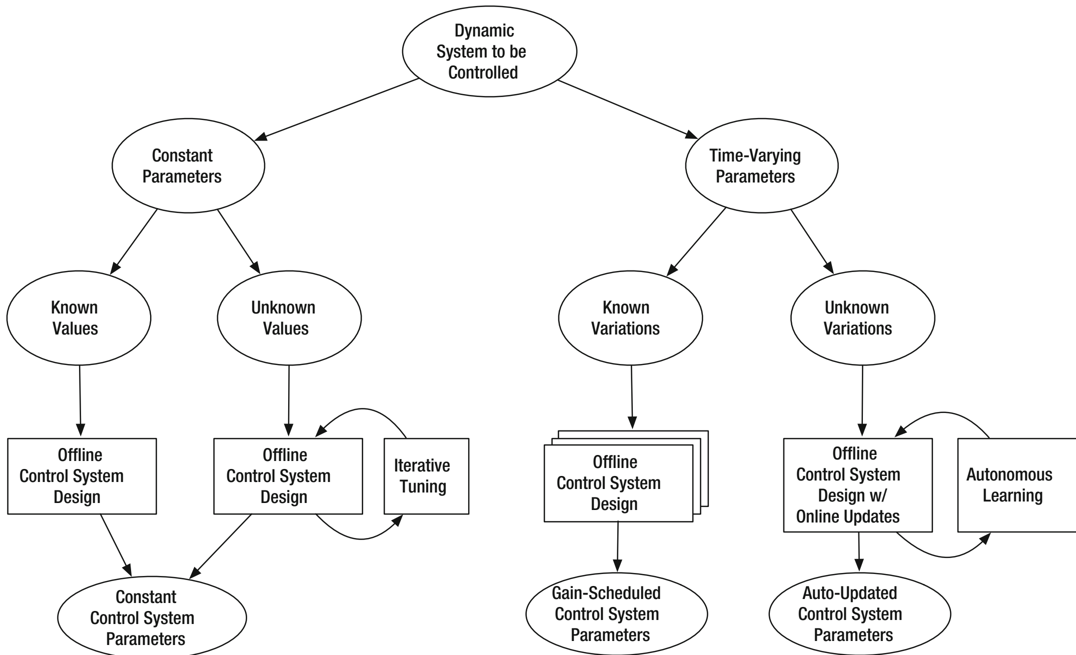


self-tuning Recipe 5.1. The next path is for systems that will vary with time. Our first example of a system with time-varying parameters applies Model Reference Adaptive Control (MRAC) for a spinning wheel. This is discussed in Section 5.3.

The next example is for ship control. Your goal is to control the heading angle. The dynamics of the ship are a function of the forward speed. Although it isn't really learning from experience, it is adapting based on information about its environment.

The last example is a spacecraft with variable inertia. This shows very simple parameter estimation.

## 5.1 Self Tuning: Modeling an Oscillator

We want to tune a damper so that we critically damp a spring system for which the spring constant changes. Our system will work by perturbing the undamped spring with a step and measuring the frequency using a fast Fourier transform (FFT). We then compute the damping using the frequency and add a damper to the simulation. We then measure the undamped natural frequency again to see that it is the correct value. Finally, we set the damping ratio to 1 and observe the response. The system in shown in Figure 5.2.

***Figure 5.2:*** Spring-mass-damper system. The mass is on the right. The spring is on the top to the left of the mass. The damper is below. $F$ is the external force, $m$ is the mass, $k$ is the stiffness, and $c$ is the damping.

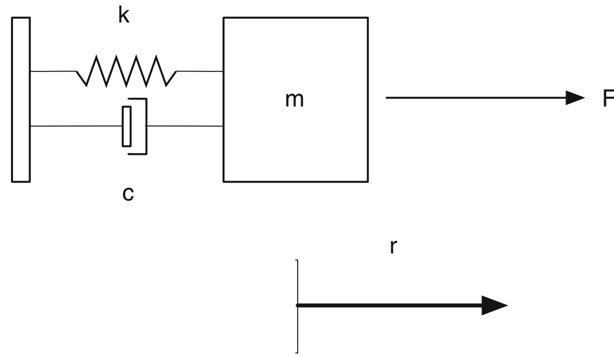


In Chapter 4, we introduced parameter identification in the context of Kalman Filters, which is another way of finding the frequency. The approach here is to collect a large sample of data and process it in batches to find the natural frequency. The equations for the system are:

$$\dot{r} = v \tag{5.1}$$
$$m\dot{v} = -cv - kr \tag{5.2}$$

$c$ is the damping and $k$ is the stiffness. The damping term causes the velocity to go to zero. The stiffness term bounds the range of motion (unless the damping is negative). The dot above the symbols means the first derivative with respect to time, that is:

$$\dot{r} = \frac{dr}{dt} \tag{5.3}$$

The equations state that the change in position with respect to time is the velocity and the mass times the change in velocity with respect to time is equal to a force proportional to its velocity and position. The second equation is Newton's Law:

$$F = ma \tag{5.4}$$

where $F$ is force, $m$ is mass, and $a$ is acceleration.

■ **TIP** Weight is mass times the acceleration of gravity.

$$F = -cv - kr \tag{5.5}$$
$$a = \frac{dv}{dt} \tag{5.6}$$

## 5.2 Self Tuning: Tuning an Oscillator

### 5.2.1 Problem

We want to identify the frequency of an oscillator and tune a control system to that frequency.

### 5.2.2 Solution

The solution is to have the control system measure the frequency of the spring. We will use an FFT to identify the frequency of the oscillation.

### 5.2.3 How It Works

The following script shows how an FFT identifies the oscillation frequency for a damped oscillator.

The function is shown in the following code. We use the `RHSOscillator` dynamical model for the system. We start with a small initial position to get it to oscillate. We also have a small damping ratio so that it will damp out. The resolution of the spectrum is dependent on the number of samples:

$$r = \frac{2\pi}{nT} \tag{5.7}$$

where $n$ is the number of samples and $T$ is the sampling period. The maximum frequency is:

$$\omega = \frac{nr}{2} \tag{5.8}$$

The following shows the simulation loop and `FFTEnergy` call.

```
%% Initialize
nSim          = 2^16;           % Number of time steps
dT            = 0.1;            % Time step (sec)
dRHS          = RHSOscillator;  % Get the default data structure
dRHS.omega    = 0.1;            % Oscillator frequency
dRHS.zeta     = 0.1;            % Damping ratio
x             = [1;0];          % Initial state [position;velocity]
y1Sigma       = 0.000;          % 1 sigma position measurement noise
%% Simulation
xPlot = zeros(3,nSim);

for k = 1:nSim
  % Measurements
  y           = x(1) + y1Sigma*randn;
  % Plot storage
  xPlot(:,k)  = [x;y];
  % Propagate (numerically integrate) the state equations
  x           = RungeKutta( @RHSOscillator, 0, x, dT, dRHS );
end
```

FFTEnergy is shown below.

```
function [e, w, wP] = FFTEnergy( y, tSamp, aPeak )

if( nargin < 3 )
  aPeak  = 0.95;
end

n = size( y, 2 );

% If the input vector is odd drop one sample
if( 2*floor(n/2) ~= n )
  n = n - 1;
  y = y(1:n,:);
end

x  = fft(y);
e  = real(x.*conj(x))/n;

hN = n/2;
e  = e(1,1:hN);
r  = 2*pi/(n*tSamp);
w  = r*(0:(hN-1));

if( nargin > 1 )
  k   = find( e > aPeak*max(e) );
  wP  = w(k);
end
```

The FFT takes the sampled time sequence and computes the frequency spectrum. We compute the FFT using MATLAB's fft function. We take the result and multiply it by its conjugate to get the energy. The first half of the result has the frequency information. aPeak is to indicate peaks for the output. It is just looking for values greater than a certain threshold.

Figure 5.3 shows the damped oscillation. Figure 5.4 shows the spectrum. We find the peak by searching for the maximum value. The noise in the signal is seen at the higher frequencies. A noise-free simulation is shown in Figure 5.5.

The tuning approach is to:

1. Excite the oscillator with a pulse
2. Run it for $2^n$ steps
3. Do an FFT
4. If there is only one peak, compute the damping gain

The script TuningSim calls FFTEnergy.m with aPeak set to 0.7. The value for aPeak is found by looking at a plot and picking a suitable number. The disturbances are Gaussian distributed accelerations and there is noise in the measurement.

***Figure 5.3:*** Simulation of the damped oscillator. The damping ratio, $\zeta$ is 0.5 and undamped natural frequency $\omega$ is 0.1 rad/s.

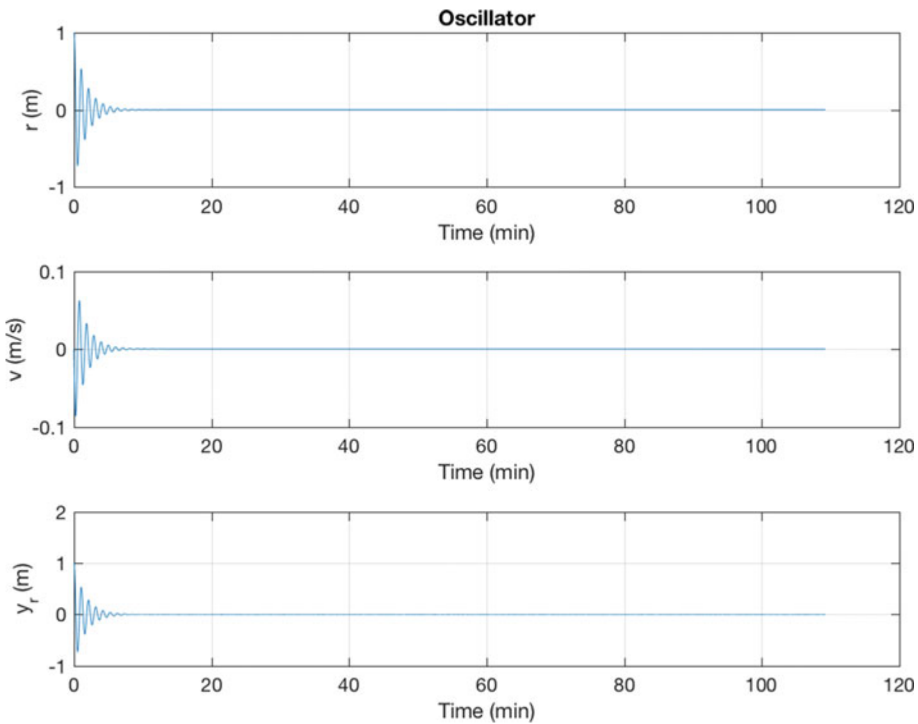


***Figure 5.4:*** The frequency spectrum. The peak is at the oscillation frequency of 0.1 rad/s.

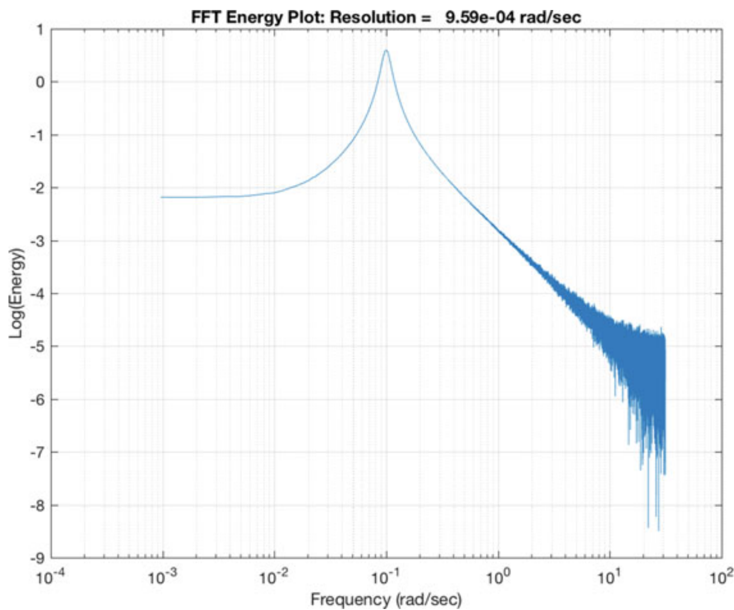

***Figure 5.5:*** The frequency spectrum without noise. The peak of the spectrum is at 0.1 rad/s in agreement with the simulation.

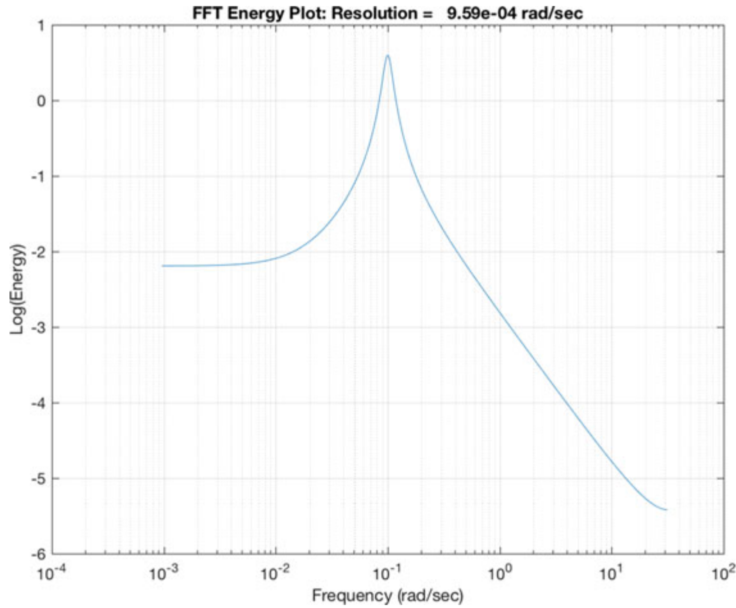


The results in the command window are:

```
TuningSim
Estimated oscillator frequency       0.0997 rad/s
Tuned
Tuned
Tuned
```

As you can see from the FFT plots in Figure 5.6, the spectra are “noisy” owing to the sensor noise and Gaussian disturbance. The criterion for determining that it is underdamped is a distinctive peak. If the noise is large enough we have to set lower thresholds to trigger the tuning. The top left FFT plot shows the 0.1 rad/s peak. After tuning, we damp the oscillator sufficiently so that the peak is diminished. The time plot in Figure 5.6 (the bottom plot) shows that initially the system is lightly damped. After tuning it oscillates very little. There is a slight transient every time the tuning is adjusted at 1.9, 3.6, and 5.5 s. The FFT plots (the top right and middle two) show the data used in the tuning.

An important point is that we must stimulate the system to identify the peak. All system identification, parameter estimation, and tuning algorithms have this requirement. An alternative to a pulse (which has a broad frequency spectrum) would be to use a sinusoidal sweep. That would excite any resonances and make it easier to identify the peak. However, care must be taken when exciting a physical system at different frequencies to ensure that it does not have an unsafe or unstable response at natural frequencies.

***Figure 5.6:*** Tuning simulation results. The first four plots are the frequency spectrums taken at the end of each sampling interval; the last shows the results over time.

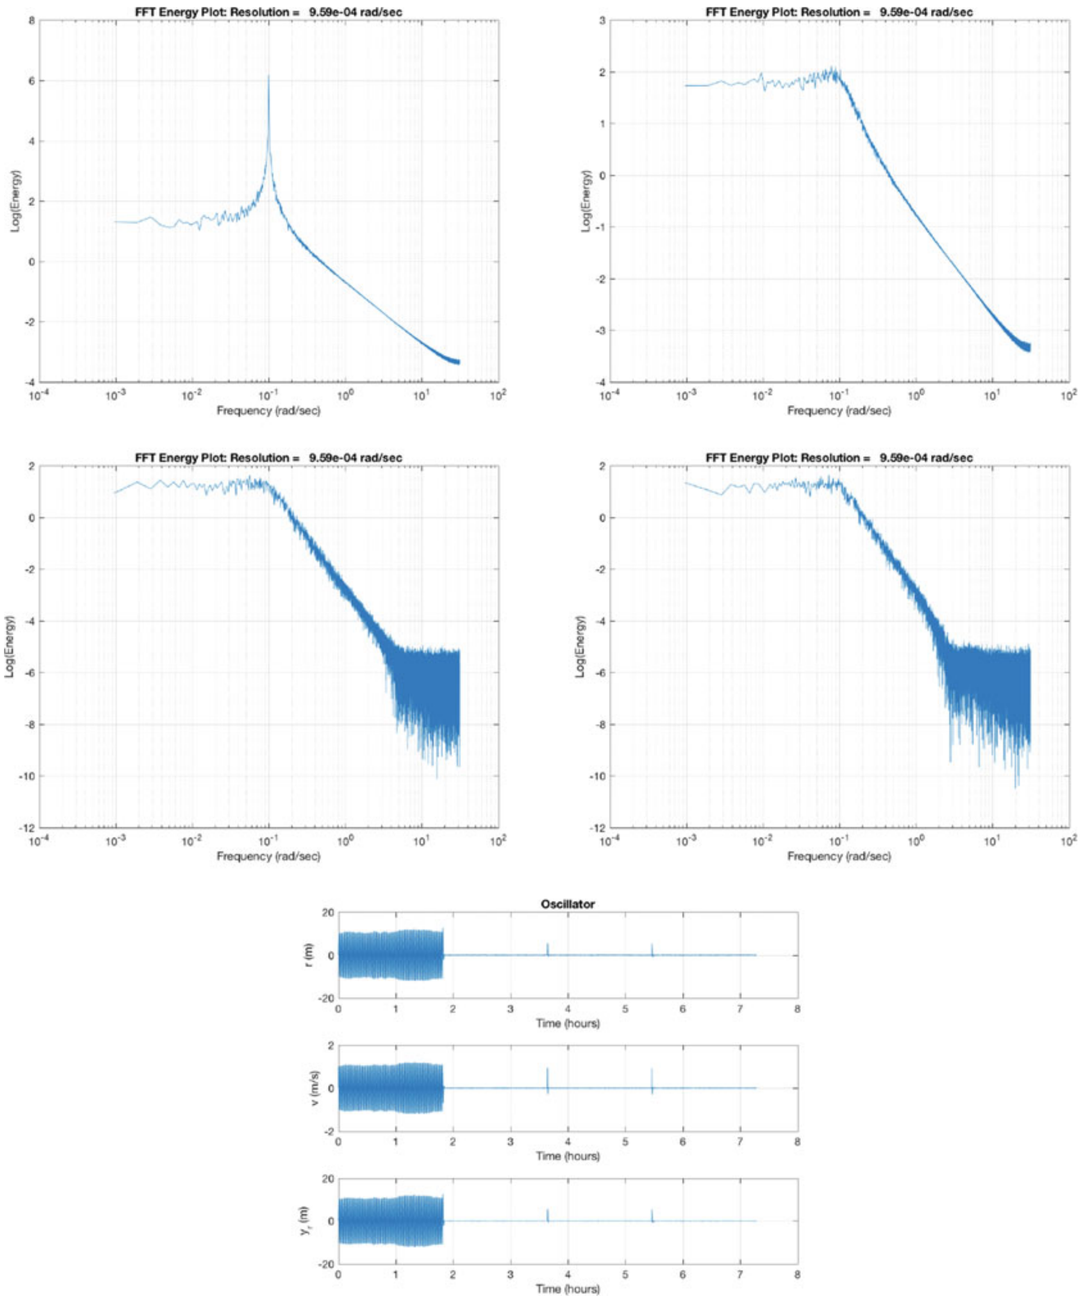

## 5.3 Implement Model Reference Adaptive Control

Our next example is to control a rotor with an unknown load so that it behaves in a desired manner. The dynamical model of the rotary joint is [2]:

$$\frac{d\omega}{dt} = -a\omega + bu_c + u_d \tag{5.9}$$

where the damping $a$ and/or input constants $b$ are unknown. $\omega$ is the angular rate. $u_c$ is the input voltage and $u_d$ is a disturbance angular acceleration. This is a first-order system, which is modeled by one first-order differential equation. We would like the system to behave like the reference model:

$$\frac{d\omega}{dt} = -a_m\omega + b_m u_c + u_d \tag{5.10}$$

***Figure 5.7:*** Speed control of a rotor for the Model Reference Adaptive Control demo.

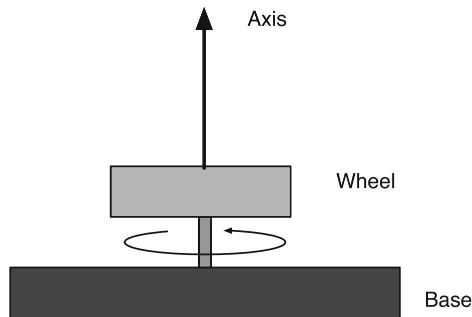


### 5.3.1 Problem

We want to control a system to behave like a particular model. Our example is a simple rotor.

### 5.3.2 Solution

The solution is to implement an MRAC function.

### 5.3.3 How It Works

The idea is to have a dynamical model that defines the behavior of your system. You want your system to have the same dynamics. This desired model is the reference, hence the name Model Reference Adaptive Control. We will use the MIT rule [3] to design the adaptation system. The MIT rule was first developed at the MIT Instrumentation Laboratory (now Draper Laboratory), which developed the NASA Apollo and Space Shuttle guidance and control systems.

Consider a closed-loop system with one adjustable parameter, $\theta$. $\theta$ is a parameter, not an angle. The desired output is $y_m$. The error is:

$$e = y - y_m \tag{5.11}$$

Define a loss function (or cost) as:

$$J(\theta) = \frac{1}{2}e^2 \tag{5.12}$$

The square removes the sign. If the error is zero, the cost is zero. We would like to minimize $J(\theta)$. To make $J$ small, we change the parameters in the direction of the negative gradient of $J$ or:

$$\frac{d\theta}{dt} = -\gamma\frac{\partial J}{\partial \theta} = -\gamma e\frac{\partial e}{\partial \theta} \tag{5.13}$$

This is the MIT rule. If the system is changing slowly, then we can assume that $\theta$ is constant as the system adapts. $\gamma$ is the adaptation gain. Our dynamical model is:

$$\frac{d\omega}{dt} = a\omega + bu_c \tag{5.14}$$

We would like it to be the model:

$$\frac{d\omega_m}{dt} = a_m\omega_m + b_m u_c \tag{5.15}$$

$a$ and $b$ are the actual unknown parameters. $a_m$ and $b_m$ are the model parameters. We would like $a$ and $b$ to be $a_m$ and $b_m$. Let the controller for our rotor be:

$$u = \theta_1 u_c - \theta_2\omega \tag{5.16}$$

The second term provides the damping. The controller has two adaptation parameters. If they are chosen to be:

$$\theta_1 = \frac{b_m}{b} \tag{5.17}$$

$$\theta_2 = \frac{a_m - a}{b} \tag{5.18}$$

the input–output relations of the system and model are the same. This is called perfect model following. This is not required. To apply the MIT rule write the error as:

$$e = \omega - \omega_m \tag{5.19}$$

With the parameters $\theta_1$ and $\theta_2$ the system is:

$$\frac{d\omega}{dt} = -(a + b\theta_2)\omega + b\theta_1 u_c \tag{5.20}$$

where $\gamma$ is the adaptation gain. To continue with the implementation, we introduce the operator $p = \frac{d}{dt}$. We then write:

$$p\omega = -(a + b\theta_2)\omega + b\theta_1 u_c \tag{5.21}$$

or

$$\omega = \frac{b\theta_1}{p + a + b\theta_2} u_c \tag{5.22}$$

We need to get the partial derivatives of the error with respect to $\theta_1$ and $\theta_2$. These are:

$$\frac{\partial e}{\partial \theta_1} = \frac{b}{p + a + b\theta_2} u_c \tag{5.23}$$

$$\frac{\partial e}{\partial \theta_2} = -\frac{b^2\theta_1}{(p + a + b\theta_2)^2} u_c \tag{5.24}$$

from the chain rule for differentiation. Noting that:

$$u_c = \frac{p + a + b\theta_2}{b\theta_1} \omega \tag{5.25}$$

the second equation becomes:

$$\frac{\partial e}{\partial \theta_2} = \frac{b}{p + a + b\theta_2} y \tag{5.26}$$

Since we don't know $a$, let's assume that we are pretty close to it. Then let:

$$p + a_m \approx p + a + b\theta_2 \tag{5.27}$$

Our adaptation laws are now:

$$\frac{d\theta_1}{dt} = -\gamma \left( \frac{a_m}{p + a_m} u_c \right) e \tag{5.28}$$

$$\frac{d\theta_2}{dt} = \gamma \left( \frac{a_m}{p + a_m} \omega \right) e \tag{5.29}$$

Let:

$$x_1 = \frac{a_m}{p + a_m} u_c \tag{5.30}$$

$$x_2 = \frac{a_m}{p + a_m} \omega \tag{5.31}$$

which are differential equations that must be integrated. The complete set is:

$$\frac{dx_1}{dt} = -a_m x_1 + a_m u_c \tag{5.32}$$
$$\frac{dx_2}{dt} = -a_m x_2 + a_m \omega \tag{5.33}$$
$$\frac{d\theta_1}{dt} = -\gamma x_1 e \tag{5.34}$$
$$\frac{d\theta_2}{dt} = \gamma x_2 e \tag{5.35}$$
$$\tag{5.36}$$

Our only measurement would be $\omega$, which would be measured with a tachometer. As noted before, the controller is

$$u = \theta_1 u_c - \theta_2 \omega \tag{5.37}$$
$$e = \omega - \omega_m \tag{5.38}$$
$$\frac{d\omega_m}{dt} = -a_m \omega_m + b_m u_c \tag{5.39}$$

The MRAC is implemented in the function MRAC shown in its entirety below. The controller has five differential equations, which are propagated. The states are $[x_1, x_2, \theta_1, \theta_2, \omega_m]$. RungeKutta is used for the propagation, but a less computationally intensive lower order integrator, such as Euler, could be used instead. The function returns the default data structure if no inputs and one output is specified. The default data structure has reasonable values. That makes it easier for a user to implement the function. It only propagates one step.

```
function d = MRAC( omega, d )

if( nargin < 1 )
  d = DataStructure;
  return
end

d.x = RungeKutta( @RHS, 0, d.x, d.dT, d, omega );
d.u = d.x(3)*d.uC - d.x(4)*omega;

%% MRAC>DataStructure
function d = DataStructure
% Default data structure

d = struct('aM',2.0,'bM',2.0,'x',[0;0;0;0;0],'uC',0,'u',0,'gamma',1,'
   dT',0.1);

%% MRAC>RHS
function xDot = RHS( ~, x, d, omega )
```

```
% RHS for MRAC

e    = omega - x(5);
xDot = [-d.aM*x(1) + d.aM*d.uC;...
        -d.aM*x(2) + d.aM*omega;...
        -d.gamma*x(1)*e;...
         d.gamma*x(2)*e;...
        -d.aM*x(5) + d.bM*d.uC];
```

Now that we have the MRAC controller done, we'll write some supporting functions and then test it all out in `RotorSim`.

# 5.4 Generating a Square Wave Input

## 5.4.1 Problem

We need to generate a square wave to stimulate the rotor in the previous recipe.

## 5.4.2 Solution

For the purposes of simulation and testing our controller we will generate a square wave with a function.

## 5.4.3 How It Works

`SquareWave` generates a square wave. The first few lines are our standard code for running a demo or returning the data structure.

```
function [v,d] = SquareWave( t, d )

if( nargin < 1 )
  if( nargout == 0 )
    Demo;
  else
    v = DataStructure;
  end
  return
end

if( d.state == 0 )
  if( t - d.tSwitch >= d.tLow )
    v         = 1;
    d.tSwitch = t;
    d.state   = 1;
  else
    v         = 0;
  end
else
  if( t - d.tSwitch >= d.tHigh )
```

```
    v         = 0;
    d.tSwitch = t;
    d.state   = 0;
  else
    v         = 1;
  end
end
```

This function uses `d.state` to determine if it is in the high or low part of a square wave. The width of the low part of the wave is set in `d.tLow`. The width in the high part of the square wave is set in `d.tHigh`. It stores the time of the last switch in `d.tSwitch`.

A square wave is shown in Figure 5.8. There are many ways to specify a square wave. This function produces a square wave with a minimum of zero and maximum of 1. You specify the time at zero and the time at 1 to create the square wave.

***Figure 5.8:*** Square wave.

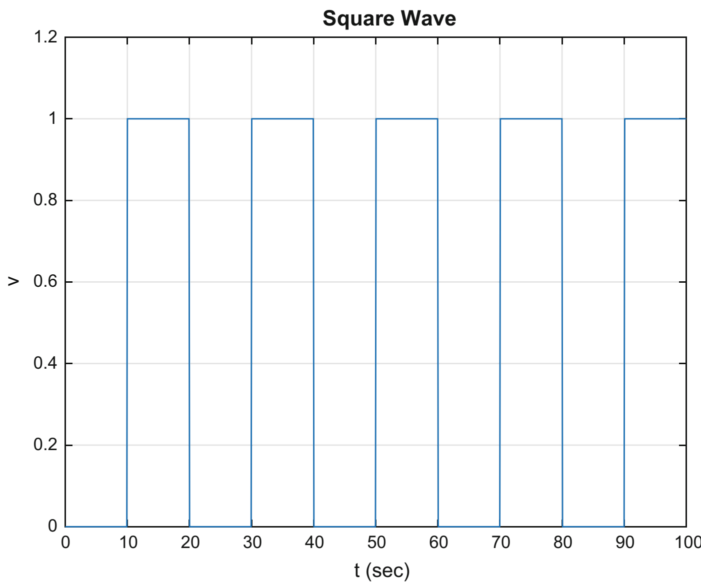


We adjusted the y-axis limit and line width using the code:

```
PlotSet(t,v,'x␣label', 't␣(sec)', 'y␣label', 'v', 'plot␣title','
    Square␣Wave',...
          'figure␣title', 'Square␣Wave');
set(gca,'ylim',[0 1.2])
h = get(gca,'children');
set(h,'linewidth',1);
```

---

■ **TIP** `h = get(gca,'children')` gives you access to the line data structure in a plot for the most recent axes.

---

## 5.5 Demonstrate MRAC for a Rotor

### 5.5.1 Problem

We want to create a recipe to control our rotor using MRAC.

### 5.5.2 Solution

The solution is to use implement our MRAC function in a MATLAB script from Recipe 5.3.

### 5.5.3 How It Works

Model Reference Adaptive Control is implemented in the script `RotorSim`. It calls `MRAC` to control the rotor. As in our other scripts, we use `PlotSet` for our 2D plots. Notice that we use two new options. One `'plot set'` allows you to put more than one line on a subplot. The other `'legend'` adds legends to each plot. The cell array argument to `'legend'` has a cell array for each plot. In this case, we have two plots each with two lines, so the cell array is:

```
{{'true', 'estimated'} ,{'Control' ,'Command'}}
```

Each plot legend is a cell entry within the overall cell array.

The rotor simulation script with MRAC is shown in the following listing. The square wave functions generates the command to the system that $\omega$ should track. `RHSRotor`, `SquareWave`, and `MRAC` all return default data structures. `MRAC` and `SquareWave` are called once per pass through the loop. The simulation right-hand side, that is the dynamics of the rotor, in `RHSRotor`, is then propagated using `RungeKutta`. Note that we pass a pointer to `RHSRotor` to `RungeKutta`.

```
%% Initialize
nSim     = 4000;      % Number of time steps
dT     = 0.1;       % Time step (sec)
dRHS     = RHSRotor;      % Get the default data structure
dC     = MRAC;
dS     = SquareWave;
x        = 0.1;       % Initial state vector

%% Simulation
xPlot = zeros(4,nSim);
theta = zeros(2,nSim);
t     = 0;
for k = 1:nSim
```

```
  % Plot storage
  xPlot(:,k)   = [x;dC.x(5);dC.u;dC.uC];
  theta(:,k)   = dC.x(3:4);
  [uC, dS]     = SquareWave( t, dS );
  dC.uC        = 2*(uC - 0.5);
  dC           = MRAC( x, dC );
  dRHS.u       = dC.u;

  % Propagate (numerically integrate) the state equations
  x            = RungeKutta( @RHSRotor, t, x, dT, dRHS );
  t            = t + dT;
end
```

■ **TIP** Pass pointers `@fun` instead of strings `'fun'` to functions whenever possible.

`RHSRotor` is shown below.

```
  dRHS.u       = dC.u;

  % Propagate (numerically integrate) the state equations
  x            = RungeKutta( @RHSRotor, t, x, dT, dRHS );
  t            = t + dT;
end

%% Plot the results
```

The dynamics is just one line of code. The remaining returns the default data structure.

The results are shown in Figure 5.9. We set the adaptation gain, $\gamma$ to 1. $a_m$ and $b_m$ are set equal to 2. $a$ is set equal to 1 and $b$ to $\frac{1}{2}$.

The first plot shows the estimated and true angular rates of the rotor on top and the control demand and actual control sent to the wheel on the bottom. The desired control is a square wave (generated by `SquareWave`). Notice the transient in the applied control at the transitions of the square wave. The control amplitude is greater than the commanded control. Notice also that the angular rate approaches the desired commanded square wave shape.

Figure 5.10 shows the convergence of the adaptive gains, $\theta_1$ and $\theta_2$. They have converged by the end of the simulation.

Model Reference Adaptive Control learns the gains of the system by observing the response to the control excitation. It requires excitation to converge. This is the nature of all learning systems. If there is insufficient stimulation, it isn't possible to observe the behavior of the system, so there is not enough information for learning. It is easy to find an excitation for a first-order system. For higher order systems, or nonlinear systems, this can be more difficult.

***Figure 5.9:*** MRAC control of a rotor.

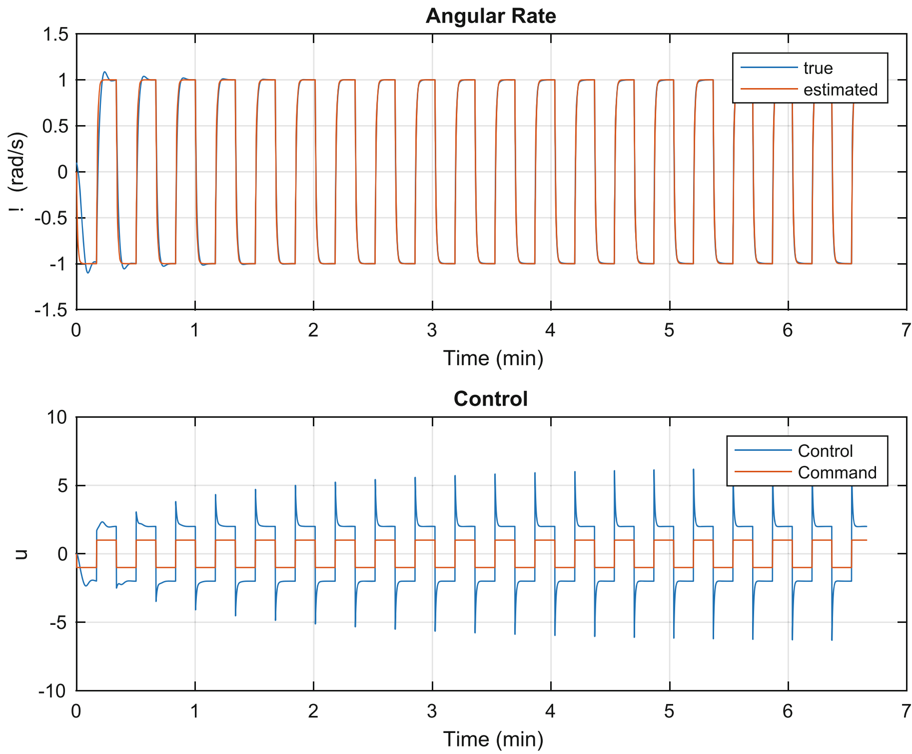


***Figure 5.10:*** Gain convergence in the MRAC controller.

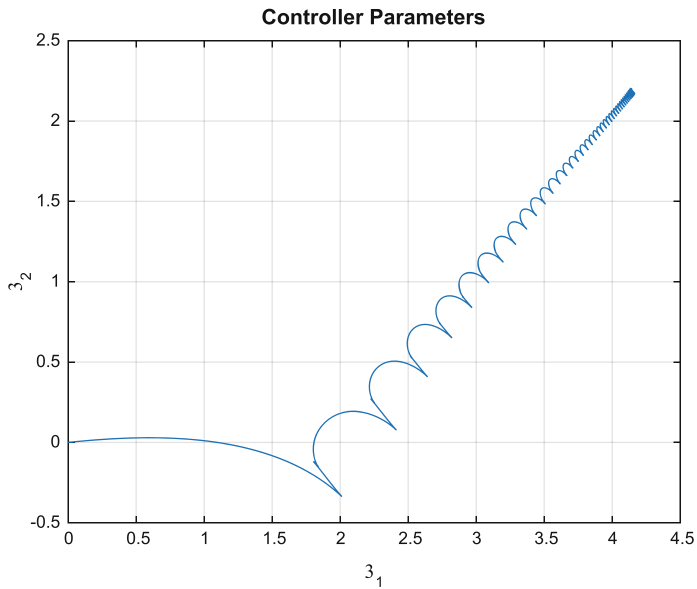

## 5.6 Ship Steering: Implement Gain Scheduling for Steering Control of a Ship

### 5.6.1 Problem

We want to steer a ship at all speeds. The problem is that the dynamics are speed dynamics, making this a nonlinear problem

### 5.6.2 Solution

The solution is to use gain scheduling to set the gains based on speeds. The gain scheduling is learned by automatically computing gains from the dynamical equations of the ship. This is similar to the self-tuning example except that we are seeking a set of gains for all speeds, not just one. In addition, we assume that we know the model of the system.

***Figure 5.11:*** Ship heading control for gain scheduling control.

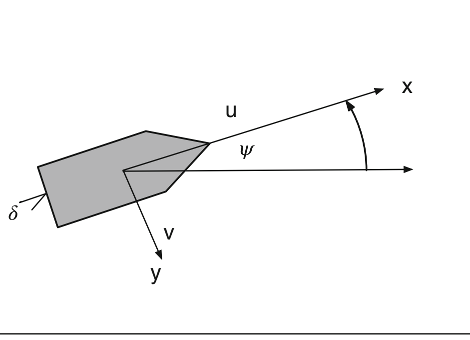


### 5.6.3 How It Works

The dynamical equations for the heading of a ship are in state space form [2]

$$\begin{bmatrix} \dot{v} \\ \dot{r} \\ \dot{\psi} \end{bmatrix} = \begin{bmatrix} \left(\frac{u}{l}\right) a_{11} & u a_{12} & 0 \\ \left(\frac{u}{l^2}\right) a_{21} & \left(\frac{u}{l}\right) a_{22} & 0 \\ 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} v \\ r \\ \psi \end{bmatrix} + \begin{bmatrix} \left(\frac{u^2}{l}\right) b_1 \\ \left(\frac{u^2}{l^2}\right) b_2 \\ 0 \end{bmatrix} \delta + \begin{bmatrix} \alpha_v \\ \alpha_r \\ 0 \end{bmatrix} \tag{5.40}$$

$v$ is the transverse speed, $u$ is the ship's speed, $l$ is the ship length, $r$ is the turning rate, and $\psi$ is the heading angle. $\alpha_v$ and $\alpha_r$ are disturbances. The ship is assumed to be moving at speed $u$. This is achieved by the propeller, which is not modeled. The control is rudder angle $\delta$. Notice that if $u = 0$, the ship cannot be steered. All of the coefficients in the state matrix are functions of $u$, except for the heading angle. Our goal is to control heading given the disturbance acceleration in the first equation and disturbance angular rate in the second.

The disturbances only affect the dynamic states, $r$ and $v$. The last state, $\psi$ is a kinematic state and does not have a disturbance.

***Table 5.1:*** Ship Parameters [3]

| Parameter | Minesweeper | Cargo | Tanker |
|---:|---:|---:|---:|
| $l$ | 55 | 161 | 350 |
| $a_{11}$ | -0.86 | -0.77 | -0.45 |
| $a_{12}$ | -0.48 | -0.34 | -0.44 |
| $a_{21}$ | -5.20 | -3.39 | -4.10 |
| $a_{22}$ | -2.40 | -1.63 | -0.81 |
| $b_1$ | 0.18 | 0.17 | 0.10 |
| $b_2$ | 1.40 | -1.63 | -0.81 |

The ship model is shown in the following code, `RHSShip`. The second and third outputs are for use in the controller. Notice that the differential equations are linear in the state and the control. Both matrices are a function of the forward velocity. We are not trying to control the forward velocity, it is an input to the system. The default parameters for the minesweeper are given in in Table 5.1.These are the same numbers that are in the default data structure.

```
function [xDot, a, b] = RHSShip( ~, x, d )

if( nargin < 1 )
  xDot = struct('l',100,'u',10,'a',[-0.86 -0.48;-5.2 -2.4],'b'
      ,[0.18;-1.4],'alpha',[0;0;0],'delta',0);
  return
end

uOL    = d.u/d.l;
uOLSq  = d.u/d.l^2;
uSqOl  = d.u^2/d.l;
a      = [  uOL*d.a(1,1) d.u*d.a(1,2) 0;...
           uOLSq*d.a(2,1) uOL*d.a(2,2) 0;...
                        0            1 0];
b      = [uSqOl*d.b(1);...
           uOL^2*d.b(2);...
           0];

xDot   = a*x + b*d.delta + d.alpha;
```

In the ship simulation, `ShipSim`, we linearly increase the forward speed while commanding a series of heading `psi` changes. The controller takes the state space model at each time step and computes new gains, which are used to steer the ship. The controller is a linear quadratic regulator. We can use full state feedback because the states are easily modeled. Such controllers will work perfectly in this case, but are a bit harder to implement when you need to estimate some of the states or have unmodeled dynamics.

```
for k = 1:nSim
  % Plot storage
  xPlot(:,k)   = x;
  dRHS.u       = u(k);
  % Control
  % Get the state space matrices
  [~,a,b]      = RHSShip( 0, x, dRHS );
  gain(k,:)    = QCR( a, b, qC, rC );
  dRHS.delta   = -gain(k,:)*[x(1);x(2);x(3) - psi(k)]; % Rudder angle
  delta(k)     = dRHS.delta;
  % Propagate (numerically integrate) the state equations
  x            = RungeKutta( @RHSShip, 0, x, dT, dRHS );
end
```

The quadratic regulator generator code is shown in the following lists. It generates the gain from the matrix Riccati equation. A Riccati equation is an ordinary differential equation that is quadratic in the unknown function. In steady state, this reduces to the algebraic Riccati equation, which is solved in this function.

```
function k = QCR( a, b, q, r )

[sinf,rr] = Riccati( [a,-(b/r)*b';-q',-a'] );

if( rr == 1 )
  disp('Repeated␣roots.␣Adjust␣q,␣r␣or␣n');
end

k = r\(b'*sinf);

function [sinf, rr] = Riccati( g )
%% Ricatti
%   Solves the matrix Riccati equation in the form
%
%   g = [a    r ]
%       [q   -a']

rg = size(g);

[w, e] = eig(g);

es = sort(diag(e));

% Look for repeated roots
j = 1:length(es)-1;

if ( any(abs(es(j)-es(j+1))<eps*abs(es(j)+es(j+1))) )
  rr = 1;
else
```

```
  rr = 0;
end

% Sort the columns of w
ws    = w(:,real(diag(e)) < 0);

sinf = real(ws(rg/2+1:rg,:)/ws(1:rg/2,:));
```

a is the state transition matrix, b is the input matrix, q is the state cost matrix, and r is the control cost matrix. The bigger the elements of q, the more cost we place on deviations of the states from zero. That leads to tight control at the expense of more control. The bigger the elements of b, the more cost we place on control. Bigger b means less control. Quadratic regulators guarantee stability if all states are measured. They are a very handy controller to get something working. The results are given in Figure 5.12. Note how the gains evolve. The gain on the angular rate $r$ is nearly constant. Notice that the $\psi$ range is very small! Normally, you would zoom out the plot. The other two gains increase with speed. This is an example of gain scheduling. The difference is that we autonomously compute the gains from perfect measurements of the ship's forward speed.

***Figure 5.12:*** Ship steering simulation. The states are shown on the left with the forward velocity. The gains and rudder angle are shown on the right. Notice the "pulses" in the rudder to make the maneuvers.

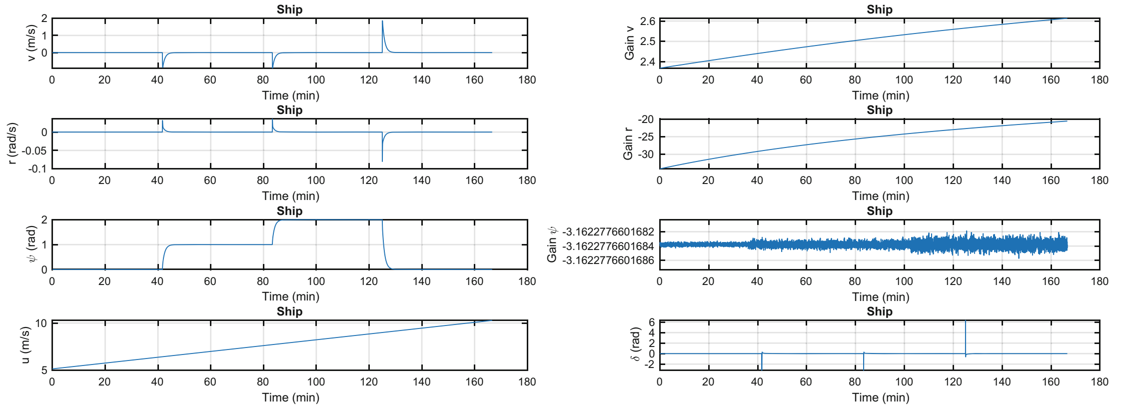

`ShipSimDisturbance` is a modified version of `ShipSim`, which is of shorter duration, with only one course change, and with disturbances in both angular rate and lateral velocity. The results are given in Figure 5.13.

***Figure 5.13:*** Ship steering simulation. The states are shown on the left with the rudder angle. The disturbances are Gaussian white noise.

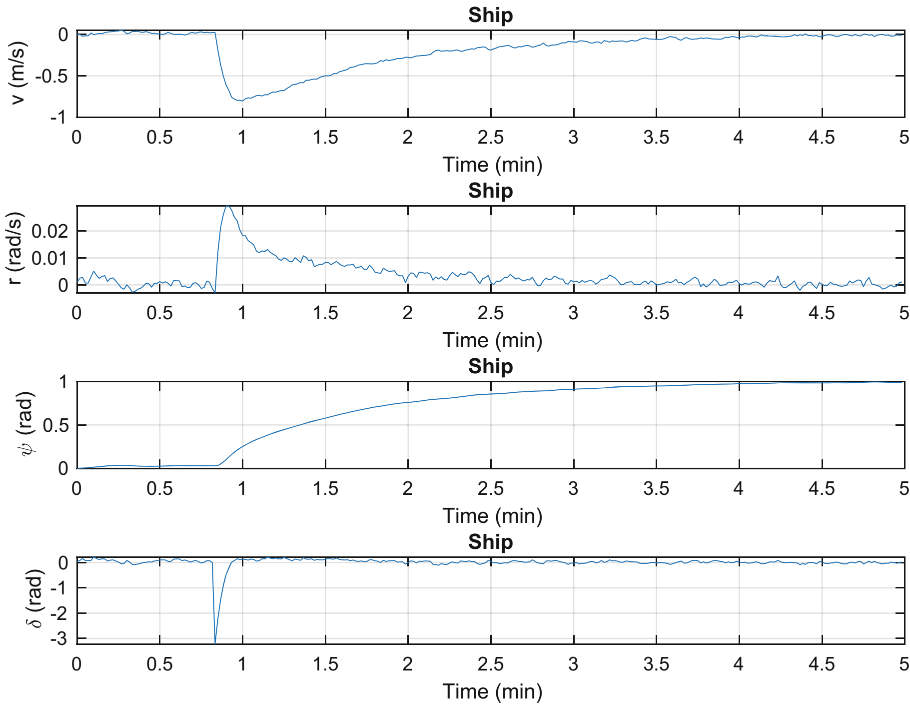


## 5.7 Spacecraft Pointing

### 5.7.1 Problem

We want to control the orientation of a spacecraft with thrusters for control.

### 5.7.2 Solution

The solution is to use a parameter estimator to estimate the inertia and feed it into the control system.

### 5.7.3 How It Works

The spacecraft model is shown in Figure 5.14.

***Figure 5.14:*** Spacecraft model.

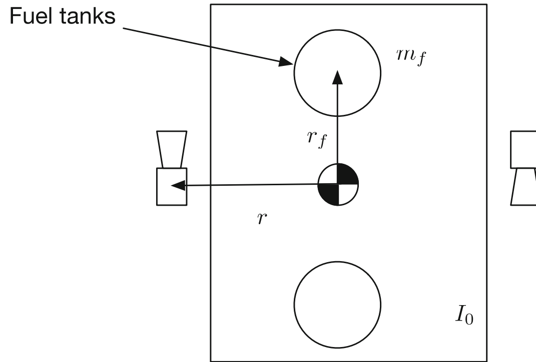


The dynamical equations are

$$I = I_0 + m_f r_f^2 \tag{5.41}$$

$$T_c + T_d = I\ddot{\theta} + \dot{m}_f r_f^2 \dot{\theta} \tag{5.42}$$

$$\dot{m}_f = -\frac{T_c}{r u_e} \tag{5.43}$$

where $I$ is the total inertia, $I_0$ is the constant inertia for everything except the fuel mass, $T_c$ is the thruster control torque, $T_d$ is the disturbance torque, $m_f$ is the total fuel mass, $r_f$ is the distance to the fuel tank center, $r$ is the vector to the thrusters, $u_e$ is the thruster exhaust velocity, and $\theta$ is the angle of the spacecraft axis. Fuel consumption is balanced between the two tanks so that the center-of-mass remains at (0,0). The second term in the second equation is the inertia derivative term, which adds damping to the system.

Our controller is a proportional derivative controller of the form:

$$T_c = Ia \tag{5.44}$$

$$a = -K(\theta + \tau\dot{\theta}) \tag{5.45}$$

$K$ is the forward gain and $\tau$ the rate constant. We design the controller for a unit inertia and then estimate the inertia so that our dynamical response is always the same. We will estimate the inertia using a very simple algorithm:

$$I_k = \frac{T_{c_{k-1}}}{\ddot{\theta}_k - \ddot{\theta}_{k-1}} \tag{5.46}$$

We will do this only when the control torque is not zero and the change in rate is not zero. This is a first difference approximation and should be good if we don't have a lot of noise. The following shows a code snippet showing the simulation loop with the control system.

```
%% Initialize
nSim       = 50;              % Number of time steps
dT         = 1;               % Time step (sec)
dRHS       = RHSSpacecraft;   % Get the default data structure
x          = [2.4;0.;1];      % [angle;rate;mass fuel]

%% Controller
kForward   = 0.1;
tau        = 10;

%% Simulation
xPlot      = zeros(6,nSim);
omegaOld   = x(2);
inrEst     = dRHS.i0 + dRHS.rF^2*x(3);
dRHS.tC    = 0;
tCThresh   = 0.01;
kI         = 0.99; % Inertia filter gain

for k = 1:nSim
  % Collect plotting information
  xPlot(:,k)   = [x;inrEst;dRHS.tD;dRHS.tC];
  % Control
  % Get the state space matrices
  dRHS.tC   = -inrEst*kForward*(x(1) + tau*x(2));
  omega     = x(2);
  omegaDot = (omega-omegaOld)/dT;
  if( abs(dRHS.tC) > tCThresh  )
    inrEst = kI*inrEst + (1-kI)*omegaDot/(dRHS.tC);
  end
  omegaOld = omega;
  % Propagate (numerically integrate) the state equations
  x            = RungeKutta( @RHSSpacecraft, 0, x, dT, dRHS );
end
```

We only estimate inertia when the control torque is above a threshold. This prevents us from responding to noise. We also incorporate the inertia estimator in a simple low pass filter. The results are shown in Figure 5.15. The threshold has it only estimating inertia at the very beginning of the simulation when it is reducing the attitude error.

This algorithm appears crude, but it is fundamentally all we can do in this situation given just angular rate measurements. More sophisticated filters or estimators could improve the performance.

***Figure 5.15:*** States and control outputs from the spacecraft simulation.

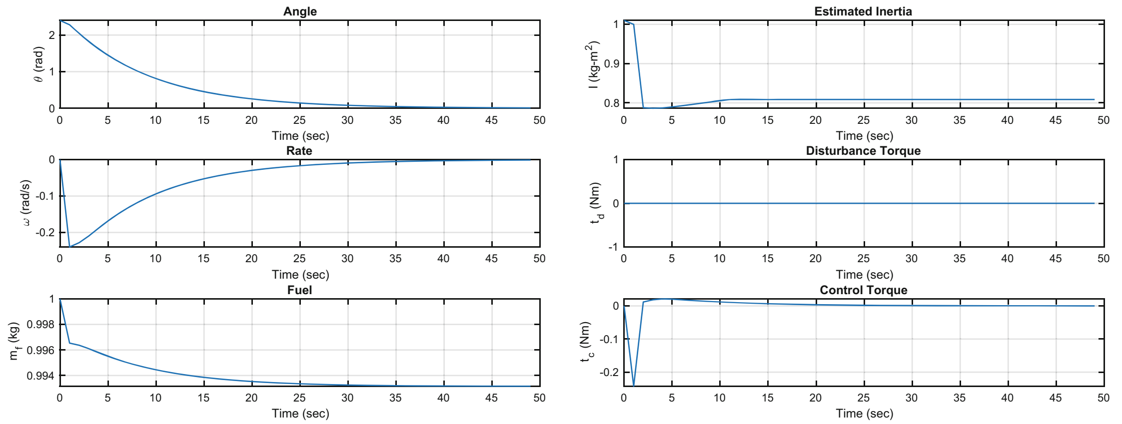


## 5.8 Summary

This chapter has demonstrated adaptive or learning control. You learned about model tuning, model reference adaptive control, adaptive control and gain scheduling. Table 5.2 lists the functions and scripts included in the companion code.

***Table 5.2:*** Chapter Code Listing

| File | Description |
|---|---|
| Combinations | Enumerates n integers for 1:n taken k at a time. |
| FFTEnergy | Generates fast Fourier transform energy. |
| FFTSim | Demonstration of the fast Fourier transform. |
| MRAC | Implement model reference adaptive control. |
| QCR | Generates a full state feedback controller. |
| RHSOscillatorControl | Right-hand side of a damped oscillator with a velocity gain. |
| RHSRotor | Right-hand side for a rotor. |
| RHSShip | Right-hand side for a ship steering model. |
| RHSSpacecraft | Right-hand side for a spacecraft model. |
| RotorSim | Simulation of model reference adaptive control. |
| ShipSim | Simulation of ship steering. |
| ShipSimDisturbance | Simulation of ship steering with disturbances. |
| SpacecraftSim | Time varying inertia demonstration. |
| SquareWave | Generate a square wave. |
| TuningSim | Controller tuning demonstration. |
| WrapPhase | Keep angles between $-\pi$ and $\pi$. |

# CHAPTER 6

# Fuzzy Logic

Fuzzy logic [26] is an alternative approach to control system design. Fuzzy logic works within the framework of set theory and is better at dealing with ambiguities. For example, three sets may be defined for a sensor: hard failure, soft failure, and no failure. The three sets may overlap and at any given time the sensor may have a degree of membership in each set. The degree of membership in each set can be used to determine what action to take. An algorithmic approach would have to assign a number to the state of the sensor. This could be problematic and not necessarily represent the actual state of the system. In effect, you would be applying a degree of fuzziness.

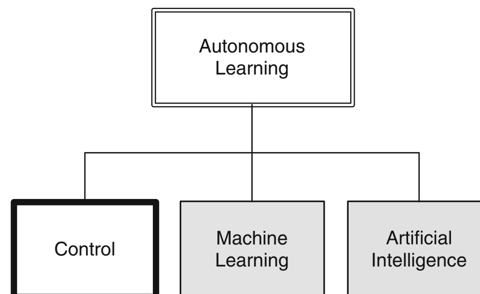


When you go to a doctor with pain the doctor will often try and get you to convert a fuzzy concept, pain, into a number from 0 to 10. As pain is personal and your impression is imprecise, you are giving a fuzzy concept or belief a hard number. As you may have experienced, this is not terribly productive or useful.

Surveys do the same thing. For example, you will be asked to rate the service in a restaurant from 0 to 5. You then rate a bunch of other things on the same scale. This allows the review to come up with a number for your overall impression of the restaurant. Does the resulting 4.8 actually mean anything? Netflix abandoned the numerical ratings of movies you have seen for thumbs up and down. It seems that they felt that a binary decision, really two sets, was a better indicator than a number.

NASA and the U.S. Department of Defense like to use technology readiness levels that go from 1 to 9 to determine where your work is in terms of readiness. Nine is a technology already

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_6

operating in a target system. One is just an idea. All the other levels are fuzzy for anything moderately complicated. Even giving a technology a 9 is not terribly informative. The M-16 rifle was deployed to Vietnam. It often jammed. In terms of TRL it was 9, but a 9 doesn't say how well it is working. Again, the readiness of the rifle, when you read soldiers' and marines' impressions, was best represented by fuzzy beliefs.

This chapter will show you how to implement a fuzzy logic control system for windshield wipers. Unlike the other chapters, we will be working with linguistic concepts, not hard numbers. Of course, when you set your wiper motor speed you need to pick a number (defuzzify your output), but all the intermediate steps employ fuzzy logic.

## 6.1 Building Fuzzy Logic Systems

### 6.1.1 Problem

We want to have a tool to build a fuzzy logic controller.

### 6.1.2 Solution

Build a MATLAB function that takes parameter pairs that define everything needed for the fuzzy controller.

### 6.1.3 How It Works

To create a fuzzy system you must create inputs, outputs, and rules. You can also choose methods for some parts of the fuzzy inference. The fuzzy inference engine has three steps:

1. Fuzzify
2. Fire
3. Defuzzify

The fuzzy system data are stored in a MATLAB data structure. This structure has the following fields:

- Input {:}
- Output {:}
- Rules {:}
- Implication
- Aggregation
- Defuzzify

The first three fields are cell arrays of structs. There is a separate structure for rules and fuzzy sets, described below. The last three fields are strings containing the names of the desired functions.

The fuzzy set structure has the following fields:

- name
- range(2) (two-element array with minimum and maximum values)
- comp {:} (cell array of label strings)
- type {:} (cell array of membership function names)
- params {:} (cell array of parameter vectors)

The fuzzy rule struct has the following fields:

- input(:) (vector of input component numbers)
- output(:) (vector of outputs)
- operator {:} (cell array of operator strings)

This is a lot of data to organize. We do it with the function `BuildFuzzySystem`. The following code snippet shows how it assigns data to the data structure using parameter pairs.

```
d = load('SmartWipers');

j = 1;

for k = 1:2:length(varargin)
  switch (lower(varargin{k}))
    case 'id'
      j = varargin{k+1};
    case 'input_comp'
      d.input(j).comp = varargin{k+1};
    case 'input_type'
      d.input(j).type = varargin{k+1};
    case 'input_name'
       d.input(j).name = varargin{k+1};
    case 'input_params'
```

This code continues with other cases. If you don't enter anything, `BuildFuzzySystem` loads the Smart Wipers demo, as shown above. and returns it unchanged. For example, if you just enter one input type you get:

```
SmartWipers = BuildFuzzySystem(...
              'id',1,...
              'input_comp',{'Dry'  'Drizzle'  'Wet'} ,...
              'input_type', {'Trapezoid'  'Triangle'  'Trapezoid'}
                 ,...
              'input_params',{[0 0 10 50]  [40 50]  [50 90 101
                 101]},...
```

```
              'input␣range',[0 100])

SmartWipers =
  struct with fields:

    SmartWipers: [1x1 struct]
          input: [1x1 struct]
```

Fuzzy sets in this context consist of a set of linguistic categories or components defining a variable. For instance, if the variable is “age,” the components might be “young,” “middle-aged,” and “old.” Each fuzzy set has a range over which it is valid, for instance, a good range for “age” may be 0 to 100. Each component has a membership function that describes the degree to which a value in the set’s range belongs to each component. For instance, a person who is 50 would rarely be described as “young,” but might be described as “middle aged” or “old,” depending on the person asked.

To build a fuzzy set, you must divide it into components. The following membership functions are provided:

1. Triangular

2. Trapezoidal

3. Gaussian

4. General bell

5. Sigmoidal

Membership functions are limited in value to between 0 and 1. The membership functions are shown in Figure 6.1.

The triangular membership function requires two parameters: the center of the triangle and the half-width of the desired triangle base. Triangular membership functions are limited to symmetrical triangles.

The trapezoid membership function requires four parameters: the left-most point, the start of the plateau, the end of the plateau, and the right-most points.

A Gaussian membership function is a continuous function with two parameters: the center of the bell and the width (standard deviation) of the bell. Gaussian membership functions are symmetrical.

A general bell function is also continuous and symmetrical, but it has three parameters to allow for a flattened top, making it similar to a smoothed trapezoid. It requires three parameters: the center of the bell, the width of the bell at the points $y = 0.5$, and the slope of the function at the points $y = 0.5$.

Just as a bell function is similar to a smoothed trapezoid, a sigmoidal membership function is similar to a smoothed step function. It takes two parameters: the point at which $y = 0.5$ and the slope of the function. As the slope approaches infinity the sigmoidal function approaches the step function.

***Figure 6.1:*** Membership functions.

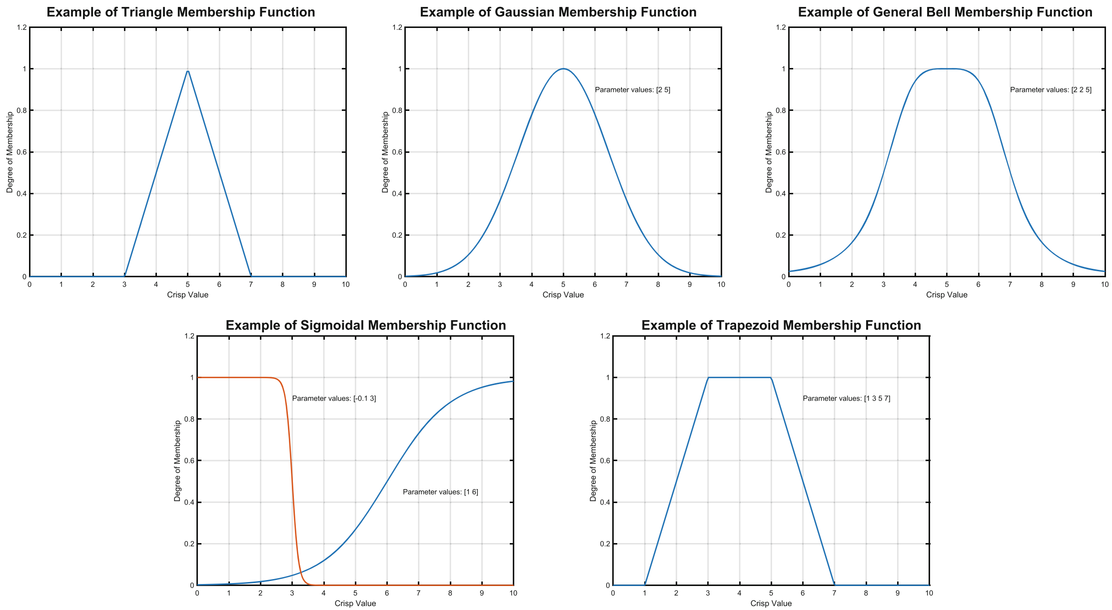


Fuzzy rules are if-then statements. For example, an air conditioner rule might say IF the room temperature IS high, THEN the blower level IS high. In this case, "room temperature" is the input fuzzy set, "high" is its component for this rule, "blower level" is the output fuzzy set, and "high" is its chosen component.

## 6.2 Implement Fuzzy Logic

### 6.2.1 Problem

We want to implement fuzzy logic.

### 6.2.2 Solution

Build a fuzzy inference engine.

### 6.2.3 How It Works

Let's repeat the three steps in fuzzy inference, adding two steps within defuzzify:

1. Fuzzify
2. Fire
3. Defuzzify
    (a) Implication
    (b) Aggregation

The control flow is in the main function, called `FuzzyInference`. It just calls `Fuzzify`, `Fire`, and `Defuzzify` in order. It calls `warndlg` if the inputs are not sensible.

```
function y = FuzzyInference( x, system )

if length(x) == length( system.input )
  fuzzyX   = Fuzzify( x, system.input );
  strength = Fire( fuzzyX, system.rules );
  y        = Defuzzify( strength, system );
else
  warndlg({'The length of x must be equal to the',...
           'number of input sets in the system.'})
end
```

You will notice the use of `eval` to evaluate function names stored as strings as the input. You could also store pointers and do the same thing. For example, for the function:

```
function y = MyFun(x)
y = x;
```

`eval` works on a string. Essentially, it applies the MATLAB parser to your own text. You can make the string as complex as you want, albeit at the expense of readability. You can also do things such as make self-modifying code.

```
>> eval(['MyFun(',sprintf('%f',2),')'])

ans =
     2
```

It is cleaner and takes less processing time to use a pointer to the function.

```
>> feval(@MyFun,2)

ans =
     2
```

`feval` works on a pointer to the function and is generally faster.

---

■ **TIP** Use `feval` instead of `eval` whenever possible.

---

The fuzzify sub-function code is shown below. It puts the data into the various input membership sets. An input may be in more than one set.

```
function fuzzyX = Fuzzify( x, sets )

m = length(sets);
fuzzyX = cell(1,m);
for i = 1:m
  n = length(sets(i).comp);
  range = sets(i).range(:);
  if range(1) <= x(i) <= range(2)
```

```
    for j = 1:n
      fuzzyX{i}(j) = eval([sets(i).type{j} 'MF(x(i),[' num2str(sets(i
         ).params{j}) '])']);
    end
  else
    fuzzyX{i}(1:n) = zeros(1,n);
  end
end
```

The fuzzy rules fire in the following code. The code applies “Fuzzy AND” or “Fuzzy OR.” “Fuzzy AND” is the minimum of a set of membership values. “Fuzzy OR” is the maximum of a set of membership values. Suppose we have a vector `[1 0 1 0]`. The maximum value is 1 and the minimum is 0.

```
>> 1 && 0 &&  1 && 0

ans =

  logical
   0

>> 1 || 0 ||  1 || 0

ans =

  logical
   1
```

This corresponds to the fuzzy logic AND and OR.

The next code snippet shows the `Fire` sub-function in `FuzzyInference`.

```
function strength = Fire( FuzzyX, rules )

m = length( rules );
n = length( FuzzyX );

strength = zeros(1,m);

for i = 1:m
  method = rules(i).operator;
  for j = 1:n
    comp = rules(i).input(j);
    if comp ~= 0
      dom(j) = FuzzyX{j}(comp);
    else
      dom(j) = inf;
    end
  end
  strength(i) = eval([method '(dom(find(dom<=1)))']);
end
```

Finally, we defuzzify the results. This function first uses the implication functions to determine membership. It aggregates the output using the aggregate function which, in this case, is max.

```
function result = Defuzzify( strength, system )

rules   = system.rules;
output  = system.output;

m       = length( output );
n       = length( rules );
imp     = system.implicate;
agg     = system.aggregate;
defuzz  = system.defuzzify;
result  = zeros(1,m);

for i = 1:m
  range = output(i).range(:);
  x = linspace( range(1),range(2),200 );
  for j = 1:n
    comp = rules(j).output(i);
    if( comp ~= 0 )
      mf        = [output(i).type{comp} 'MF'];
      params    = output(i).params{comp};
      mem(j,:)  = eval([ imp 'IMP(' mf '(x,params),strength(j))']);
    else
      mem(j,:)  = zeros(size(x));
    end
  end
  aggregate = eval([ agg '(mem)' ]);
  result(i) = eval([ defuzz 'DF(aggregate,␣x)']);
end
```

## 6.3 Demonstrate Fuzzy Logic

### 6.3.1 Problem

We want a control system to select window wiper speed and interval based on rainfall.

### 6.3.2 Solution

Build a fuzzy logic control system using the tools we've developed.

### 6.3.3 How It Works

To call a fuzzy system, use the function `y = FuzzyInference( x, system )`.

The script `SmartWipersDemo` implements the rainfall demo. We only show the code that calls the inference engine. The fuzzy system is loaded using `SmartWipers = BuildFuzzySystem()`, as discussed above.

```
% Generate regularly space arrays in the 2 inputs
x = linspace(SmartWipers.input(1).range(1),SmartWipers.input(1).range
    (2),n);
y = linspace(SmartWipers.input(2).range(1),SmartWipers.input(2).range
    (2),n);

PlotSet(1:n,[x;y],'x label','Input','y label',{'Wetness','Intensity'
    },...
         'figure title','Inputs','Plot Title',{'Wetness','Intensity'})

h = waitbar(0,'Smart Wipers Demo: plotting the rule base');
z1 = zeros(n,n);
z2 = zeros(n,n);
for k = 1:n
  for j = 1:n
    temp = FuzzyInference([x(k),y(j)], SmartWipers);
    z1(k,j) = temp(1);
    z2(k,j) = temp(2);
  end
  waitbar(k/n)
end
close(h);

NewFigure('Wiper Speed from Fuzzy Logic')
```

Smart wipers is a control system for an automatic windshield wiper [7]. First, the demo will plot the input and output fuzzy variables. Fuzzy inference is performed on each set of crisp inputs plotted. Figure 6.2 shows the inputs to the fuzzy logic system. Figure 6.3 shows the outputs.

The inputs that are tested in the fuzzy logic system are given in Figure 6.4.

Figure 6.5 gives surface plots to show how the outputs relate to the inputs. The surface plots are generated by the code below. We add a `colorbar` to make the plot more readable. The color is related to z-value. We use `view` in the second plot to make it easier to read in the figure. You can use `rotate3d on` to allow you to rotate the figure with the mouse.

```
NewFigure('Wiper Speed from Fuzzy Logic')
surf(x,y,z1)
xlabel('Raindrop Wetness')
ylabel('Droplet Frequency')
zlabel('Wiper Speed')
colorbar
```

***Figure 6.2:*** Rain wetness and intensity are the inputs for the smart wiper control system.

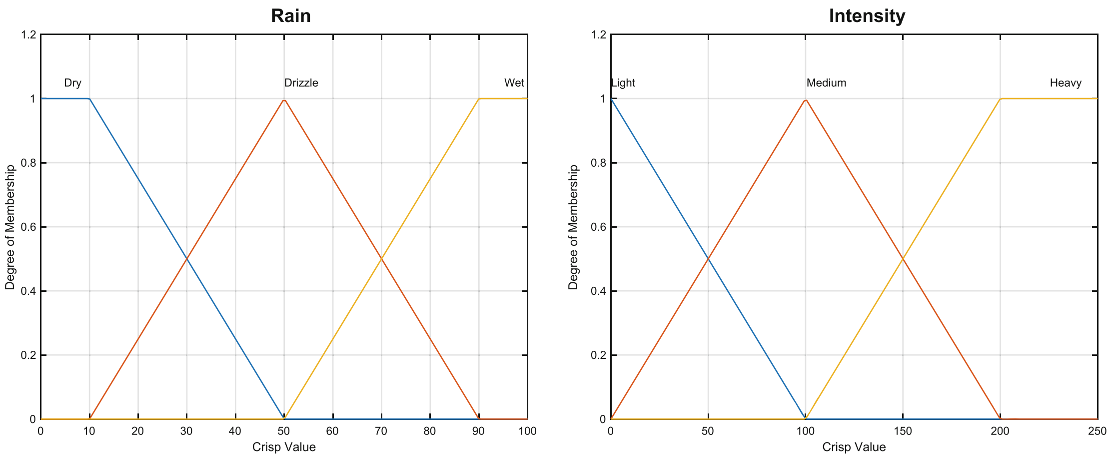


***Figure 6.3:*** Wiper speed and interval are the outputs for the smart wiper control system.

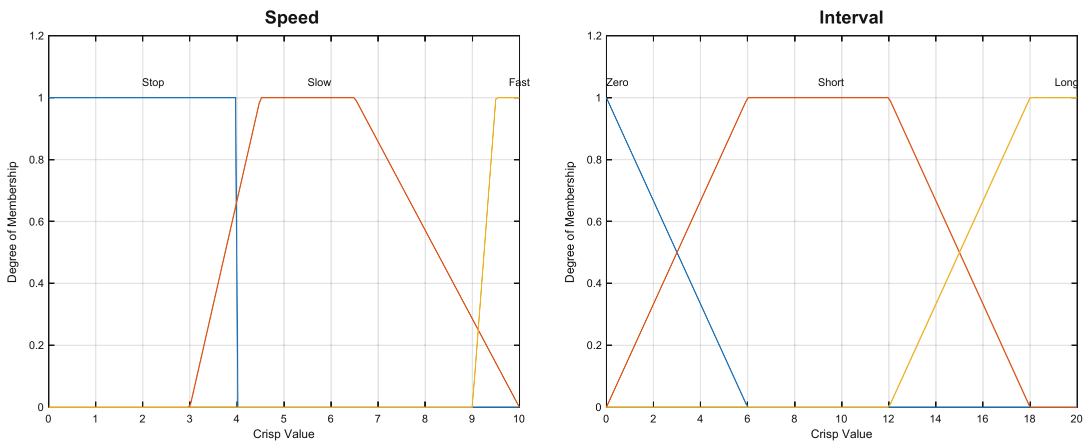


```
NewFigure('Wiper␣Interval␣from␣Fuzzy␣Logic')
surf(x,y,z2)
xlabel('Raindrop␣Wetness')
ylabel('Droplet␣Frequency')
zlabel('Wiper␣Interval')
view([142.5 30])
```

■ **TIP** Use `rotate3d on` to rotate a figure with the mouse.

***Figure 6.4:*** Rain wetness and intensity input numbers.

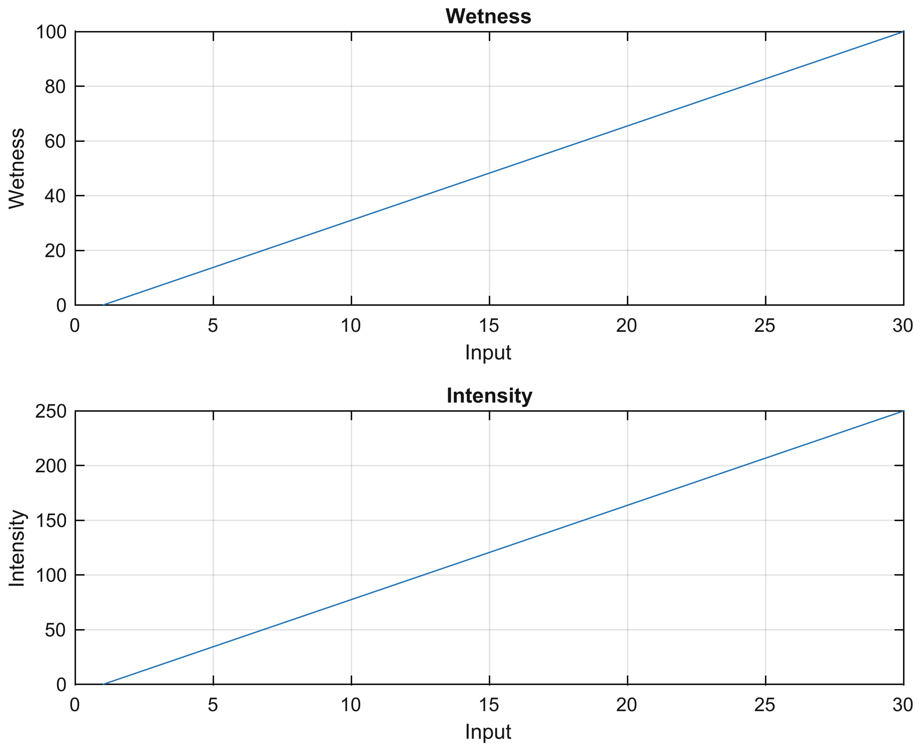


***Figure 6.5:*** Wiper speed and interval versus droplet frequency and wetness.

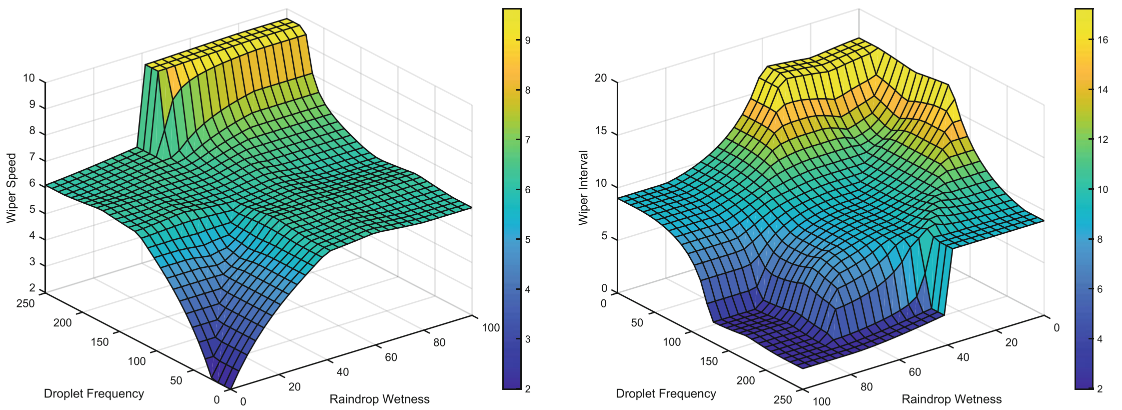

## 6.4 Summary

This chapter demonstrated fuzzy logic. A windshield wipers demonstration gives an example of how it is used. Table 6.1 lists the functions and scripts included in the companion code.

***Table 6.1:*** Chapter Code Listing

| File | Description |
|---|---|
| BuildFuzzySystem | Builds a fuzzy logic system (data structure) using parameter pairs. |
| SmartWipersDemo | Demonstrates a fuzzy logic control system for windshield wipers. |
| FuzzyPlot | Plots a fuzzy set. |
| TriangleMF | Triangle membership function. |
| TrapezoidMF | Trapezoid membership function. |
| SigmoidalMF | Display a neural net with multiple layers. |
| ScaleIMP | Scale implication function. |
| ClipIMP | Clip implication function. |
| GeneralBellMF | General bell membership function. |
| GaussianMF | Gaussian membership function. |
| FuzzyOR | Fuzzy Or (maximum of membership values). |
| FuzzAND | Fuzzy And (minimum of membership values). |
| FuzzyInference | Performs fuzzy inference given a fuzzy system and crisp data x. |
| CentroidDF | Centroid defuzzification. |

CHAPTER 7

# Data Classification with Decision Trees

In this chapter, we will develop the theory for binary decision trees. Decision trees can be used to classify data, and fall into the Learning category in our Autonomous Learning taxonomy. Binary trees are easiest to implement because each node branches to two other nodes, or none. We will create functions for the Decision Trees and to generate sets of data to classify. Figure 7.1 shows a simple binary tree. Point "a" is in the upper left quadrant. The first binary test finds that its $x$ value is greater than 1. The next test finds that its $y$ value is greater than 1 and puts it in set 2. Although the boundaries show square regions, the binary tree really tests for regions that go to infinity in both $x$ and $y$.

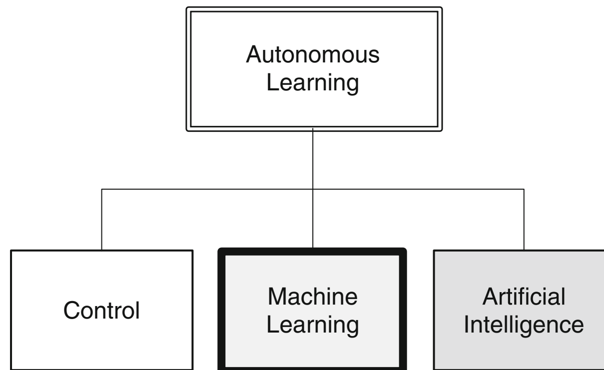


A binary decision tree is a decision tree in which at each decision node there are only two decisions to make. Once you make a decision, the next decision node provides you with two additional options. Each node accepts a binary value of 0 or 1. 0 sends you down one path, and 1 the other. At each decision node, you are testing a new variable. When you get to the bottom, you will have found a path where all of the values are true. The problem with a binary tree of $n$ variables is that it will have $2^n - 1$ nodes. Four variables would require 15 decision nodes. Eight would require 65 decision nodes, and so forth. If the order of testing variables is fixed, we call it an ordered tree.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_7

***Figure 7.1:*** A simple binary tree with one point to classify.

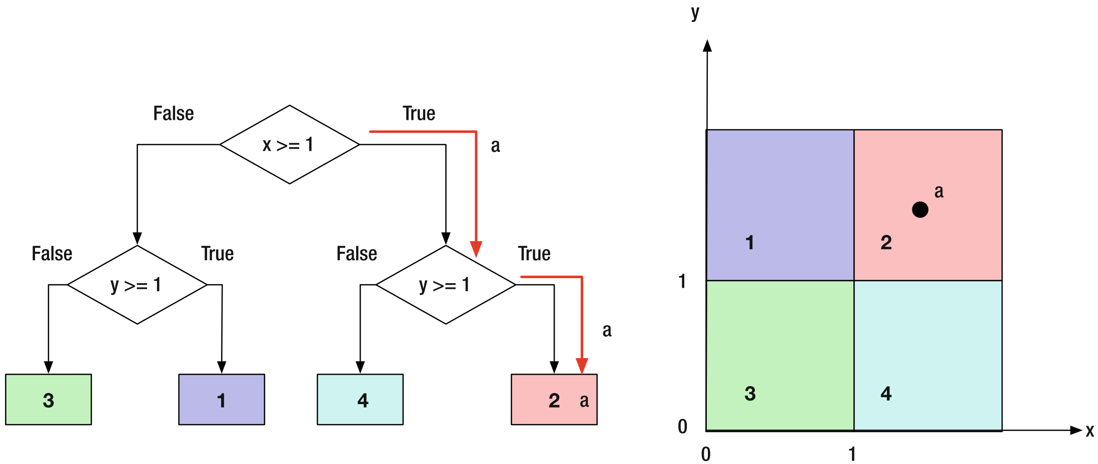


For classification, we are assuming that we can make a series of binary decisions to classify something. If we can, we can implement the reasoning in a binary tree.

# 7.1 Generate Test Data

## 7.1.1 Problem

We want to generate a set of training and testing data for classification.

## 7.1.2 Solution

Write a function using `rand` to generate data over a selected range in two dimensions, $x$ and $y$.

## 7.1.3 How It Works

The function `ClassifierSet` generates random data and assigns them to classes. The function call is:

```
function p = ClassifierSets( n, xRange, yRange, name, v, f, setName )
```

The first argument to `ClassifierSets` is the square root of the number of points. The second, `xRange`, gives the `x` range for the data and the third, `yRange`, gives the `y` range. The $n^2$ points will be placed randomly in this region. The next argument is a cell array with the names of the sets, `name`. These are used for plot labels. The remaining inputs are a list of vertices, `v`, and the faces, `f`. The faces select the vertices to use in each polygon. The faces connect the vertices into specific polygons. `f` is a cell array, since each face array can be of any length. A triangle has a length of 3, a hexagon a length of 6. Triangles, rectangles, and hexagons can be easily meshed so that there are no gaps.

Classes are defined by adding polygons that divide the data into regions. Any polygon can be used. You should pick polygons so that there are no gaps. Rectangles are easy, but you could

also use uniformly sized hexagons. The following code is the built-in demo. The demo is the last subfunction in the function. This specifies the vertices and faces.

```
function Demo

v = [0 0;0 4; 4 4; 4 0; 0 2; 2 2; 2 0;2 1;4 1;2 1];
f = {[5 6 7 1] [5 2 3 9 10 6] [7 8 9 4]};
ClassifierSets( 5, [0 4], [0 4], {'width', 'length'}, v, f );
```

In this demo, there are three polygons. All points are defined in a square ranging from 0 to 4 in both the $x$ and $y$ directions.

The other subfunctions are `PointInPolygon` and `Membership`. `Membership` determines if a point is in a polygon. `Membership` calls `PointInPolygon` to assign points to sets. `ClassifierSets` randomly puts points in the regions. It figures out which region each point is in using this code in the function, `PointInPolygon`.

```
function r = PointInPolygon( p, v )

m = size(v,2);

% All outside
r = 0;

% Put the first point at the end to simplify the looping
v = [v v(:,1)];

for i = 1:m
  j   = i + 1;
  v2J   = v(2,j);
  v2I = v(2,i);
  if (((v2I > p(2)) ~= (v2J > p(2))) && ...
      (p(1) < (v(1,j) - v(1,i)) * (p(2) - v2I) / (v2J - v2I) + v(1,i)
         ))
    r = ~r;
  end
end
```

This code can determine if a point is inside a polygon defined by a set of vertices. It is used frequently in computer graphics and in games when you need to know if one object's vertex is in another polygon. You could correctly argue that this function could replace our decision tree logic for this type of problem. However, a decision tree can compute membership for more complex sets of data. Our classifier set is simple and makes it easy to validate the results.

Run `ClassifierSets` to see the demo. Given the input ranges, it determines the membership of randomly selected points. `p` is a data structure that holds the vertices and the membership. It plots the points after creating a new figure using `NewFigure`. It then uses `patch` to create the rectangular regions.

```
p.x     = (xRange(2) - xRange(1))*(rand(n,n)-0.5) + mean(xRange);
p.y     = (yRange(2) - yRange(1))*(rand(n,n)-0.5) + mean(yRange);
p.m     = Membership( p, v, f );

NewFigure(setName);
i = 0;
drawNum = n^2 < 50;
for j = 1:n
  for k = 1:n
    i = i + 1;
    plot(p.x(k,j),p.y(k,j),'marker','o','MarkerEdgeColor','k')
    if( drawNum )
      text(p.x(k,j),p.y(k,j),sprintf(' %3d',i));
    end
    hold on
  end
end

m = length(f);
a = linspace(0,2*pi-2*pi/m,m)';
c = abs(cos([a a+pi/6 a+3*pi/5]));

for k = 1:m
  patch('vertices',v,'faces',f{k},'facecolor',c(k,:),'facealpha',0.1)
end

xlabel(name{1});
ylabel(name{2});
grid on
```

The function shows the data numbers if there are fewer than 50 points. The MATLAB-function `patch` is used to generate the polygons. The code shows a range of graphics coding including the use of graphics parameters. Notice the way we create `m` colors.

---

■ **TIP** You can create an unlimited number of colors for plots using `linspace` and `cos`.

---

`ClassifierTestSet` can generate test sets or demonstrate a trained decision tree. The drawing shows that the classification regions are regions with sides parallel to the $x$- or $y$-axes. The regions should not overlap.

***Figure 7.2:*** Classifier set with three regions from the demo. Two are rectangles and one is L-shaped.

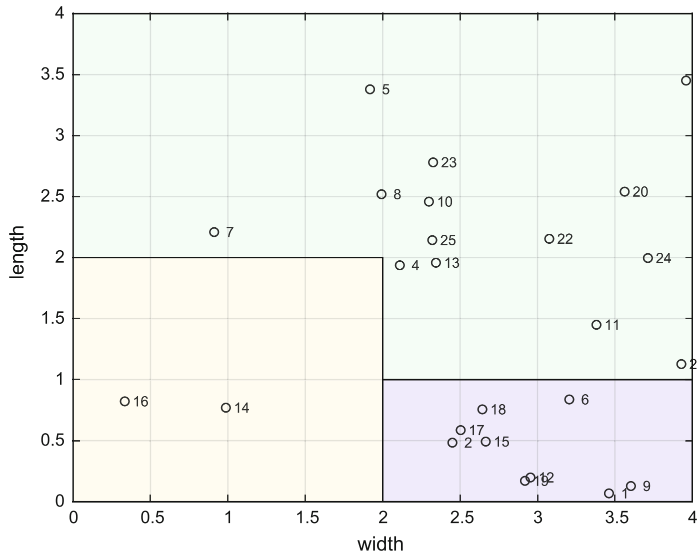


## 7.2 Drawing Decision Trees

### 7.2.1 Problem

We want to draw a binary decision tree to show decision tree thinking.

### 7.2.2 Solution

The solution is to use MATLAB graphics functions, `patch`, `text`, and `line` to draw a tree.

### 7.2.3 How It Works

The function `DrawBinaryTree` draws any binary tree. The function call is

```
function d = DrawBinaryTree( d, name )
```

You pass it a data structure, `d`, with the decision criteria in a cell array. The `name` input is optional. It has a default option for the name. The boxes start from the left and go row by row. In a binary tree, the number of rows is related to the number of nodes through the formula for a geometric series:

$$m = \log_2(n) \tag{7.1}$$

where $m$ is the number of rows and $n$ is the number of boxes. Therefore, the function can compute the number of rows.

The function starts by checking the number of inputs and either runs the demo or returns the default data structure. When you write a function you should always have defaults for anything where one is possible.

■ **TIP** Whenever possible, have default inputs for function arguments.

It immediately creates a new figure with that name. It then steps through the boxes assigning them to rows based on it being a binary tree. The first row has one box, the next two boxes, the following four boxes, etc. As this is a geometric series, it will soon get unmanageable! This points to a problem with decision trees. If they have a depth of more than four, even drawing them is impossible. As it draws the boxes, it computes the bottom and top points, which will be the anchors for the lines between the boxes. After drawing all the boxes, it draws all the lines.

All of the drawing functionality is in the subfunction `DrawBox`.

```
v = [x y 0;x y+h 0; x+w y+h 0;x+w y 0];

patch('vertices',v,'faces',[1 2 3 4],'facecolor',[1;1;1]);

text(x+w/2,y + h/2,t,'fontname',d.font,'fontsize',...
  d.fontSize,'HorizontalAlignment','center');

%% DrawBinaryTree>DefaultDataStructure
```

This draws a box using the `patch` function and the text using the `text` function. `'facecolor'` is white. Red green blue (RGB) numbers go from 0 to 1. Setting `'facecolor'` to `[1 1 1]` makes the face white and leaves the edges black. As with all MATLAB graphics, there are dozens of properties that you can edit to produce beautiful graphics. Notice the extra arguments in `text`. The most interesting is `'HorizontalAlignment'` in the last line. It allows you to center the text in the box. MATLAB does all the figuring of font sizes for you.

The following listing shows the code in `DrawBinaryTree`, for drawing the tree, starting after checking for demos. The function returns the default data structure if one output and no inputs are specified. The first part of the code creates a new figure and draws the boxes at each node. It also creates arrays for the box locations for use in drawing the lines that connect the boxes. It starts off with the default argument for name. The first set of loops draws the boxes for the trees. `rowID` is a cell array. Each row in the cell is an array. A cell array allows each cell to be different. This makes it easy to have different length arrays in the cell. If you used a standard matrix, you would need to resize rows as new rows were added.

```
if( nargin < 2 )
  name = 'Binary Tree';
end

NewFigure(name);
m       = length(d.box);
nRows   = ceil(log2(m+1));
w       = d.w;
h       = d.h;
i       = 1;
```

```
x       = -w/2;
y       =  1.5*nRows*h;
nBoxes  = 1;
bottom  = zeros(m,2);
top     = zeros(m,2);
rowID   = cell(nRows,1);
% Draw a box at each node
for k = 1:nRows
  for j = 1:nBoxes
    bottom(i,:)   = [x+w/2 y ];
    top(i,:)      = [x+w/2 y+h];
    DrawBox(d.box{i},x,y,w,h,d);
    rowID{k}      = [rowID{k} i];
    i             = i + 1;
    x             = x + 1.5*w;
    if( i > length(d.box) )
      break;
    end
  end
  nBoxes   = 2*nBoxes;
  x        = -(0.25+0.5*(nBoxes/2-1))*w - nBoxes*w/2;
  y        = y - 1.5*h;
end
```

The remaining code draws the lines between the boxes.

```
for k = 1:length(rowID)-1
  iD = rowID{k};
  i0 = 0;
  % Work from left to right of the current row
  for j = 1:length(iD)
    x(1)  = bottom(iD(j),1);
    y(1)  = bottom(iD(j),2);
    iDT   = rowID{k+1};
    if( i0+1 > length(iDT) )
      break;
    end
    for i = 1:2
      x(2)  = top(iDT(i0+i),1);
      y(2)  = top(iDT(i0+i),2);
      line(x,y);
    end
    i0 = i0 + 2;
  end
end
axis off
```

The following built-in demo draws a binary tree. The demo creates three rows. It starts with the default data structure. You only have to add strings for the decision points. The boxes are in a flat list.

```
function Demo
% Draw a simple binary data treea

d           = DefaultDataStructure;
d.box{1}    = 'a > 0.1';
d.box{2}    = 'b > 0.2';
d.box{3}    = 'b > 0.3';
d.box{4}    = 'a > 0.8';
d.box{5}    = 'b > 0.4';
d.box{6}    = 'a > 0.2';
d.box{7}    = 'b > 0.3';

DrawBinaryTree( d );
```

Notice that it calls the subfunction `DefaultDataStructure` to initialize the demo.

```
%% DrawBinaryTree>DefaultDataStructure
function d = DefaultDataStructure
% Default data structure

d           = struct();
d.fontSize  = 12;
d.font      = 'courier';
d.w         = 1;
d.h         = 0.5;
d.box       = {};
```

---

■ **TIP** Always have the function return its default data structure. The default should have values that work.

---

It starts off with a default argument for name. The loops draw the boxes for the trees. `rowID` is a cell array. Each row in the cell is an array. A cell array allows each cell to be different. This makes it easy to have different length arrays in the cell. If you used a standard matrix you would need to resize rows as new rows were added. The binary tree resulting from the demo is shown in Figure 7.3. The text in the boxes could be anything you want.

The inputs for `box` could have been done in a loop. You could create them using `sprintf`. For example, for the first box you could write:

```
d.box{1}= sprintf('%s %s %3.1f','a','>',0.1);
```

and put similar code in a loop.

***Figure 7.3:*** Binary tree from the demo in `DrawBinaryTree`.

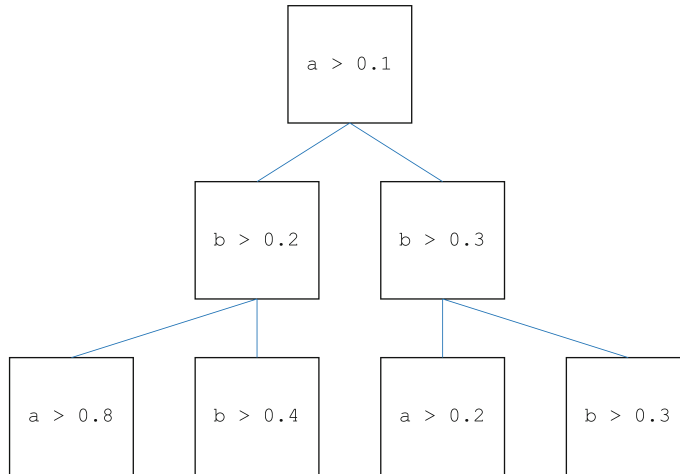


## 7.3 Implementation

Decision trees are the main focus of this chapter. We'll start with looking at how we determine if our decision tree is working correctly. We'll then hand-build a decision tree and finally write learning code to generate the decisions for each block of the tree.

### 7.3.1 Problem

We need to measure the homogeneity of a set of data at different nodes on the decision tree. A data set is homogeneous if the points are similar to each other. For example, if you were trying to study grade points in a school with an economically diverse population, you would want to know if your sample was all children from wealthy families. Our goal in the decision tree is to end up with homogeneous sets.

### 7.3.2 Solution

The solution is to implement the Gini impurity measure for a set of data. The function will return a single number as the homogeneity measure.

### 7.3.3 How It Works

The homogeneity measure is called the information gain (IG). The IG is defined as the increase in information by splitting at the node. This is:

$$\Delta I = I(p) - \frac{N_{c_1}}{N_p}I(c_1) - \frac{N_{c_2}}{N_p}I(c_2) \tag{7.2}$$

where $I$ is the impurity measure and $N$ is the number of samples at that node. If our tree is working, it should go down, eventually to zero or a very small number. In our training set, we know the class of each data point. Therefore, we can determine the IG. Essentially, we have

gained information if the mixing decreases in the child nodes. For example, in the first node in a decision tree, all the data are mixed together. There are two child nodes for the first node. After the decision in the first node, we expect that each child node will have more of one class than does the other child node. We look at the percentages of classes in each node and look for the maximum increase in nonhomogeneity.

There are three impurity measures:

- Gini impurity
- Entropy
- Classification error

Gini impurity, $I_G$, is the criterion to minimize the probability of misclassification. We don't want to push a sample into the wrong category.

$$I_G = 1 - \sum_{1}^{c} p(i|t)^2 \tag{7.3}$$

$p(i|t)$ is the proportion of the samples in class $c_i$ at node $t$. For a binary class entropy, $I_E$, is either zero or one.

$$I_E = 1 - \sum_{1}^{c} p(i|t) \log_2 p(i|t) \tag{7.4}$$

Classification error, $I_C$, is:

$$I_C = 1 - \max p(i|t) \tag{7.5}$$

We will use Gini impurity in the decision tree. The following code implements the Gini measure. The first part just decides whether it is initializing the function or updating. All data are saved in the data structure d. This is often easier than using global data. One advantage is that you can use the function multiple times in the same script or function without mixing up the persistent data in the function.

```
function [i, d] = HomogeneityMeasure( action, d, data )

if( nargin == 0 )
  if( nargout == 1 )
    i = DefaultDataStructure;
  else
    Demo;
  end
  return
end

switch lower(action)
  case 'initialize'
    d = Initialize( d, data );
```

```
    i = d.i;
  case 'update'
    d = Update( d, data );
    i = d.i;
  otherwise
    error('%s is not an available action',action);
end
```

Initialize initializes the data structure and computes the impurity measures for the data. There is one class for each different value of the data. For example, [1 2 3 3] would have three classes.

```
function d = Initialize( d, data )
%% HomogeneityMeasure>Initialize

m       = reshape(data,[],1);
c       = 1:max(m);
n       = length(m);
d.dist  = zeros(1,c(end));
d.class = c;
if( n > 0 )
  for k = 1:length(c)
    j          = find(m==c(k));
    d.dist(k) = length(j)/n;
  end
end
d.i = 1 - sum(d.dist.^2);
```

The demo is shown below. We try four different sets of data and get the measures. 0 is homogeneous. 1 means there is no data.

```
function d = Demo
% Demonstrate the homogeniety measure for a data set.

data    = [1 2 3 4 3 1 2 4 4 1 1 1 2 2 3 4]; fprintf(1,'%2.0f',data);
d       = HomogeneityMeasure;
[i, d]  = HomogeneityMeasure( 'initialize', d, data );
fprintf(1,'\nHomogeneity Measure %6.3f\n',i);
fprintf(1,'Classes              [%1d %1d %1d %1d]\n',d.class);
fprintf(1,'Distribution         [%5.3f %5.3f %5.3f %5.3f]\n',d.dist);

data = [1 1 1 2 2]; fprintf(1,'%2.0f',data);
[i, d] = HomogeneityMeasure( 'update', d, data );
fprintf(1,'\nHomogeneity Measure %6.3f\n',i);
fprintf(1,'Classes              [%1d %1d %1d %1d]\n',d.class);
fprintf(1,'Distribution         [%5.3f %5.3f %5.3f %5.3f]\n',d.dist);
```

```
data = [1 1 1 1]; fprintf(1,'%2.0f',data);
[i, d] = HomogeneityMeasure( 'update', d, data );
fprintf(1,'\nHomogeneity Measure %6.3f\n',i);
fprintf(1,'Classes              [%1d %1d %1d %1d]\n',d.class);
fprintf(1,'Distribution         [%5.3f %5.3f %5.3f %5.3f]\n',d.dist);

data = []; fprintf(1,'%2.0f',data);
[i, d] = HomogeneityMeasure( 'update', d, data );
fprintf(1,'\nHomogeneity Measure %6.3f\n',i);
fprintf(1,'Classes              [%1d %1d %1d %1d]\n',d.class);
fprintf(1,'Distribution         [%5.3f %5.3f %5.3f %5.3f]\n',d.dist);
```

`i` is the homogeneity measure. `d.dist` is the fraction of the data points that have the value of that class. The class is the distinct values. The outputs of the demo are shown below.

```
>> HomogeneityMeasure
 1 2 3 4 3 1 2 4 4 1 1 1 2 2 3 4
Homogeneity Measure  0.742
Classes             [1 2 3 4]
Distribution        [0.312 0.250 0.188 0.250]
 1 1 1 2 2
Homogeneity Measure  0.480
Classes             [1 2 3 4]
Distribution        [0.600 0.400 0.000 0.000]
 1 1 1 1
Homogeneity Measure  0.000
Classes             [1 2 3 4]
Distribution        [1.000 0.000 0.000 0.000]

Homogeneity Measure  1.000
Classes             [1 2 3 4]
Distribution        [0.000 0.000 0.000 0.000]
```

The second to last set has a zero, which is the desired value. If there are no inputs it returns 1, since by definition for a class to exist it must have members.

## 7.4 Creating a Decision tree

### 7.4.1 Problem

We want to implement a decision tree for classifying data with two parameters.

### 7.4.2 Solution

The solution is to write a binary decision tree function in MATLAB called `DecisionTree`.

### 7.4.3 How It Works

A decision tree [20] breaks down data by asking a series of questions about the data. Our decision trees will be binary in that there will be a yes or no answer to each question. For each feature in the data, we ask one question per decision node. This always splits the data into two child nodes. We will be looking at two parameters that determine class membership. The parameters will be numerical measurements.

At the following nodes, we ask additional questions, further splitting the data. Figure 7.4 shows the parent/child structure. We continue this process until the samples at each node are in one of the classes. At each node we want to ask the question that provides us with the most information about the class in which our samples reside. In constructing our decision tree for a two-parameter classification, we have two decisions at each node:

- Which parameter ($x$ or $y$) to check.
- What value of the parameter to use in the decision.

***Figure 7.4:*** Parent/child nodes.

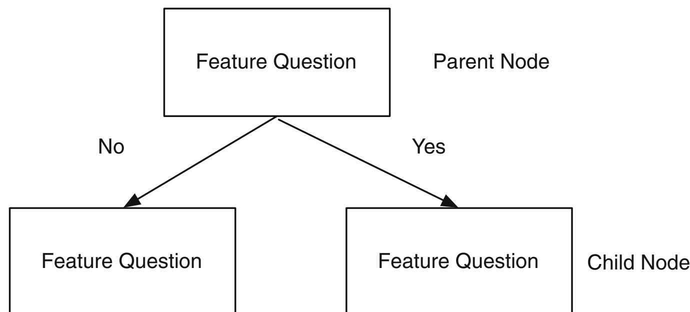


Training is done using the Gini values given in the previous recipe. We use the MATLAB-function `fminbnd` at each node, once for each of the two parameters. `fminbnd` is a one-dimensional local minimizer that finds the minimum of a function between two specified end-points. If you know the range of interest, then this is a very effective way to find the minimum.

$$\min_x f(x) \text{ such that } x_1 < x < x_2 \tag{7.6}$$

There are two actions, "train" and "test." "train" creates the decision tree and "test" runs the generated decision tree. You can also input your own decision tree. `FindOptimalAction` finds the parameter that minimizes the inhomogeneity on both sides of the division. The function called by `fminbnd` is `RHSGT`. We only implement the greater than action. The function call is:

```
function [d, r] = DecisionTree( action, d, t )
```

action is a string that is either "train" or "test." d is the data structure that defines the tree. t are the inputs for either training or testing. The outputs are the updated data structure and r with the results.

The function is first called with training data and the action is "train." The main function is short.

```
switch lower(action)
  case 'train'
    d = Training( d, t );
    d.box(1)

  case 'test'
    for k = 1:length(d.box)
      d.box(k).id = [];
    end
    [r, d] = Testing( d, t );
    for k = 1:length(d.box)
      d.box(k)
    end
  otherwise
    error('%s is not an available action',action);
end
```

We added the error case otherwise for completeness. Note that we use lower to eliminate case sensitivity. Training creates the decision tree. A decision tree is a set of boxes connected by lines. A parent box has two child boxes if it is a decision box. A class box has no children. The subfunction Training trains the tree. It adds boxes at each node.

```
%% DecisionTree>Training
function d = Training( d, t )
[n,m]    = size(t.x);
nClass   = max(t.m);
box(1)   = AddBox( 1, 1:n*m, [] );
box(1).child = [2 3];
[~, dH] = HomogeneityMeasure( 'initialize', d, t.m );

class    = 0;
nRow     = 1;
kR0      = 0;
kNR0     = 1; % Next row;
kInRow   = 1;
kInNRow = 1;
while( class < nClass )
  k    = kR0 + kInRow;
  idK    = box(k).id; % Data that is in the box and to use to compute
      the next action
  % Enter this loop if it not a non-decision box
  if( isempty(box(k).class) )
```

```
    [action, param, val, cMin]  = FindOptimalAction( t, idK, d.xLim,
       d.yLim, dH );
    box(k).value                = val;
    box(k).param                = param;
    box(k).action               = action;
    x                           = t.x(idK);
    y                           = t.y(idK);
    if( box(k).param == 1 ) % x
      id  = find(x >    d.box(k).value );
      idX = find(x <=   d.box(k).value );
    else % y
      id  = find(y >  d.box(k).value );
      idX = find(y <=   d.box(k).value );
    end
    % Child boxes
    if( cMin < d.cMin) % Means we are in a class box
      class          = class + 1;
      kN             = kNR0 + kInNRow;
      box(k).child   = [kN kN+1];
      box(kN)        = AddBox( kN, idK(id), class  );
      class          = class + 1;
      kInNRow        = kInNRow + 1;
      kN             = kNR0 + kInNRow;
      box(kN)        = AddBox( kN, idK(idX), class );
      kInNRow        = kInNRow + 1;
    else
      kN             = kNR0 + kInNRow;
      box(k).child   = [kN kN+1];
      box(kN)        = AddBox( kN, idK(id)  );
      kInNRow        = kInNRow + 1;
      kN             = kNR0 + kInNRow;
      box(kN)        = AddBox( kN, idK(idX) );
      kInNRow        = kInNRow + 1;
    end
  end
      % Update current row
  kInRow   = kInRow + 1;
  if( kInRow > nRow )
    kR0        = kR0 + nRow;
    nRow       = 2*nRow; % Add two rows
    kNR0       = kNR0 + nRow;
    kInRow     = 1;
    kInNRow    = 1;
  end
end

for k = 1:length(box)
```

```
  if( ~isempty(box(k).class) )
    box(k).child = [];
  end
  box(k).id = [];
  fprintf(1,'Box %3d action %2s Value %4.1f\n',k,box(k).action,box(k)
      .value);
end

d.box = box;
```

We use `fminbnd` to find the optimal switch point. We need to compute the homogeneity on both sides of the switch and sum the values. The sum is minimized by `fminbnd` in the subfunction `FindOptimalAction`. This code is designed for rectangular region classes. Other boundaries won't necessarily work correctly. The code is fairly involved. It needs to keep track of the box numbering to make the parent child connections. When the homogeneity measure is low enough, it marks the boxes as containing the classes.

The data structure `box` has multiple fields. One is the action to be taken in a decision box. The `param` is 1 for `x` and anything else for `y`. That determines if it is making the decision based on `x` or `y`. The `value` is the value used in the decision. `child` are indexes to the box children. The remaining code determines which row the box is in. `class` boxes have no children. The fields are shown in Table 7.1.

***Table 7.1:*** Box Data Structure Fields

| Field | Decision Box | Class Box |
|---|---|---|
| `action` | String | Not used |
| `value` | Value to be used in the decision | Not used |
| `param` | x or y | Not used |
| `child` | Array with two children | Empty |
| `id` | Empty | ID of data in the class |
| `class` | Class ID | Not used |

## 7.5 Creating a Handmade Tree

### 7.5.1 Problem

We want to test a handmade decision tree.

### 7.5.2 Solution

The solution is to write a script to test a handmade decision tree.

### 7.5.3 How It Works

We write the test script `SimpleClassifierDemo` shown below. It uses the `'test'` action for `DecisionTree`. It generates $5^2$ points. We create rectangular regions so that the face arrays have four elements for each polygon. `DrawBinaryTree` draws the tree.

```
d = DecisionTree;

% Vertices for the sets
v = [ 0 0; 0 4; 4 4; 4 0; 2 4; 2 2; 2 0; 0 2; 4 2];

% Faces for the sets
f = { [6 5 2 8] [6 7 4 9] [6 9 3 5] [1 7 6 8] };

% Generate the testing set
pTest  = ClassifierSets( 5, [0 4], [0 4], {'width', 'length'}, v, f,
   'Testing Set' );

% Test the tree
[d, r] = DecisionTree( 'test',  d, pTest  );

q = DrawBinaryTree;
c = 'xy';
for k = 1:length(d.box)
  if( ~isempty(d.box(k).action) )
    q.box{k} = sprintf('%c %s %4.1f',c(d.box(k).param),d.box(k).
        action,d.box(k).value);
  else
    q.box{k} = sprintf('Class %d',d.box(k).class);
  end
end
DrawBinaryTree(q);

m = reshape(pTest.m,[],1);

for k = 1:length(r)
  fprintf(1,'Class %d\n',m(r{k}(1)));
  for j = 1:length(r{k})
    fprintf(1,'%d ',r{k}(j));
  end
  fprintf(1,'\n')
end
```

`SimpleClassifierDemo` uses the hand-built example in `DecisionTree`.

```
function d = DefaultDataStructure
%% DecisionTree>DefaultDataStructure
% Generate a default data structure
d.tree          = DrawBinaryTree;
```

```
d.threshold      = 0.01;
d.xLim           = [0 4];
d.yLim           = [0 4];
d.data       = [];
d.cMin           = 0.01;
d.box(1)     = struct('action','>','value',2,'param',1,'child',[2 3],'
    id',[],'class',[]);
d.box(2)     = struct('action','>','value',2,'param',2,'child',[4 5],'
    id',[],'class',[]);
d.box(3)     = struct('action','>','value',2,'param',2,'child',[6 7],'
    id',[],'class',[]);

for k = 4:7
  d.box(k) = struct('action','','value',0,'param',0,'child',[],'id'
      ,[],'class',[]);
end
```

Figure 7.5 shows the results from `SimpleClassifierDemo`. There are four rectangular areas, which are our sets.

***Figure 7.5:*** Data and classes in the test set.

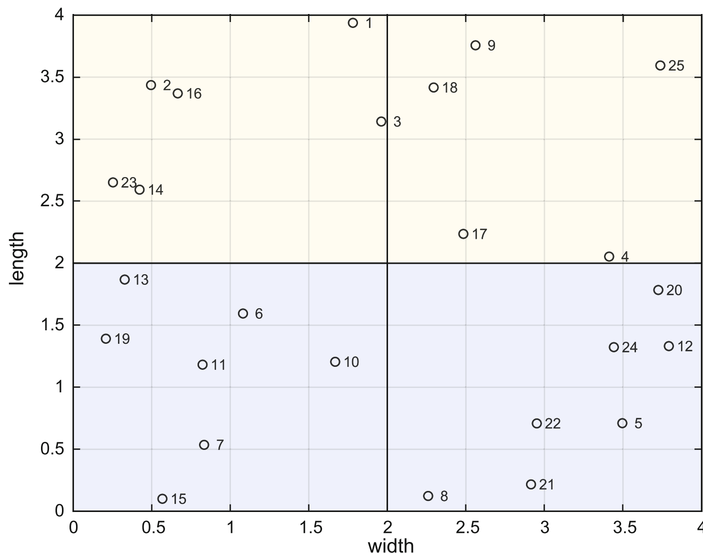


We can create a decision tree by hand as shown Figure 7.6.

The decision tree sorts the samples into the four sets. In this case, we know the boundaries and can use them to write the inequalities. In software, we will have to determine what values provide the shortest branches. The following is the output of `SimpleClassifierDemo`. The decision tree properly classifies all of the data.

***Figure 7.6:*** A manually created decision tree. The drawing is generated by `DecisionTree`. The last row of boxes is the data sorted into the four classes. The last nodes are the classes. Each box is a decision tree node.

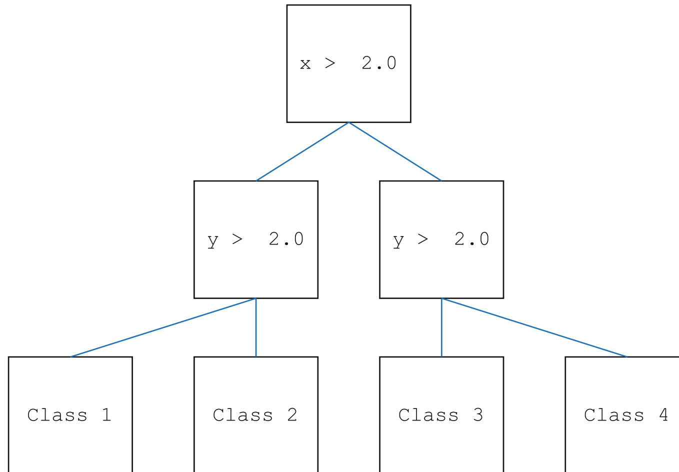


```
>> SimpleClassifierDemo

Class 3
4 6 9 13 18
Class 2
7 14 17 21
Class 1
1 2 5 8 10 11 12 23 25
Class 4
3 15 16 19 20 22 24
```

## 7.6 Training and Testing

### 7.6.1 Problem

We want to train our decision tree and test the results.

### 7.6.2 Solution

We replicated the previous recipe, only this time we have `DecisionTree` create the decision tree instead of creating it by hand.

### 7.6.3 How It Works

`TestDecisionTree` trains and tests the decision tree. It is very similar to the code for the hand-built decision tree demo, `SimpleClassifierDemo`. Once again, we use rectangles for the regions.

```
% Vertices for the sets
v = [ 0 0; 0 4; 4 4; 4 0; 2 4; 2 2; 2 0; 0 2; 4 2];

% Faces for the sets
f = { [6 5 2 8] [6 7 4 9] [6 9 3 5] [1 7 6 8] };

% Generate the training set
pTrain = ClassifierSets( 40, [0 4], [0 4], {'width', 'length'},...
  v, f, 'Training Set' );

% Create the decision tree
d      = DecisionTree;
d      = DecisionTree( 'train', d, pTrain );

% Generate the testing set
pTest  = ClassifierSets( 5, [0 4], [0 4], {'width', 'length'},...
  v, f, 'Testing Set' );

% Test the tree
[d, r] = DecisionTree( 'test',  d, pTest  );

q = DrawBinaryTree;
c = 'xy';
for k = 1:length(d.box)
  if( ~isempty(d.box(k).action) )
    q.box{k} = sprintf('%c %s %4.1f',c(d.box(k).param),...
      d.box(k).action,d.box(k).value);
  else
    q.box{k} = sprintf('Class %d',d.box(k).class);
  end
end
DrawBinaryTree(q);

m = reshape(pTest.m,[],1);

for k = 1:length(r)
  fprintf(1,'Class %d\n',m(r{k}(1)));
  for j = 1:length(r{k})
    fprintf(1,'%d ',r{k}(j));
  end
```

It uses `ClassifierSets` to generate the training data. The output includes the coordinates and the sets in which they fall. We then create the default data structure and call `DecisionTree` in training mode.

The tree is shown in Figure 7.9. The training data are shown in Figure 7.7 and the testing data in Figure 7.8. We need enough testing data to fill the classes. Otherwise, the decision tree generator may draw the lines to encompass just the data in the training set.

***Figure 7.7:*** The training data. A large amount of data is needed to fill the classes.

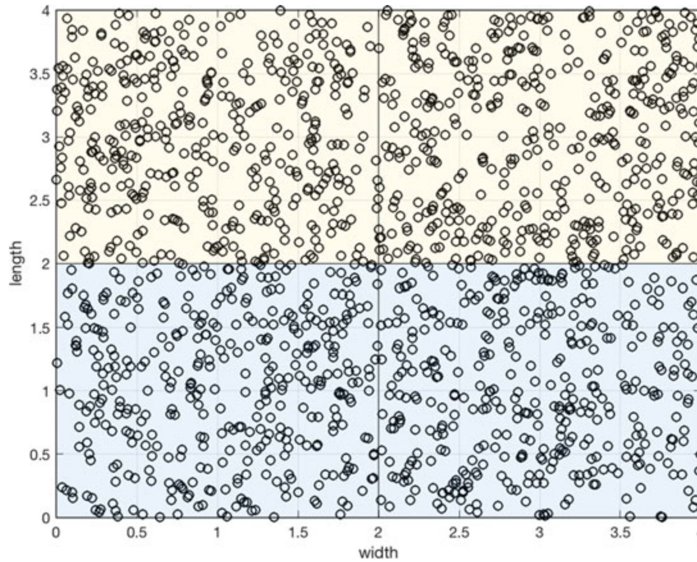


The results are similar to the simple test.

```
Class 3
1 14 16 21 23
Class 2
2 4 5 6 9 13 17 18 19 20 25
Class 1
3 7 8 10 11 15 24
Class 4
12 22
```

The generated tree separates the data effectively.

***Figure 7.8:*** The testing data.

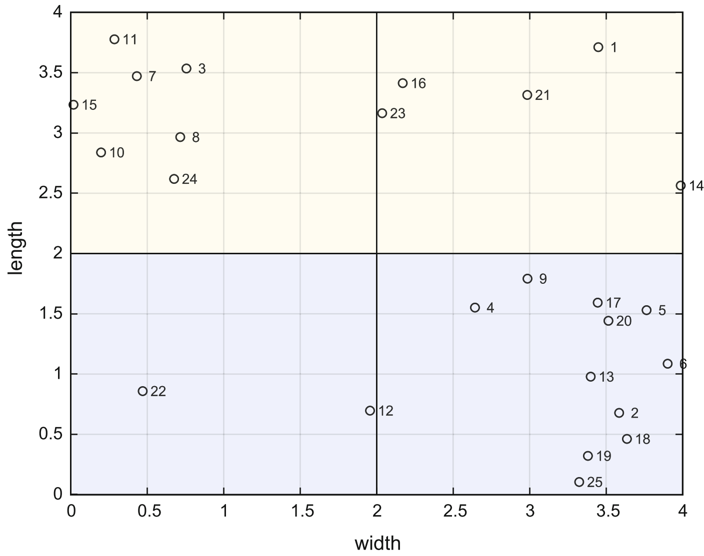


***Figure 7.9:*** The tree derived from the training data. It is essentially the same as the hand-derived tree. The values in the generated tree are not exactly 2.0.

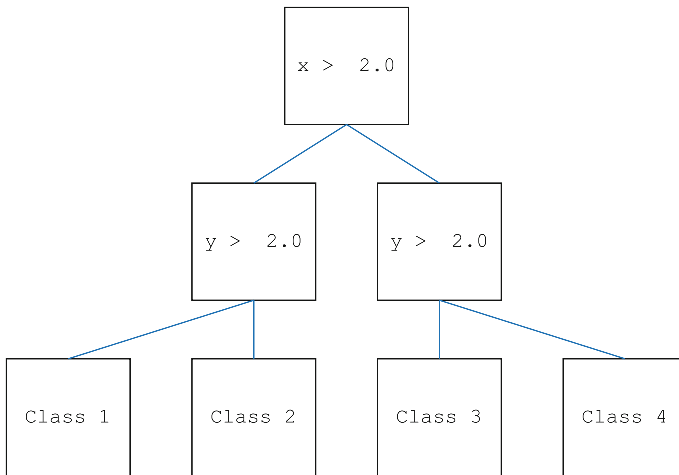

## 7.7 Summary

This chapter has demonstrated data classification using decision trees in MATLAB. We also wrote a new graphics function to draw decision trees. The decision tree software is not general purpose, but can serve as a guide to more general purpose code. Table 7.2 lists the functions and scripts included in the companion code.

***Table 7.2:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| ClassifierSets | Generate data for classification or training. |
| DecisionTree | Implements a decision tree to classify data. |
| DrawBinaryTree | Generates data for classification or training. |
| HomogeneityMeasure | Computes Gini impurity. |
| SimpleClassifierDemo | Demonstrates decision tree testing. |
| SimpleClassifierExample | Generates data for a simple problem. |
| TestDecisionTree | Tests a decision tree. |

# CHAPTER 8

# Introduction to Neural Nets

Neural networks, or neural nets, are a popular way of implementing machine "intelligence." The idea is that they behave like the neuron in a brain. In our taxonomy, neural nets fall into the category of true machine learning, as shown on the right.

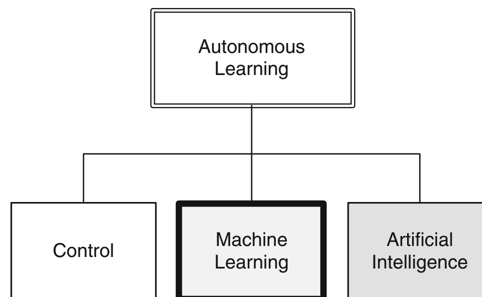


In this chapter, we will explore how neural nets work, starting with the most fundamental idea with a single neuron and working our way up to a multi-layer neural net. Our example for this will be a pendulum. We will show how a neural net can be used to solve the prediction problem. This is one of the two uses of a neural net, prediction and categorization. We'll start with a simple categorization example. We'll do more sophisticated categorization neural nets in chapters 9 and 10.

## 8.1 Daylight Detector

### 8.1.1 Problem

We want to use a simple neural net to detect daylight.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_8

### 8.1.2 Solution

Historically, the first neuron was the perceptron. This is a neural net with an activation function that is a threshold. Its output is either 0 or 1. This is not really useful for problems such as the pendulum angle estimation covered in the remaining recipes of this chapter. However, it is well suited to categorization problems. We will use a single perceptron in this example.

### 8.1.3 How It Works

Suppose our input is a light level measured by a photo cell. If you weight the input so that 1 is the value defining the brightness level at twilight, you get a sunny day detector.

This is shown in the following script, `SunnyDay`. The script is named after the famous neural net that was supposed to detect tanks, but instead detected sunny days; this was due to all the training photos of tanks being taken, unknowingly, on a sunny day, whereas all the photos without tanks were taken on a cloudy day. The solar flux is modeled using a cosine and scaled so that it is 1 at noon. Any value greater than 0 is daylight.

```
%% The data
t = linspace(0,24);         % time, in hours
d = zeros(1,length(t));
s = cos((2*pi/24)*(t-12)); % solar flux model

%% The activation function
% The nonlinear activation function which is a threshold detector
j    = s < 0;
s(j) = 0;
j    = s > 0;
d(j) = 1;

%% Plot the results
PlotSet(t,[s;d],'x␣label','Hour', 'y␣label',...
  {'Solar␣Flux', 'Day/Night'}, 'figure␣title','Daylight␣Detector',...
  'plot␣title', 'Daylight␣Detector');
set([subplot(2,1,1) subplot(2,1,2)],'xlim',[0 24],'xtick',[0 6 12 18
   24]);
```

Figure 8.1 shows the detector results. The `set(gca,...)` code sets the x-axis ticks to end at exactly 24 h. This is a really trivial example, but does show how categorization works. If we had multiple neurons with thresholds set to detect sun light levels within bands of solar flux, we would have a neural net sun clock.

***Figure 8.1:*** The daylight detector.

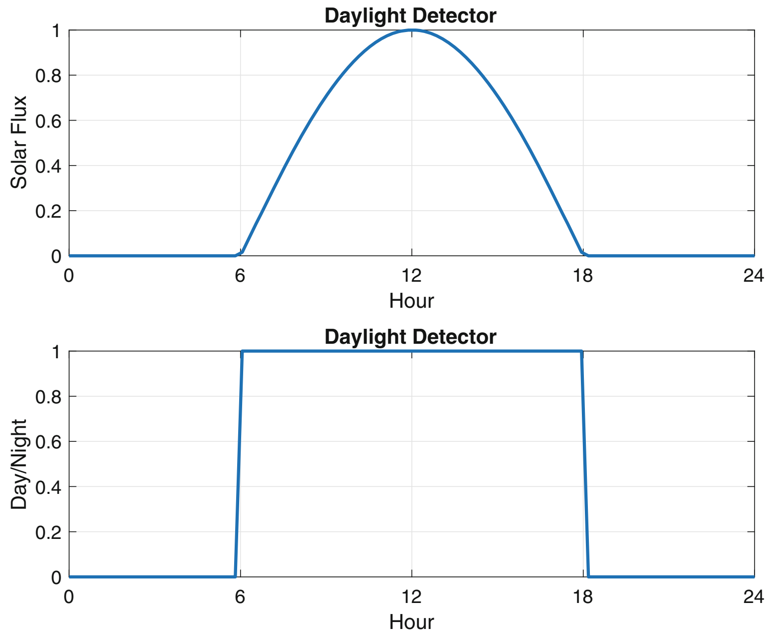


## 8.2 Modeling a Pendulum

### 8.2.1 Problem

We want to implement the dynamics of a pendulum as shown in Figure 8.2. The pendulum will be modeled as a point mass with a rigid connection to its pivot. The rigid connection is a rod that cannot contract or expand.

***Figure 8.2:*** A pendulum. The motion is driven by the acceleration of gravity.

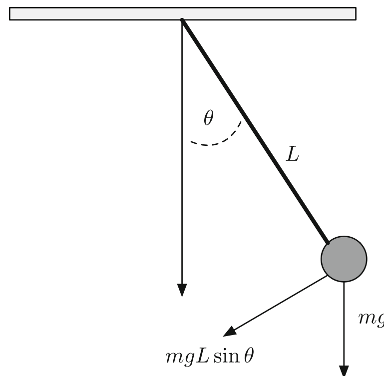

### 8.2.2 Solution

The solution is to write a pendulum dynamics function in MATLAB. The dynamics will be written in torque form, that is, we will model it as rigid body rotation. Rigid body rotation is what happens when you spin a wheel. It will use the `RungeKutta` integration routine in the General folder of the included toolbox to integrate the equations of motion.

### 8.2.3 How It Works

Figure 8.2 shows the pendulum. The easiest way to get the equations is to write it as a torque problem, that is, as rigid body rotation. When you look at a two-dimensional pendulum, it moves in a plane and its location has $x$ and $y$ coordinates. However, these two coordinates are constrained by the fixed pendulum of length $L$. We can write:

$$L^2 = x^2 + y^2 \tag{8.1}$$

where $L$ is the length of the rod and a constant and $x$ and $y$ are the coordinates in the plane. They are also the degrees of freedom in the problem. This shows that $x$ is uniquely determined by $y$. If we write:

$$x = L \sin\theta \tag{8.2}$$

$$y = L \cos\theta \tag{8.3}$$

where $\theta$ is the angle from vertical, i.e., it is zero when the pendulum is hanging straight down, we see that we need only one degree of freedom, $\theta$, to model the motion. So our force problem becomes a rigid body rotational motion problem. The torque is related to the angular acceleration by the inertia as:

$$T = I\frac{d^2\theta}{dt^2} \tag{8.4}$$

where $I$ is the inertia and $T$ is the torque. The inertia is constant and depends on the square of the pendulum length and the mass $m$:

$$I = mL^2 \tag{8.5}$$

The torque is produced by the component of the gravitational force, $mg$, which is perpendicular to the pendulum, where $g$ is the acceleration of gravity. Recall that torque is the applied force, $mg\sin\theta$, times the moment arm, in this case $L$. The torque is therefore:

$$T = -mgL\sin\theta \tag{8.6}$$

The equations of motion are then:

$$-mgL\sin\theta = mL^2\frac{d^2\theta}{dt^2} \tag{8.7}$$

or simplifying:

$$\frac{d^2\theta}{dt^2} + \left(\frac{g}{mL}\right)\sin\theta = 0 \tag{8.8}$$

We set:

$$\frac{g}{mL} = \Omega^2 \tag{8.9}$$

where $\Omega$ is the frequency of the pendulum's oscillation. This equation is nonlinear because of the $\sin\theta$. We can linearize it about small angles, $\theta$, about vertical. For small angles:

$$\sin\theta \approx \theta \tag{8.10}$$

$$\cos\theta \approx 1 \tag{8.11}$$

to get the linear constant coefficient equation. The linear version of sine comes from the Taylor's series expansion:

$$\sin\theta = \theta - \frac{\theta^3}{6} + \frac{\theta^5}{120} - \frac{\theta^7}{5040} + \cdots \tag{8.12}$$

You can see that the first term is a pretty good approximation around $\theta = 0$, which is when the pendulum is hanging vertically. We can actually apply this to any angle. Let $\theta = \theta + \theta_k$, where $\theta_k$ is our current angle and $\theta$ is now small. We can expand the sine term:

$$\sin(\theta + \theta_k) = \sin\theta\cos\theta_k + \sin\theta_k\cos\theta \approx \theta\cos\theta_k + \sin\theta_k \tag{8.13}$$

We get a linear equation with a new torque term and a different coefficient for $\theta$.

$$\frac{d^2\theta}{dt^2} + \cos\theta_k\Omega^2\theta = -\Omega^2\sin\theta_k \tag{8.14}$$

This tells us that a linear approximation may be useful, regardless of the current angle.

Our final equations (nonlinear and linear) are:

$$\frac{d^2\theta}{dt^2} + \Omega^2\sin\theta = 0 \tag{8.15}$$

$$\frac{d^2\theta}{dt^2} + \Omega^2\theta \approx 0 \tag{8.16}$$

The dynamical model is in the following code, with an excerpt from the header. This can be called by the MATLAB Recipes `RungeKutta` function or any MATLAB integrator. There is an option to use either the full nonlinear dynamics or the linearized form of the dynamics, using a Boolean field called `linear`. The state vector has the angle as the first element and the angle derivative, or angular velocity $\omega$, as the second element. Time, the first input, is not used because it only appears in the equations as $dt$, so it is replaced with a tilde. The output is the derivative, `xDot`, of the state `x`. If no inputs are specified, the function will return the default data structure `d`.

```
%  x        (2,1) State vector [theta;theta dot]
%  d        (.)   Data structure
%                 .linear  (1,1) If true use a linear model
%                 .omega   (1,1) Input gainss

function xDot = RHSPendulum( ~, x, d )

if( nargin < 1 )
  xDot = struct('linear',false,'omega',0.5);
  return
end

if( d.linear )
  f = x(1);
else
  f = sin(x(1));
end

xDot = [x(2);-d.omega^2*f];
```

The code for xDot has two elements. The first element is just the second element of the state, because the derivative of the angle is the angular velocity. The second term is the angular acceleration computed using our equations. The set of differential equations that is implemented is a set of first-order differential equations:

$$\frac{d\theta}{dt} = \omega \tag{8.17}$$

$$\frac{d\omega}{dt} = -\Omega^2 \sin\theta \tag{8.18}$$

First order means there are only first derivatives on the left-hand side.

The script PendulumSim, shown below, simulates the pendulum by integrating the dynamical model. Setting the data structure field linear to true gives the linear model. Note that the state is initialized with a large initial angle of 3 radians to highlight the differences between the models.

```
%% Pendulum simulation
%% Initialize the simulation
n              = 1000;          % Number of time steps
dT             = 0.1;           % Time step (sec)
dRHS           = RHSPendulum;     % Get the default data structure
dRHS.linear    = false;         % true for linear model

%% Simulation
xPlot          = zeros(2,n);
theta0         = 3;            % radians
x              = [theta0;0];   % [angle;velocity]
```

```
for k = 1:n
  xPlot(:,k)   = x;
  x            = RungeKutta( @RHSPendulum, 0, x, dT, dRHS );
end

%% Plot the results
yL       = {'\theta␣(rad)' '\omega␣(rad/s)'};
[t,tL]   = TimeLabel(dT*(0:n-1));

PlotSet( t, xPlot, 'x␣label', tL, 'y␣label', yL, ...
        'plot␣title', 'Pendulum', 'figure␣title', 'Pendulum␣State' );
```

Figure 8.3 shows the results of the two models. The period of the nonlinear model is not the same as that of the linear model.

***Figure 8.3:*** A pendulum modeled by the linear and nonlinear equations. The period for the nonlinear model is not the same as for the linear model. The left-hand plot is linear and the right nonlinear.

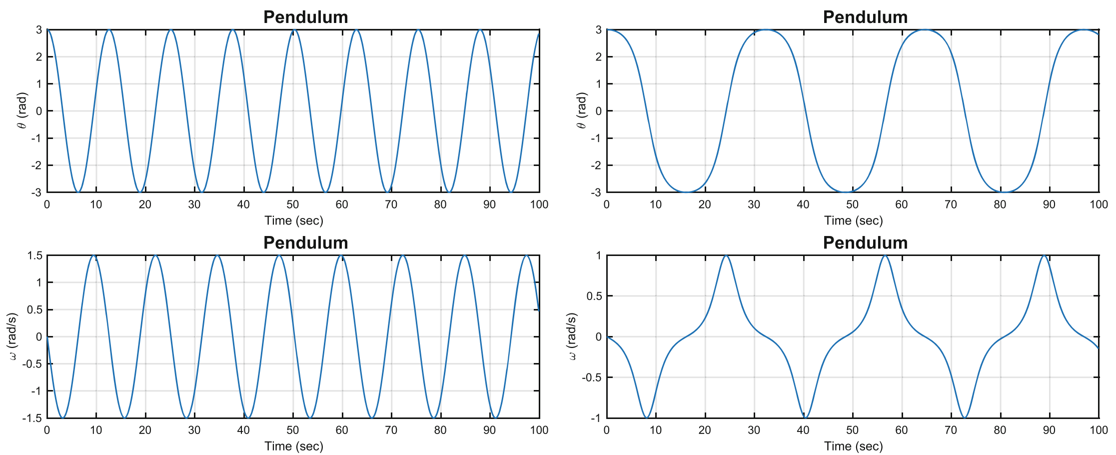


## 8.3 Single Neuron Angle Estimator

### 8.3.1 Problem

We want to use a simple neural net to estimate the angle between the rigid pendulum and vertical.

### 8.3.2 Solution

We will derive the equations for a linear estimator and then replicate it with a neural net consisting of a single neuron.

### 8.3.3 How It Works

Let's first look at a single neuron with two inputs. This is shown in Figure 8.4. This neuron has inputs $x_1$ and $x_2$, a bias $b$, weights $w_1$ and $w_2$, and a single output $z$. The activation function $\sigma$ takes the weighted input and produces the output.

$$z = \sigma(w_1 x_1 + w_2 x_2 + b) \tag{8.19}$$

***Figure 8.4:*** A two input neuron.

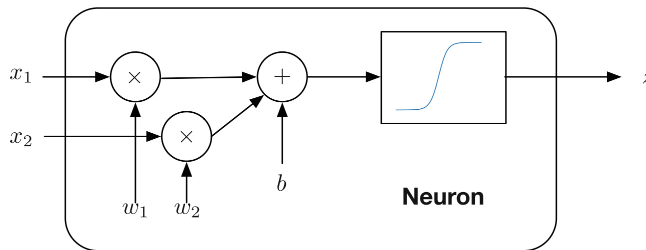


Let's compare this with a real neuron as shown in Figure 8.5. A real neuron has multiple inputs via the dendrites. Some of these branch, which means that multiple inputs can connect to the cell body through the same dendrite. The output is via the axon. Each neuron has one output. The axon connects to a dendrite through the synapse. Signals pass from the axon to the dendrite via a synapse.

***Figure 8.5:*** A real neuron can have 10,000 inputs!

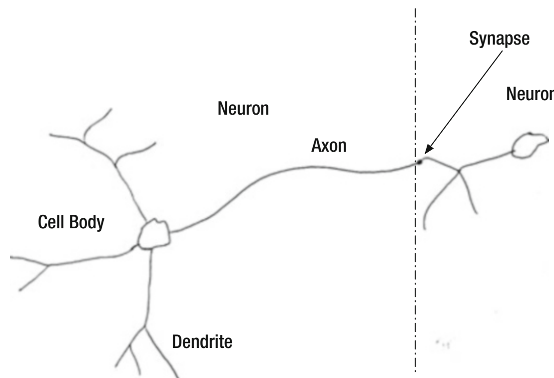


There are numerous commonly used activation functions. We show three:

$$\sigma(y) = \tanh(y) \tag{8.20}$$

$$\sigma(y) = \frac{2}{1 - e^{-y}} - 1 \tag{8.21}$$

$$\sigma(y) = y \tag{8.22}$$

The exponential one is normalized and offset from zero so that it ranges from -1 to 1. The following code in the script OneNeuron computes and plots these three activation functions for an input q.

```
%% Look at the activation functions
q        = linspace(-4,4);
v1       = tanh(q);
v2       = 2./(1+exp(-q)) - 1;

PlotSet(q,[v1;v2;q],'x␣label','Input', 'y␣label',...
  'Output', 'figure␣title','Activation␣Functions','plot␣title', '
     Activation␣Functions',...
  'plot␣set',{[1 2 3]},'legend',{{'Tanh','Exp','Linear'}});
```

Figure 8.6 shows the three activation functions on one plot.

***Figure 8.6:*** The three activation functions.

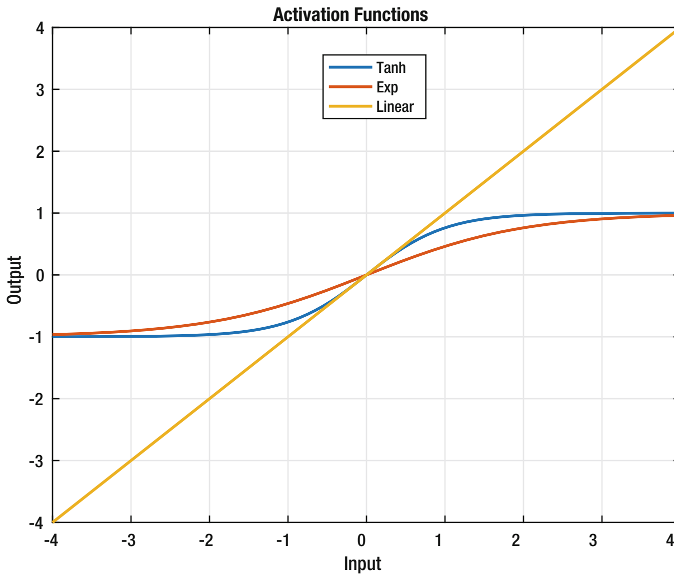


Activation functions that saturate model a biological neuron that has a maximum firing rate. These particular functions also have good numerical properties that are helpful in learning.

Now that we have defined our neuron model, let's return to the pendulum dynamics. The solution to the linear pendulum equation is:

$$\theta = a \sin \Omega t + b \cos \Omega t \tag{8.23}$$

Given the initial angle $\theta_0$ and angular rate $\dot{\theta}_0$, we get the angle as a function of time:

$$\theta(t) = \frac{\dot{\theta}_0}{\Omega} \sin \Omega t + \theta_0 \cos \Omega t \tag{8.24}$$

For small $\Omega t$:

$$\theta(t) = \dot{\theta}_0 t + \theta_0 \tag{8.25}$$

which is a linear equation. Change this to a discrete time problem:

$$\theta_{k+1} = \dot{\theta}_k \Delta t + \theta_k \tag{8.26}$$

where $\Delta t$ is the time step between measurements, $\theta_k$ is the current angle, and $\theta_{k+1}$ is the angle at the next step. The linear approximation to the angular rate is:

$$\dot{\theta}_k = \frac{\theta_k - \theta_{k-1}}{\Delta t} \tag{8.27}$$

so combining Eqs. 8.26 and 8.27, our "estimator" is

$$\theta_{k+1} = 2\theta_k - \theta_{k-1} \tag{8.28}$$

This is quite simple. It does not need to know the time step.

Let's do the same thing with a neural net. Our neuron inputs are $x_1$ and $x_2$. If we set:

$$x_1 = \theta_k \tag{8.29}$$

$$x_2 = \theta_{k-1} \tag{8.30}$$

$$w_1 = 2 \tag{8.31}$$

$$w_2 = -1 \tag{8.32}$$

$$b = 0 \tag{8.33}$$

we get

$$z = \sigma(2\theta_k - \theta_{k-1}) \tag{8.34}$$

which is, aside from the activation function $\sigma$, our estimator.

Continuing through `OneNeuron`, the following code implements the estimators. We input a pure sine wave that is only valid for small pendulum angles. We then compute the neuron with the linear activation function and then the `tanh` activation function. Note that the variable `thetaN` is equivalent to using the linear activation function.

```
%% Look at the estimator for a pendulum
omega   = 1;                     % pendulum frequency in rad/s
t       = linspace(0,20);
theta   = sin(omega*t);
thetaN  = 2*theta(2:end) - theta(1:end-1); % linear estimator for "
   next" theta
truth   = theta(3:end);
tOut    = t(3:end);
thetaN  = thetaN(1:end-1);
```

```
% Apply the activation function
z = tanh(thetaN);

PlotSet(tOut,[truth;thetaN;z],'x␣label','Time␣(s)', 'y␣label',...
  'Next␣angle', 'figure␣title','One␣neuron','plot␣title', 'One␣neuron
      ',...
  'plot␣set',{[1 2 3]},'legend',{{'True','Estimate','Neuron'}});
```

Figure 8.7 shows the two neuron outputs, linear and `tanh`, compared with the truth. The one with the linear activation function matches the truth very well. The `tanh` does not, but that is to be expected as it saturates.

***Figure 8.7:*** The true pendulum dynamics compared with the linear and `tanh` neuron output.

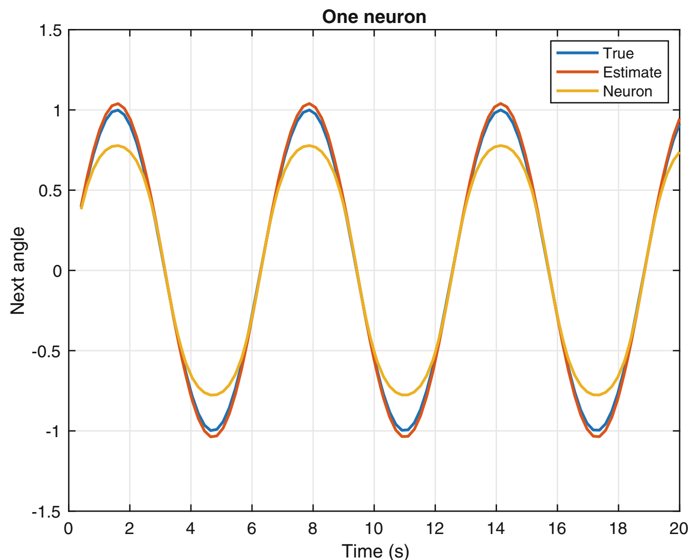


The one neuron function with the linear activation function is the same as the estimator by itself. Usually output nodes, and this neural net has only an output node, have linear activation functions. This makes sense, otherwise the output would be limited to the saturation value of the activation functions, as we have seen with `tanh`. With any other activation function, the output does not produce the desired result. This particular example is one in which a neural net doesn't really give us any advantage and was chosen because it reduces to a simple linear estimator. For more general problems, with more inputs and nonlinear dependencies among the inputs, activation functions that have saturation may be valuable.

For this, we will need a multi-neuron net to be discussed in the last section of the chapter. Note that even the neuron with the linear activation function does not quite match the truth value. If we were to actually use the linear activation function with the nonlinear pendulum, it would not work very well. A nonlinear estimator would be complicated, but a neural net with multiple layers (deep learning) could be trained to cover a wider range of conditions.

## 8.4 Designing a Neural Net for the Pendulum

### 8.4.1 Problem

We want to estimate angles for a nonlinear pendulum.

### 8.4.2 Solution

We will use `NeuralNetMLFF` to build a neural net from training sets. (MLFF stands for multi-layer, feed-forward). We will run the net using `NeuralNetMLFF`. The code for `NeuralNetMLFF` is included with the neural net developer GUI in the next chapter.

### 8.4.3 How It Works

The script for this recipe is `NNPendulumDemo`. The first part generates the test data running the same simulation as `PendulumSim.m` in Recipe 8.2. We calculate the period of the pendulum in order to set the simulation time step at a small fraction of the period. Note that we will use `tanh` as the activation function for the net.

```
% Demo parameters
nSamples     = 800;          % Samples in the simulation
nRuns        = 2000;         % Number of training runs
activation   = 'tanh';       % activation function

omega        = 0.5;          % frequency in rad/s
tau          = 2*pi/omega;   % period in secs
dT           = tau/100;      % sample at a rate of 20*omega

rng(100);              % consistent random number generator

%% Initialize the simulation RHS
dRHS         = RHSPendulum;  % Get the default data structure
dRHS.linear  = false;
dRHS.omega   = omega;

%% Simulation
nSim    = nSamples + 2;
x       = zeros(2,nSim);
theta0  = 0.1;               % starting position (angle)
x(:,1)  = [theta0;0];
for k = 1:nSim-1
  x(:,k+1) = RungeKutta( @RHSPendulum, 0, x(:,k), dT, dRHS );
end
```

The next block defines the network and trains it using `NeuralNetTraining`. `NeuralNetTraining` and `NeuralNetMLFF` are described in the next chapter. Briefly, we define a first layer with three neurons and a second output layer with a single neuron; the network has two inputs, which are the previous two angles.

```
        'plot␣title', 'Pendulum', 'figure␣title', 'Pendulum␣State' );

%% Define a network with two inputs, three inner nodes, and one
   output
layer              = struct;
layer(1,1).type    = activation;
layer(1,1).alpha = 1;
layer(2,1).type    = 'sum'; %'sum';
layer(2,1).alpha = 1;

% Thresholds
layer(1,1).w0 = rand(3,1) - 0.5;
layer(2,1).w0 = rand(1,1) - 0.5;

% Weights w(i,j) from jth input to ith node
layer(1,1).w  = rand(3,2) - 0.5;
layer(2,1).w  = rand(1,3) - 0.5;

%% Train the network
% Order the samples using a random list
kR          = ceil(rand(1,nRuns)*nSamples);
thetaE      = x(1,kR+2); % Angle to estimate
theta       = [x(1,kR);x(1,kR+1)]; % Previous two angles
e           = thetaE - (2*theta(1,:) - theta(2,:));
[w,e,layer] = NeuralNetTraining( theta, thetaE, layer );

PlotSet(1:length(e), e.^2, 'x␣label','Sample', 'y␣label','Error^2'
   ,...
  'figure␣title','Training␣Error','plot␣title','Training␣Error','plot
     ␣type','ylog');

% Assemble a new network with the computed weights
layerNew             = struct;
layerNew(1,1).type   = layer(1,1).type;
```

The training data structure includes the weights to be computed. It defines the number of layers and the type of activation function. The initial weights are random. Training returns the new weights and the training error. We pass the training data in a random order to the function using the index array k. This gives better results than if we passed it in the original order. We also send the same training data multiple times using the parameter nRuns. Figure 8.8 shows the training error. It looks good. To see the weights that were calculated, just display w at the command line. For example, the weights of the output node are now:

```
>> w(2)
ans =
  struct with fields:
```

***Figure 8.8:*** Training error.

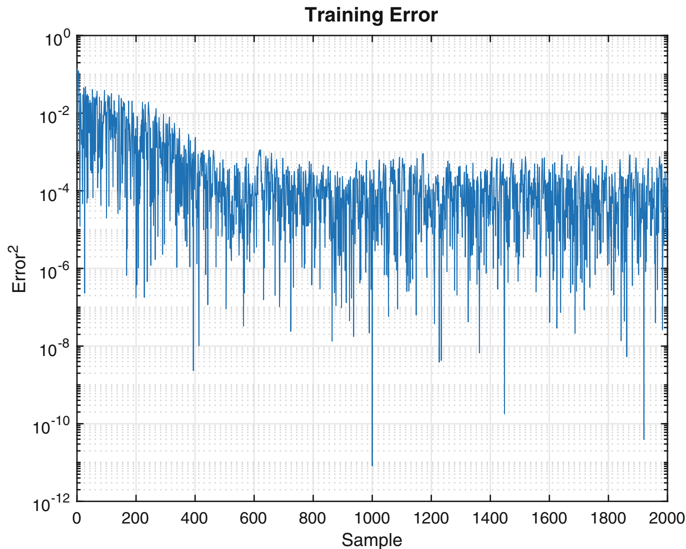


```
       w: [-0.67518 -0.21789 -0.065903]
      w0: -0.014379
    type: 'tanh'
```

We test the neural net in the last block of code. We rerun the simulation and then run the neural net using `NeuralNetMLFF`. Note that you may choose to initialize the simulation with a different starting point than in the training data by changing the value of `thetaD`.

```
layerNew                    = struct;
layerNew(1,1).type          = layer(1,1).type;
layerNew(1,1).w             = w(1).w;
layerNew(1,1).w0            = w(1).w0;
layerNew(2,1).type          = layer(2,1).type; %'sum';
layerNew(2,1).w             = w(2).w;
layerNew(2,1).w0            = w(2).w0;
network.layer               = layerNew;

%% Simulate the pendulum with a different starting point
x(:,1)          = [0.1;0];

%% Simulate the pendulum and test the trained network
% Choose the same or a different starting point and simulate
thetaD = 0.5;
x(:,1) = [thetaD;0];
for k = 1:nSim-1
  x(:,k+1) = RungeKutta( @RHSPendulum, 0, x(:,k), dT, dRHS );
end
```

***Figure 8.9:*** Neural net results: the simulated state, the testing error, and the truth angles compared with the neural net's estimate.

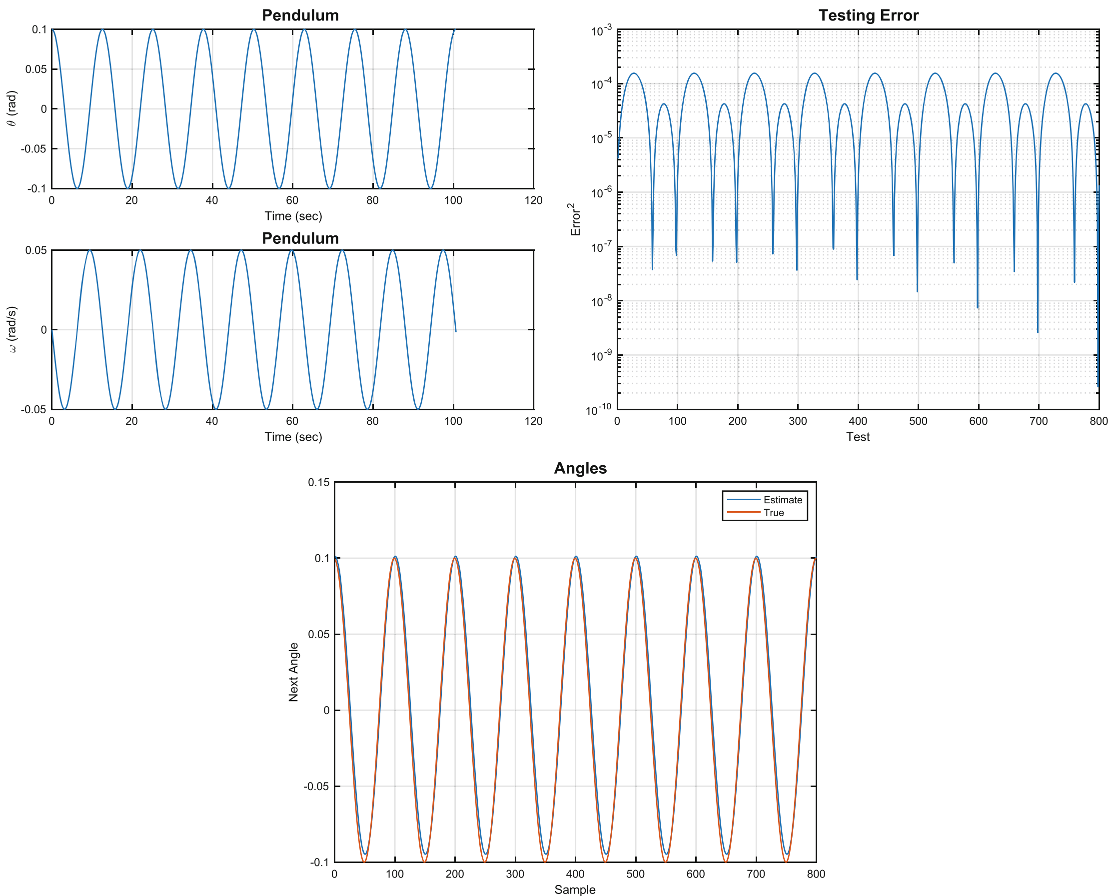


```
% Test the new network
theta    = [x(1,1:end-2);x(1,2:end-1)];
thetaE   = NeuralNetMLFF( theta, network );
eTSq     = (x(1,3:end)-thetaE).^2;
```

The results in Figure 8.9 look good. The neural net estimated angle is quite close to the true angle. Note, however, that we ran exactly the same magnitude pendulum oscillation (`thetaD = theta0`), which is exactly what we trained it to recognize. If we run the test with a different starting point, such as 0.5 radians compared with the 0.1 of the training data, there is more error in the estimated angles, as shown in Figure 8.10.

If we want the neural net to predict angles for other magnitudes, it needs to be trained with a diverse set of data that models all conditions. When we trained the network we let it see the same oscillation magnitude several times. This is not really productive. It might also be necessary to add more nodes to the net or more layers to make a more general purpose estimator.

***Figure 8.10:*** Neural estimated angles for a different magnitude oscillation.

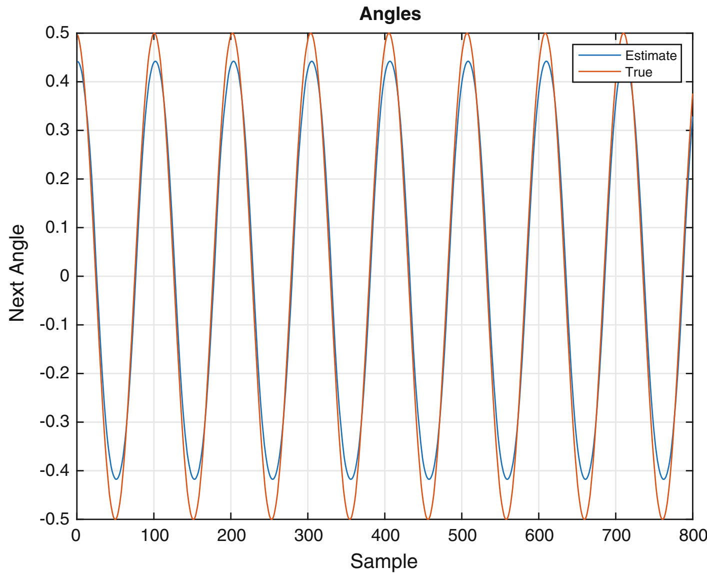


## 8.5 Summary

This chapter has demonstrated neural learning to predict pendulum angles. It introduces the concept of a neuron. It demonstrates a one-neuron network for a pendulum and shows how it compares with a linear estimator. A perceptron example and a multi-layer pendulum angle estimator are also given. Table 8.1 lists the functions and scripts included in the companion code. The last two functions are borrowed from the next chapter, which will cover multi-layer neural nets in more depth.

***Table 8.1:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| NNPendulumDemo | Train a neural net to track a pendulum. |
| OneNeuron | Explore a single neuron. |
| PendulumSim | Simulate a pendulum. |
| RHSPendulum | Right-hand side of a nonlinear pendulum. |
| SunnyDay | Recognize daylight. |
| Chapter 9 Functions | |
| NeuralNetMLFF | Compute the output of a multi-layer, feed-forward neural net. |
| NeuralNetTraining | Training with back propagation. |

## CHAPTER 9

# Classification of Numbers Using Neural Networks

Pattern recognition in images is a classic application of neural nets. This chapter builds upon the previous one by exploring multi-layer networks, which fall into the Machine Learning branch of our Autonomous Learning taxonomy. In this case, we will look at images of computer-generated digits, and the problem of identifying the digits correctly. These images will represent numbers from scanned documents. Attempting to capture the variation in digits with algorithmic rules, considering fonts and other factors, quickly becomes impossibly complex, but with a large number of examples, a neural net can readily perform the task. We allow the weights in the net to perform the job of inferring rules about how each digit may be shaped, rather than codifying them explicitly.

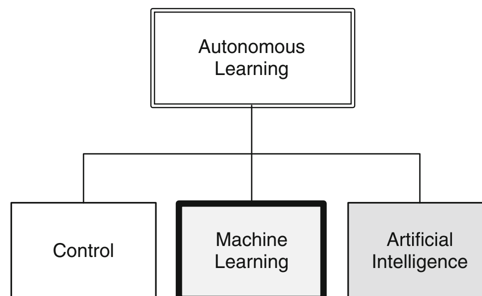


For the purposes of this chapter, we will limit ourselves to images of a single digit. The process of segmenting a series of digits into individual images is one that may be solved by many techniques, not just neural nets.

© Michael Paluszek and Stephanie Thomas 2019
M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_9

## 9.1 Generate Test Images with Defects

### 9.1.1 Problem

The first step in creating our classification system is to generate sample data. In this case, we want to load in images of numbers for 0 to 9 and generate test images with defects. For our purposes, defects will be introduced with simple Poisson or shot noise (a random number with a standard deviation of the square root of the pixel values).

### 9.1.2 Solution

We will generate the images in MATLAB by writing a digit to an axis using `text`, then creating an image using `print`. There is an option to capture the pixel data directly from print without creating an interim file, which we will utilize. We will extract the 16x16 pixel area with our digit, and then apply the noise. We will also allow the font to be an input. See Figure 9.1 for examples.

***Figure 9.1:*** A sample image of the digits 0 and 1 with noise added.

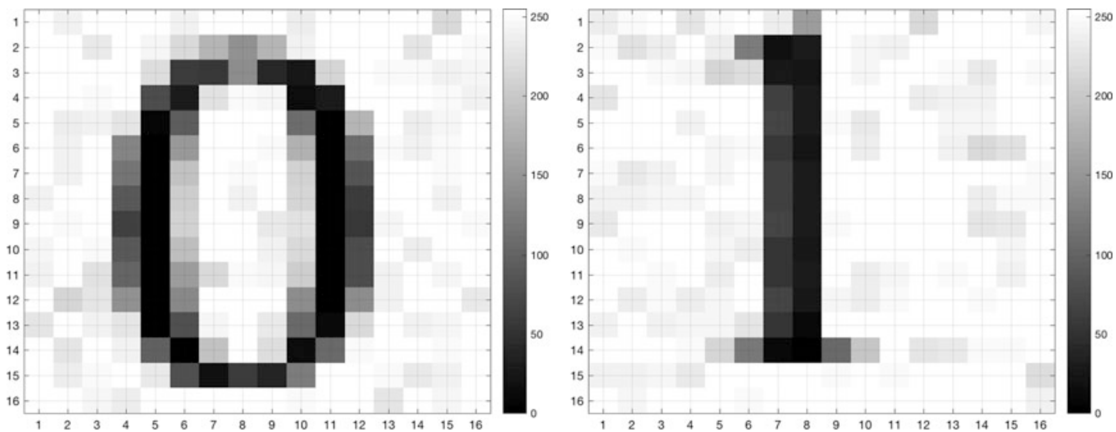


### 9.1.3 How It Works

The code listing for the `CreateDigitImage` function is below. The inputs are the digit and the desired font. It creates a 16x16 pixel image of a single digit. The intermediate figure used to display the digit text is invisible. We will use the `'RGBImage'` option for `print` to get the pixel values without creating an image file. The function has options for a built-in demo that will create pixels for the digit 0 and display the image in a figure if no inputs or outputs are given. The default font if none is given is Courier.

```
function pixels = CreateDigitImage( num, fontname )

if nargin < 1
  num = 0;
  CreateDigitImage( num );
  return;
```

```
end
if nargin < 2
  fontname = 'courier';
end

fonts = listfonts;
avail = strcmpi(fontname,fonts);
if ~any(avail)
  error('MachineLearning:CreateDigitImage',...
    'Sorry,␣the␣font␣''%s''␣is␣not␣available.',fontname);
end

f = figure('Name','Digit','visible','off');
a1 = axes( 'Parent', f, 'box', 'off', 'units', 'pixels', 'position',
    [0 0 16 16] );

% 20 point font digits are 15 pixels tall (on Mac OS)
% text(axes,x,y,string)
text(a1,4,10,num2str(num),'fontsize',19,'fontunits','pixels','unit','
    pixels',...
 'fontname',fontname)

% Obtain image data using print and convert to grayscale
cData = print('-RGBImage','-r0');
iGray = rgb2gray(cData);

% Print image coordinate system starts from upper left of the figure,
     NOT the
% bottom, so our digit is in the LAST 16 rows and the FIRST 16
    columns
pixels = iGray(end-15:end,1:16);

% Apply Poisson (shot) noise; must convert the pixel values to double
     for the
% operation and then convert them back to uint8 for the sum. the
    uint8 type will
% automatically handle overflow above 255 so there is no need to
    apply a limit.
noise = uint8(sqrt(double(pixels)).*randn(16,16));
pixels = pixels - noise;

close(f);

if nargout == 0
  h = figure('name','Digit␣Image');
  imagesc(pixels);
  colormap(h,'gray');
```

```
  grid on
  set(gca,'xtick',1:16)
  set(gca,'ytick',1:16)
  colorbar
end
```

■ **TIP** Note that we check that the font exists using `listfonts` before trying to use it, and throw an error if it's not found.

Now, we can create the training data using images generated with our new function. In the recipes below we will use data for both a single-digit identification and a multiple-digit identification net. We use a `for` loop to create a set of images and save them to a MAT-file using the helper function `SaveTS`. This saves the training sets with their input and output, and indices for training and testing, in a special structure format. Note that we scale the pixel values, which are nominally integers with a value from 0 to 255, to have values between 0 and 1.

Our data generating script `DigitTrainingData` uses a `for` loop to create a set of noisy images for each desired digit (between 0 and 9). It saves the data along with indices for data to use for training. The pixel output of the images is scaled from 0 (black) to 1 (white), so it is suitable for neuron activation in the neural net. It has two flags at the top, one for a one-digit mode and a second to automatically change fonts.

```
%% Generate the training data

% Control switches
oneDigitMode = true;  % the first digit is the desired output
changeFonts = true;   % randomly select a font

% Number of training data sets
digits     = 0:5;
nImagesPer = 20;

% Prepare data
nDigits    = length(digits);
nImages    = nDigits*nImagesPer;
input      = zeros(256,nImages);
output     = zeros(1,nImages);
trainSets = [];
testSets  = [];
if (changeFonts)
  fonts = {'times','helvetica','courier'};
else
  fonts = 'times';
  kFont = 1;
end
```

```
% Loop through digits
kImage = 1;
for j = 1:nDigits
  fprintf('Digit␣%d\n', digits(j));
  for k = 1:nImagesPer
    if (changeFonts)
      % choose a font randomly
      kFont = ceil(rand*3);
    end
    pixels = CreateDigitImage( digits(j), fonts{kFont} );
    % scale the pixels to a range 0 to 1
    pixels = double(pixels);
    pixels = pixels/255;
    input(:,kImage) = pixels(:);
    if (oneDigitMode)
      if (j == 1)
        output(j,kImage) = 1;
      end
    else
      output(j,kImage) = 1;
    end
    kImage = kImage + 1;
  end
  sets = randperm(10);
  trainSets = [trainSets (j-1)*nImages+sets(1:5)]; %#ok<AGROW>
  testSets = [testSets (j-1)*nImages+sets(6:10)]; %#ok<AGROW>
end

% Use 75% of the images for training and save the rest for testing
trainSets = sort(randperm(nImages,floor(0.75*nImages)));
testSets = setdiff(1:nImages,trainSets);

% Save the training set to a MAT-file (dialog window will open)
SaveTS( input, output, trainSets, testSets );
```

The helper function will ask for a filename and save the training set. You can load it at the command line to verify the fields. Here's an example with the training and testing sets truncated:

```
>> trainingData = load('Digit0TrainingTS')
trainingData =
  struct with fields:
    Digit0TrainingTS: [1x1 struct]

>> trainingData.Digit0TrainingTS
ans =
  struct with fields:
        inputs: [256x120 double]
```

```
desOutputs: [1x120 double]
 trainSets: [1 3 4 5 6 8 9 ...  115 117 118 120]
  testSets: [2 7 16 20 28 33 37 ... 112 114 116 119]
```

Note that the output field is a Boolean with a value of 1 when the image is of the desired digit and 0 when it is not. In the single-digit data sets, selected by using the Boolean flag `oneDigitMode`, the output is a single row. In a multi-digit set, it has as many rows as there are digits in the set. The images use a randomly selected font from among Times, Helvetica, and Courier if the `changeFonts` Boolean is true. Table 9.1 shows the three training sets created using this script.

We have created the following sets for use in these recipes:

***Table 9.1:*** Digit Training Sets

| | |
|---|---|
| `'Digit0TrainingTS'` | Single-digit set with 120 images of the digits 0 through 5, all in the same font |
| `'Digit0FontsTS'` | Single-digit set of 0 through 5 with random fonts |
| `'DigitTrainingTS'` | Multi-digit set with 200 images of the digits 0 through 9, same font |

Figure 9.2 shows example images of the digit 2 in the three different fonts, from `Digit0TrainingTS`.

***Figure 9.2:*** Images of the digit 2 in different fonts.

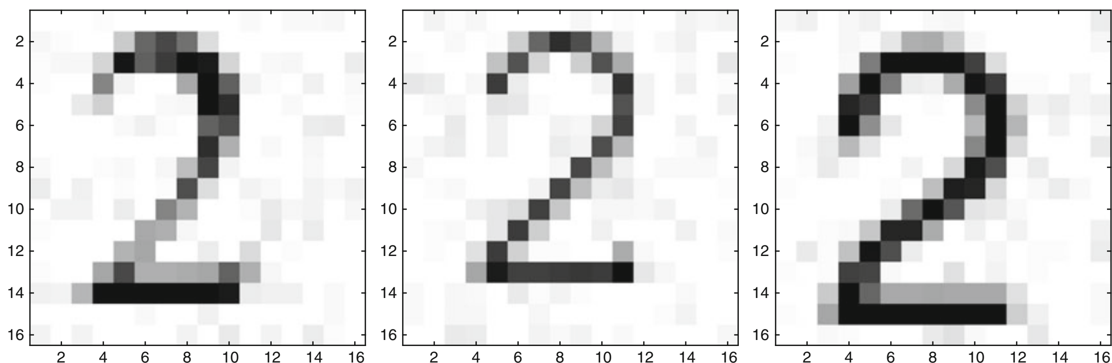


## 9.2 Create the Neural Net Functions

### 9.2.1 Problem

We want to create a neural net tool that can be trained to identify the digits. In this recipe we will discuss the functions underlying the `NeuralNetDeveloper` tool, shown in the next recipe. This interface does not use the latest graphic user interface (GUI)-building features of MATLAB, so we will not get into detail about the GUI code itself although the full GUI is available in the companion code.

### 9.2.2 Solution

The GUI uses a multi-layer feed-forward (MLFF) neural network function to classify digits. In this type of network, each neuron depends only on the inputs it receives from the previous layer. We will discuss the function that implements the neuron.

### 9.2.3 How It Works

The basis of the neural net is the `Neuron` function. Our neuron function provides six different activation types: sign, sigmoid mag, step, logistic, tanh, and sum [22]. This can be seen in Figure 9.3.

The default type of activation function is `tanh`. Two other functions useful in multi-layer networks are exponential (sigmoid logistic function):

$$\frac{1}{1+e^{-x}} \tag{9.1}$$

or sigmoid magnitude:

$$\frac{x}{1+|x|} \tag{9.2}$$

where "sigmoid" refers to a function with an S-shape.

It is a good idea to try different activation functions for any new problem. The activation function is what distinguishes a neural network, and machine learning, from curve fitting. The input `x` would be the sum of all inputs plus a bias.

■ **TIP** The sum activation function is linear and the output is just the sum of the inputs.

The following code shows `Neuron`, which implements a single neuron in the neural net. It has as an input the type, or activation function, and the outputs include the derivative of this function. A default type of `log` is enabled (for the sigmoid logistic function).

```
function [y, dYDX] = Neuron( x, type, t )

% Input processing
if( nargin < 1 )
  x = [];
end
if( nargin < 2 )
  type = [];
end
if( nargin < 3 )
  t = 0;
end
if( isempty(type) )
  type = 'log';
end
if( isempty(x) )
```

***Figure 9.3:*** Available neuron activation functions: sign, sigmoid mag, step, logistic (log), tanh, and sum.

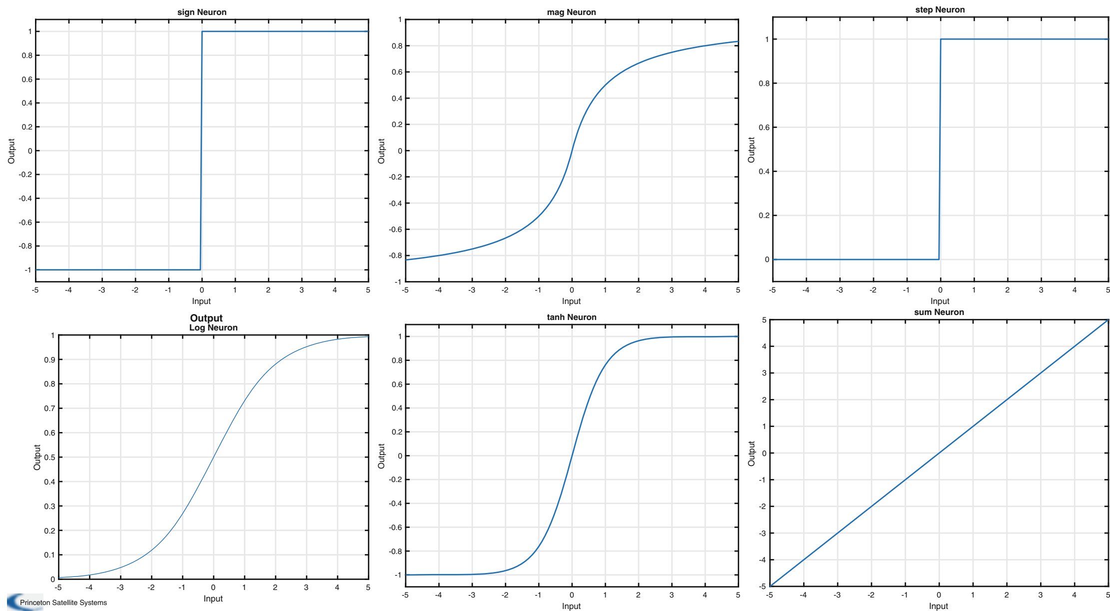


```
  x = sort( [linspace(-5,5) 0 ]);
end

% Compute the function value and the derivative
switch lower( deblank(type) )
  case 'tanh'
    yX   = tanh(x);
    dYDX = sech(x).^2;

  case 'log'
    % sigmoid logistic function
    yX   = 1./(1 + exp(-x));
    dYDX = yX.*(1 - yX);

  case 'mag'
    % sigmoid magnitude function
    d    = 1 + abs(x);
    yX   = x./d;
    dYDX = 1./d.^2;

  case 'sign'
    yX          = ones(size(x));
    yX(x < 0)   = -1;
    dYDX        = zeros(size(yX));
```

```
    dYDX(x == 0) = inf;

  case 'step'
    yX           = ones(size(x));
    yX(x < t)    = 0;
    dYDX         = zeros(size(yX));
    dYDX(x == t) = inf;

  case 'sum'
    yX   = x;
    dYDX = ones(size(yX));

  otherwise
    error([type '␣is␣not␣recognized'])
end

% Output processing
if( nargout == 0 )
  PlotSet( x, yX, 'x␣label', 'Input', 'y␣label', 'Output',...
    'plot␣title', [type '␣Neuron'] );
  PlotSet( x, dYDX, 'x␣label','Input', 'y␣label','dOutput/dX',...
    'plot␣title',['Derivative␣of␣' type '␣Function'] );
else
  y = yX;
end
```

Neurons are combined into the feed-forward neural network using a simple data structure of layers and weights. The input to each neuron is a combination of the signal $y$, the weight $w$, and the bias $w_0$, as in this line:

```
y = Neuron( w*y - w0, type );
```

The output of the network is calculated by the function `NeuralNetMLFF`. This computes the output of a MLFF neural net. Note that this also outputs the derivatives as obtained from the neuron activation functions, for use in training. The function is described below:

```
%% NEURALNETMLFF Computes the output of a multilayer feed-forward
   neural net.
% The input layer is a data structure that contains the network data.
% This data structure must contain the weights and activation
   functions
% for each layer. Calls the Neuron function.
%
% The output layer is the input data structure augmented to include
% the inputs, outputs, and derivatives of each layer for each run.
%% Form
%   [y, dY, layer] = NeuralNetMLFF( x, network )
```

The input and output layers are data structures containing the weights and activation functions for each layer. Our network will use back propagation as a training method [19]. This is

a gradient descent method and it uses the derivatives output by the network directly. Because of this use of derivatives, any threshold functions such as a step function are substituted with a sigmoid function for the training to make it continuous and differentiable. The main parameter is the learning rate $\alpha$, which multiplies the gradient changes applied to the weights in each iteration. This is implemented in `NeuralNetTraining`.

The `NeuralNetTraining` function performs training, that is, it computes the weights in the neurons, using back propagation. If no inputs are given, it will do a demo for the network where node 1 and node 2 use `exp` functions for the activation functions. The function form is given below.

```
%% NEURALNETTRAINING Training using back propagation.
% Computes the weights for a neural net using back propagation. If no
    inputs are
% given it will do a demo for the network where node 1 and node 2 use
    exp
% functions. Calls NeuralNetMLFF which implements the network.
%
%   sin(     x) -- node 1
%                \ /      \
%                 \        ---> Output
%                / \      /
%   sin(0.2*x) -- node 2
%
%% Form
%  [w, e, layer] = NeuralNetTraining( x, y, layer )
```

The back propagation is performed by calling `NeuralNetMLFF` in a loop for the number of runs requested. A wait bar is displayed, since training can take some time. Note that this can handle any number of intermediate layers. The field `alpha` contains the learning rate for the method.

```
% Perform back propagation
h = waitbar(0, 'Neural Net Training in Progress' );
for j = 1:nRuns
  % Work backward from the output layer
  [yN, dYN,layerT] = NeuralNetMLFF( x(:,j), temp );
  e(:,j)            = y(:,j) - yN(:,1); % error

  for k = 1:nLayers
    layer(k,j).w  = temp.layer(k,1).w;
    layer(k,j).w0 = temp.layer(k,1).w0;
    layer(k,j).x  = layerT(k,1).x;
    layer(k,j).y  = layerT(k,1).y;
    layer(k,j).dY = layerT(k,1).dY;
  end

  % Last layer delta is calculated first
  layer(nLayers,j).delta = e(:,j).*dYN(:,1);
```

```
  % Intermediate layers use the subsequent layer's delta
  for k  = (nLayers-1):-1:1
    layer(k,j).delta = layer(k,j).dY.*(temp.layer(k+1,1).w'*layer(k
        +1,j).delta);
  end
  % Now that we have all the deltas, update the weights (w) and
     biases (w0)
  for k = 1:nLayers
    temp.layer(k,1).w  = temp.layer(k,1).w  + layer(k,1).alpha*layer(
        k,j).delta*layer(k,j).x';
    temp.layer(k,1).w0 = temp.layer(k,1).w0 - layer(k,1).alpha*layer(
        k,j).delta;
  end

  waitbar(j/nRuns);
end
w = temp.layer;
close(h);
```

## 9.3 Train a Network with One Output Node

### 9.3.1 Problem

We want to train the neural network to classify numbers. A good first step is identifying a single number. In this case, we will have a single output node, and our training data will include our desired digit, starting with 0, plus a few other digits (1–5).

### 9.3.2 Solution

We can create this neural network with our GUI, shown in Figure 9.4. The network flows from left to right in the graphic. We can try training the net with the output node having different types, such as `sign` and `logistic`. In our case, we start with a `sigmoid` function for the hidden layer and a `step` function for the output node.

The box on the upper left of the GUI lets you set up the network with the number of inputs, in this case one per pixel, the number of outputs, one because we want to identify one digit, and the number of hidden layers. The box to the right lets us design each layer. All neurons in a layer are identical. The box on the far right lets us set the weight for each input to the node and the bias for the node. The path is the path to the training data. The display shows the resulting network. The graphic is useful, but the number of nodes in the hidden layer make it hard to read.

Our GUI has a separate training window, Figure 9.5. It has buttons for loading and saving training sets, training, and testing the trained neural net. It will plot results automatically based on preferences selected. In this case, we have loaded the training set from Recipe 9.1, which uses multiple fonts, `Digit0FontsTS`, which is displayed in the top of the figure window.

***Figure 9.4:*** A neural net with 256 inputs, one per pixel, an intermediate layer with 30 nodes, and one output.

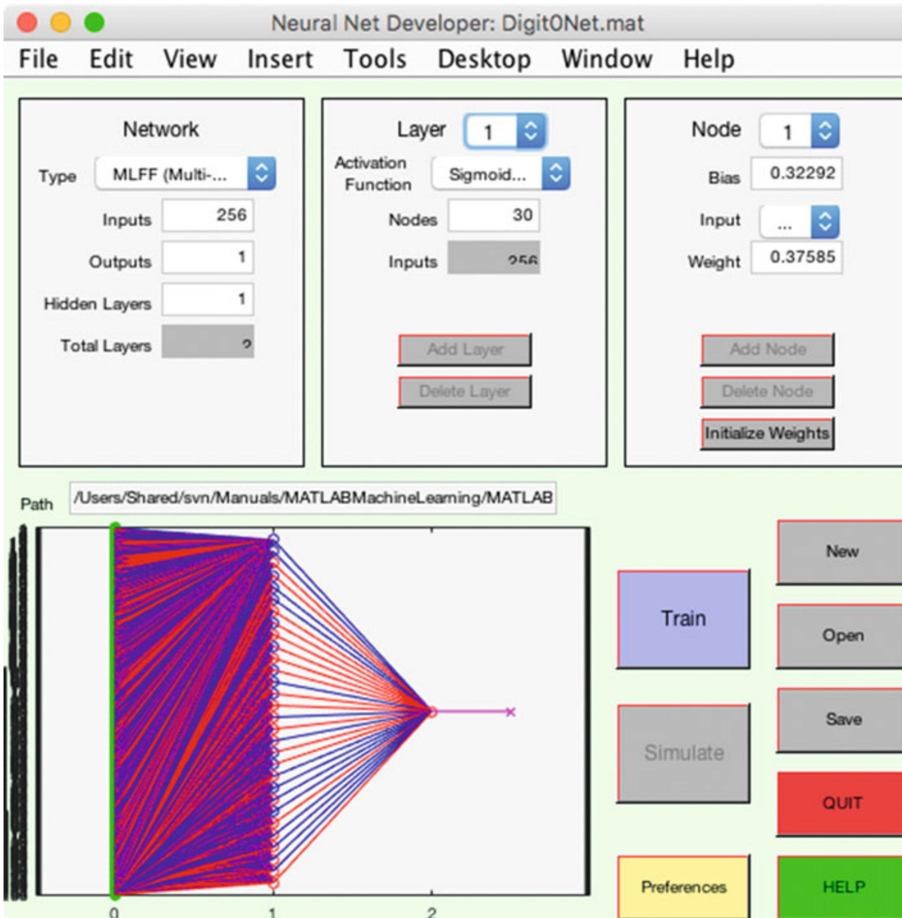


## 9.3.3 How It Works

We build the network using the GUI with 256 inputs, one for each pixel; 30 nodes in one hidden layer; and 1 output node. We load the training data from the first recipe into the Trainer GUI, and must select the number of training runs. 2000 runs should be sufficient if our neuron functions are selected properly. We have an additional parameter to select, the learning rate for the back propagation; it is reasonable to start with a value of 1.0. Note that our training data script assigned 75% of the images for training and reserved the remainder for testing, using `randperm` to extract a random set of images. The training records the weights and biases for each run and generates plots on completion. We can easily plot these for the output node, which has just 30 nodes and one bias. See Figure 9.6.

The training function also outputs the training error as the net evolves and the root mean square error (RMSE), which has dropped off to near 1e-2 by about run 1000.

***Figure 9.5:*** The neural net training GUI opens when the train button is clicked in the developer.

Since we have a large number of input neurons, a line plot is not very useful for visualizing the evolution of the weights for the hidden layer. However, we can view the weights at any given iteration as an image. Figure 9.8 shows the weights for the network with 30 nodes after training visualized using `imagesc`. We may wonder if we really need all 30 nodes in the hidden layer, or if we could extract the necessary number of features identifying our chosen digit with fewer. In the image on the right, the weights are shown sorted along the dimension of the input pixels for each node; we can clearly see that only a few nodes seem to have much variation from the random values they are initialized with, especially nodes 14, 18, and 21. That is, many of our nodes seem to be having no impact.

Since this visualization seems helpful, we add the code to the training GUI after the generation of the weights line plots. We create two images in one figure, the initial value of the weights on the left and the training values on the right. The HSV colormap looks more striking

***Figure 9.6:*** Layer 2 node weights and biases evolution during training.

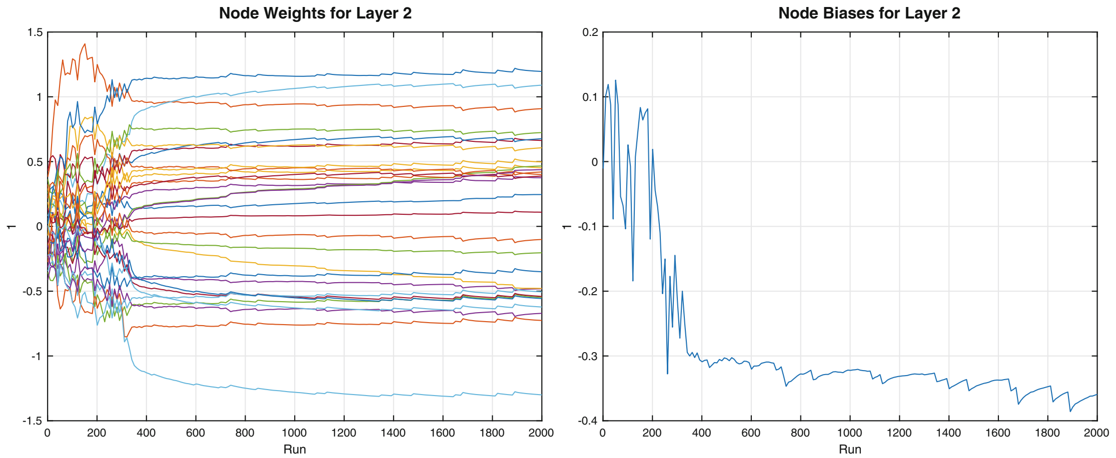


***Figure 9.7:*** Single digit training error and RMSE

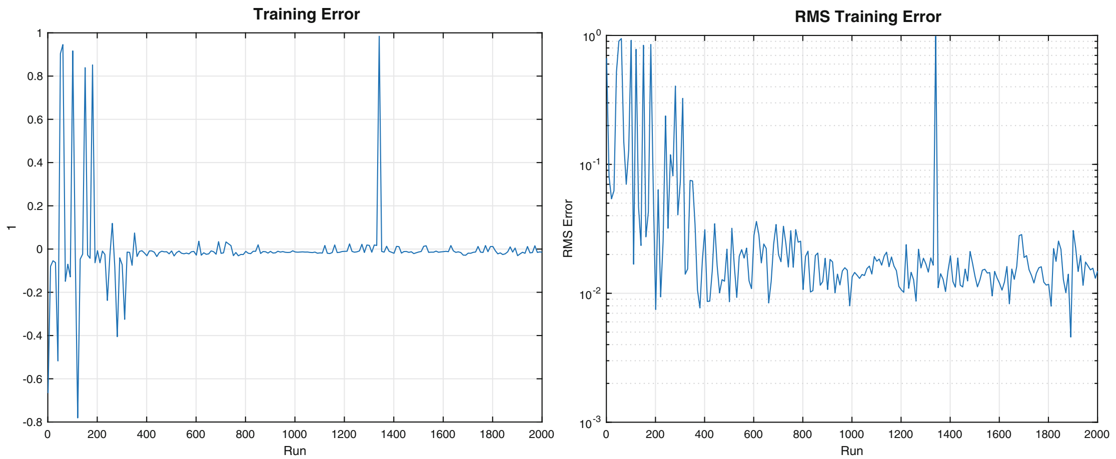


here than the default parula map. The code that generates the images in `NeuralNetTrainer` looks like this:

```
% New figure: weights as image
newH = figure('name',['Node␣Weights␣for␣Layer␣' num2str(j)]);
endWeights = [h.train.network(j,1).w(:);h.train.network(j,end).w(:)];
minW = min(endWeights);
maxW = max(endWeights);
subplot(1,2,1)
imagesc(h.train.network(j,1).w,[minW maxW])
colorbar
ylabel('Output␣Node')
xlabel('Input␣Node')
title('Weights␣Before␣Training')
```

***Figure 9.8:*** Single digit network, 30 node hidden layer weights. The plot on the left shows the weight value. The plot on the right shows the weights sorted by pixel for each node.

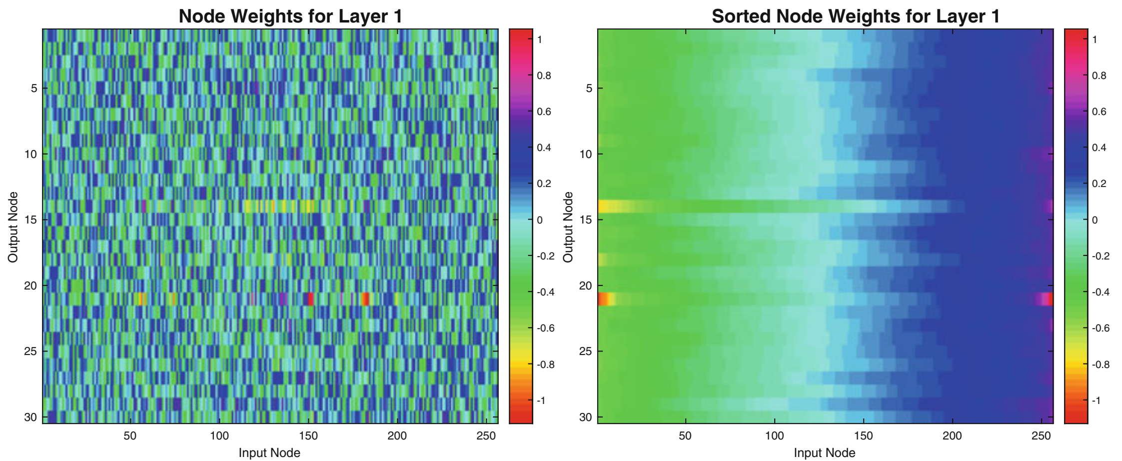


```
subplot(1,2,2)
imagesc(h.train.network(j,end).w, [minW maxW])
colorbar
xlabel('Input␣Node')
title('Weights␣After␣Training')
        colormap hsv
h.resultsFig = [newH; h.resultsFig];
```

Note that we compute the minimum and maximum weight values among both the initial and final iterations, to scale the two color maps the same. Now, since many of our 30 initial nodes seemed unneeded, we reduce the number of nodes in that layer to 10, reinitialize the weights (randomly), and train again. Now we get our new figure with the weights displayed as an image before and after the training, Figure 9.9.

Now we can see more patches of colors that have diverged from the initial random weights in the images for the 256 pixels weights, and we see clear variation in the weights for the second layer as well. The GUI allows you to save the trained net for future use.

***Figure 9.9:*** Single digit network, 10-node hidden layer weights before and after training. The first row shows the data for the first layer, and the second for the second layer, which has just one output.

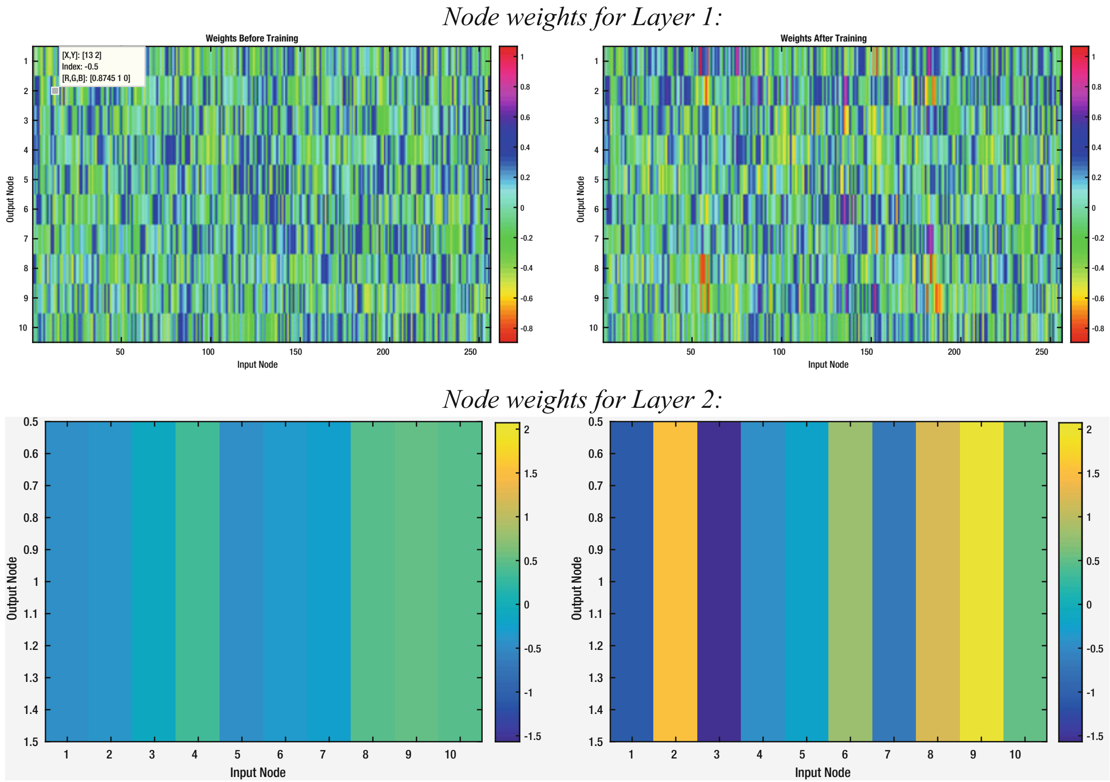


## 9.4 Testing the Neural Network

### 9.4.1 Problem

We want to test the single-digit neural net that we trained in the previous recipe.

### 9.4.2 Solution

We can test the network with inputs that were not used in training. This is explicitly allowed in the GUI, as it has separate indices for the training data and testing data. We selected 75% of our sample images for training and saved the remaining images for testing in our `DigitTrainingData` script from Recipe 9.1.

### 9.4.3 How It Works

In the case of our GUI, simply click the test button to run the neural network with each of the cases selected for testing.

Figure 9.10 shows the results for a network with the output node using the sigmoid magnitude function and another case with the output node using a step function, i.e., the output is limited to 0 or 1. Note that the first 20 images in the data set are the digit 0, with an output value of 1, and the rest are the digits 1 to 5, with an output value of 0. For the step function, the output is 1 for the first 20 sets and zero for all other sets, as desired. The sigmoid is similar, except that instead of being 0 after 20 sets, the output varies between +0.1 and -0.1. Between 20 and 120, it almost averages to zero, the same as the result from the step function. This shows that the activation functions are interpreting the data in a similar fashion.

***Figure 9.10:*** Neural net results with sigmoid (left) and step (right) activation functions.

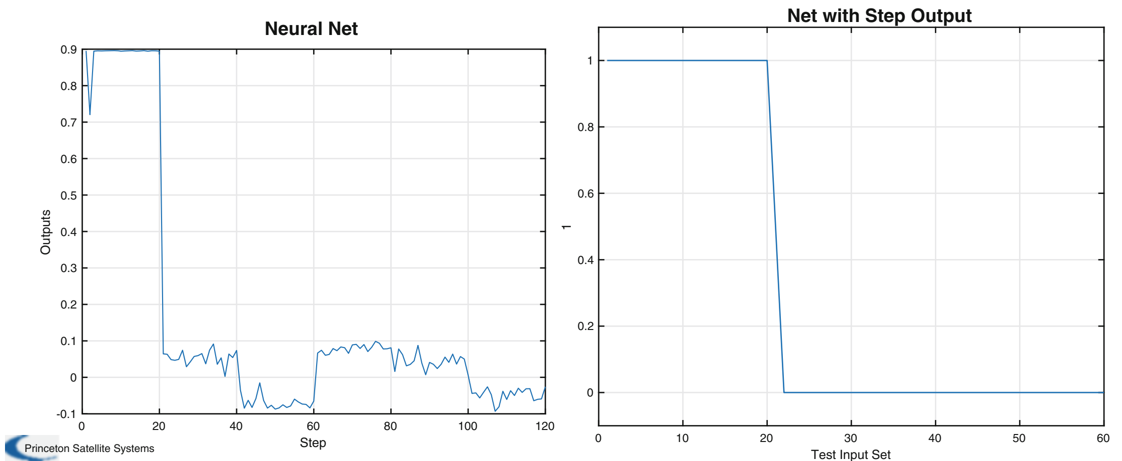


## 9.5 Train a Network with Many Outputs

### 9.5.1 Problem

We want to build a neural net that can detect all ten digits separately.

### 9.5.2 Solution

Add nodes so that the output layer has ten nodes, each of which will be 0 or 1 when the representative digit (0–9) is input. Try the output nodes with different functions, such as logistic and step. Now that we have more digits, we will go back to having 30 nodes in the hidden layer.

### 9.5.3 How It Works

Our training data now consist of all 10 digits, with a binary output of zeros with a 1 in the correct slot. For example, the digit 1 will be represented as

$$[0\ 1\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0]$$

The digit 3 would have a 1 in the fourth element. We follow the same procedure for training. We initialize the net, load the training set into the GUI, and specify the number of training runs for the back propagation.

***Figure 9.11:*** Net with multiple outputs.

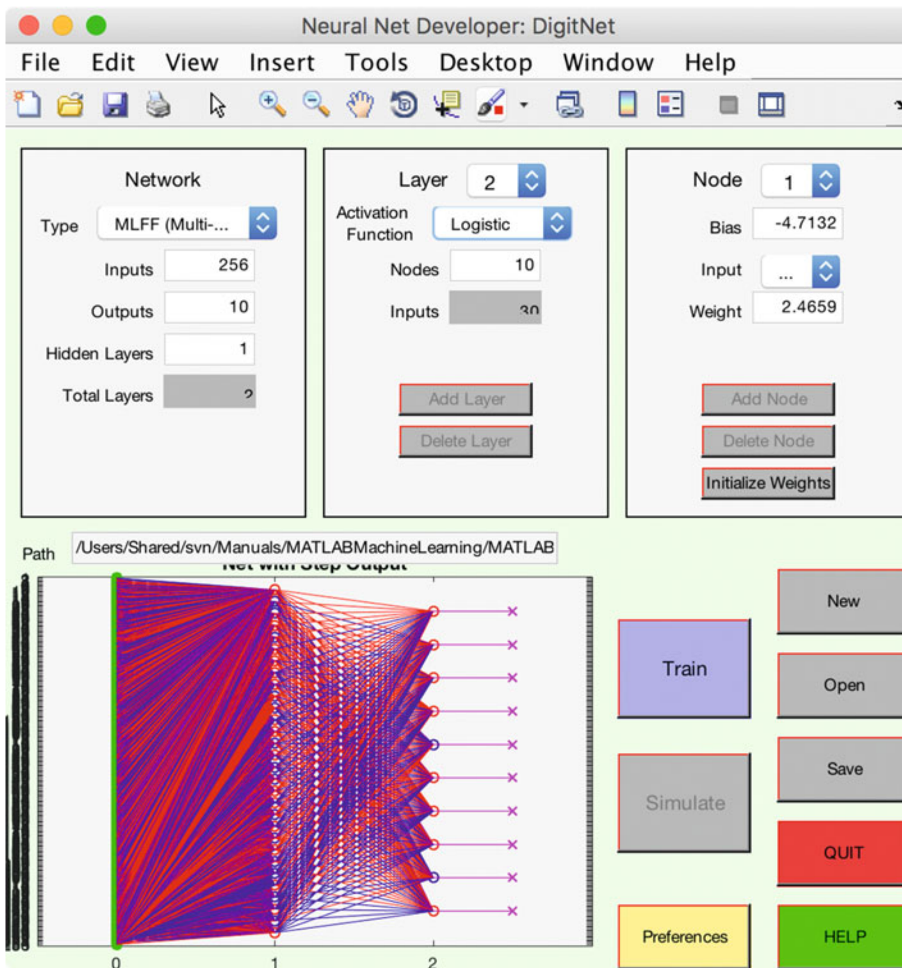


The training data, in Figure 9.12, shows that much of the learning is achieved in the first 3000 runs.

The test data, in Figure 9.13, show that each set of digits (in sets of 20 in this case, for 200 total tests) is correctly identified.

***Figure 9.12:*** Training RMSE for a multiple-digit neural net.

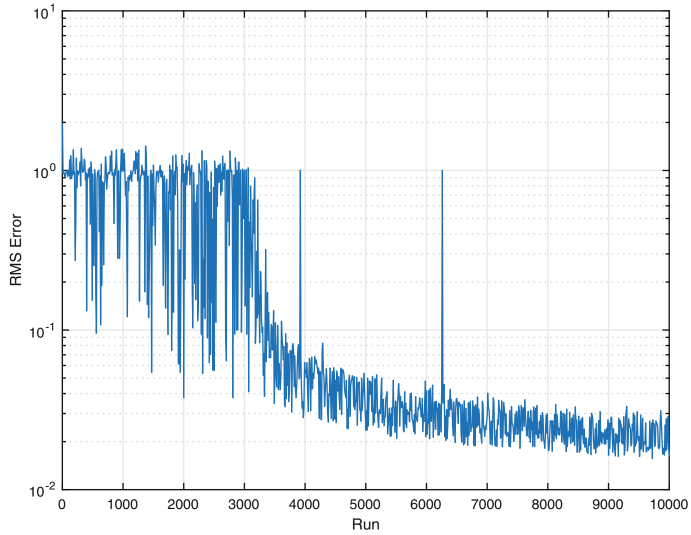


***Figure 9.13:*** Test results for a multiple-digit neural net.

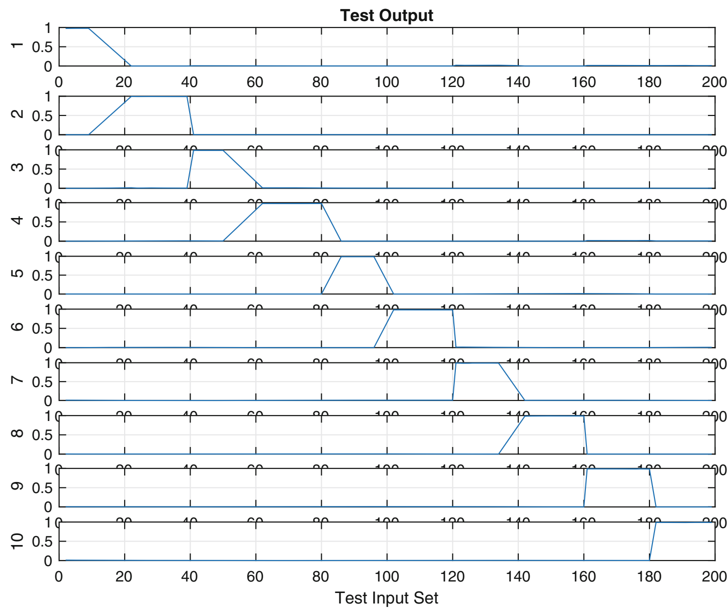

Once you have saved a net that is working well to a MAT-file, you can call it with new data using the function `NeuralNetMLFF`.

```
>> data = load('NeuralNetMat');
>> network = data.DigitsStepNet;
>> y = NeuralNetMLFF( DigitTrainingTS.inputs(:,1), data.DigitsStepNet
    )
y =
     1
     0
     0
     0
     0
     0
     0
     0
     0
     0
```

Again, it is fun to play with visualization of the neural net weights, to gain insight into the problem, and our problem is small enough that we can do so with images. We can view a single set of 256 weights for one hidden neuron as a 16x16 image, and view the whole set with each neuron in its own row as before (Figure 9.14), to see the patterns emerging.

***Figure 9.14:*** Multiple-digit neural net weights.

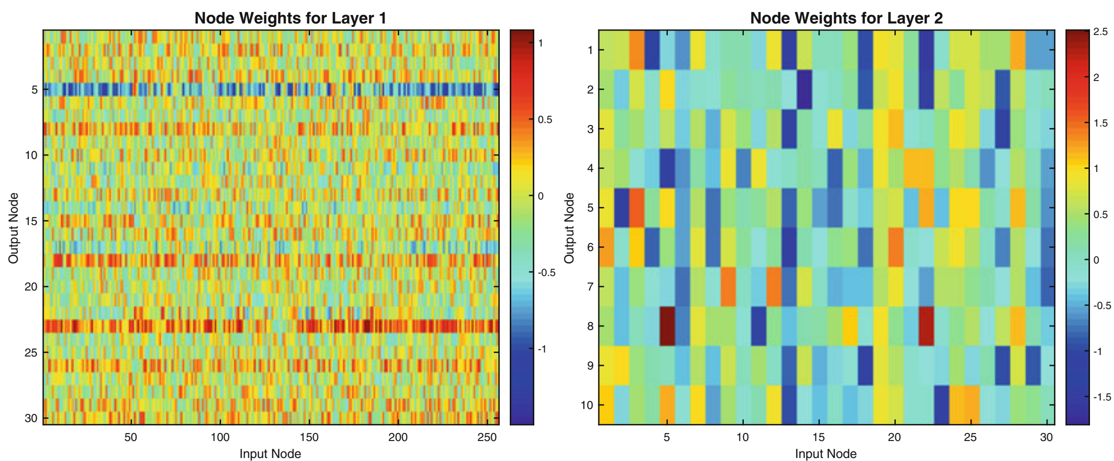


You can see parts of digits as mini-patterns in the individual node weights. Simply use `imagesc` with `reshape` like this:

```
>> figure;
>> imagesc(reshape(net.DigitsStepNet.layer(1).w(23,:),16,16));
>> title('Weights␣to␣Hidden␣Node␣23')
```

and see images as in Figure 9.15. These three nodes (chosen at random) show a 1, 2, and 3. We would expect the 30 nodes to each have "noisy" replicas of the digits.

***Figure 9.15:*** Multiple-digit neural net weights.

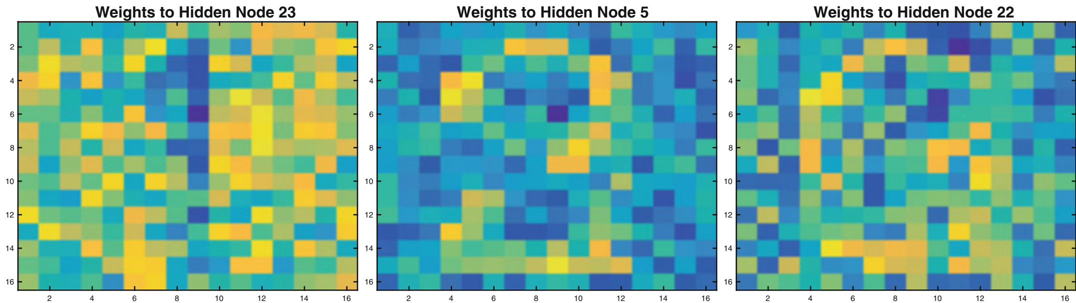


## 9.6 Summary

This chapter has demonstrated neural learning to classify digits. An interesting extension to our tool would be the use of image data stores, rather than a matrix representation of the input data. Table 9.2 lists the functions and scripts included in the companion code.

***Table 9.2:*** Chapter Code Listing

| File | Description |
|---|---|
| DigitTrainingData | Create a training set of digit images. |
| CreateDigitImage | Create a noisy image of a single digit. |
| Neuron | Model an individual neuron with multiple activation functions. |
| NeuralNetMLFF | Compute the output of a MLFF neural net. |
| NeuralNetTraining | Training with back propagation. |
| DrawNeuralNet | Display a neural net with multiple layers. |
| SaveTS | Save a training set MAT-file with index data. |

CHAPTER 10

# Pattern Recognition with Deep Learning

Neural nets fall into the Learning category of our taxonomy. In this chapter, we will expand our neural net toolbox with convolution and pooling layers. A general neural net is shown in Figure 10.1. This is a "deep learning" neural net because it has multiple internal layers. Each layer may have a distinct function and form. In the previous chapter, our multi-layer network had multiple layers, but they were all functionally similar and fully connected.

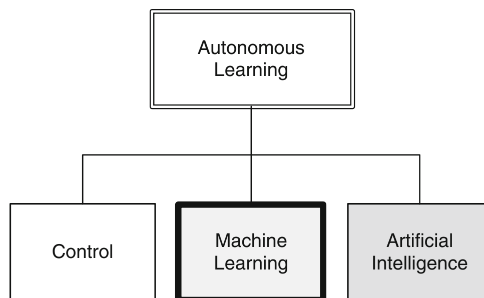


***Figure 10.1:*** Deep learning neural net.

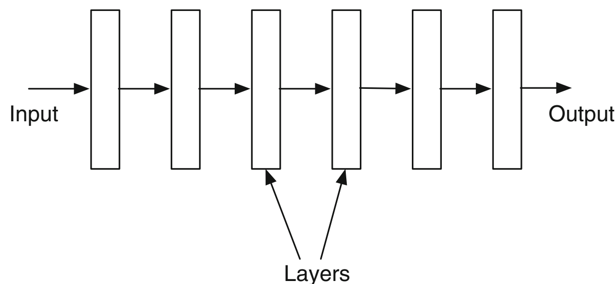


© Michael Paluszek and Stephanie Thomas 2019
M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_10

A convolutional neural network is a type of deep learning network that is a pipeline with multiple stages [18]. There are three types of layers:

- Convolutional layers (hence the name) – convolves a feature with the input matrix so that the output emphasizes that feature. This finds patterns.
- Pooling layers – these reduce the number of inputs to be processed in layers further down the chain.
- Fully connected layers

A convolutional neural net is shown in Figure 10.2. This is also a “deep learning” neural net because it has multiple internal layers, but now the layers are of the three types described above.

***Figure 10.2:*** Deep learning convolutional neural net [13].

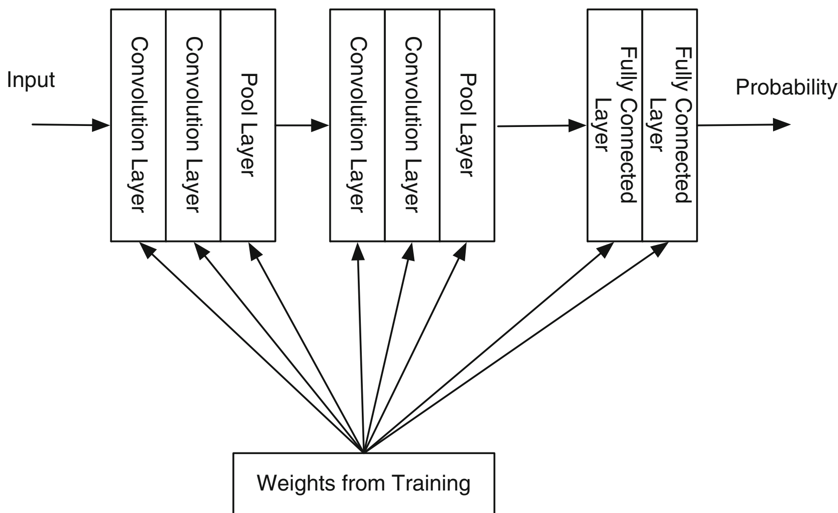


We can have as many layers as we want. The following recipes will detail each step in the chain. We will start by showing how to gather image data online. We won’t actually use online data, but the process may be useful for your work.

We will then describe the convolution process. The convolution process helps to accent features in an image. For example, if a circle is a key feature, convolving a circle with an input image will emphasize circles.

The next recipe will implement pooling. This is a way of condensing the data. For example, if you have an image of a face, you may not need every pixel. You need to find the major features, mouth and eyes, for example, but may not need details of the person’s iris. This is the reverse of what people do with sketching. A good artist can use a few strokes to clearly represent a face. She then fills in detail in successive passes over the drawing. Pooling, at the risk of losing information, reduces the number of pixels to be processed.

We will then demonstrate the full network using random weights. Finally, we will train the network using a subset of our data and test it on the remaining data, as before.

For this chapter, we are going to use pictures of cats. Our network will produce a probability that a given image is a picture of a cat. We will train networks using cat images and also reuse some of our digit images from the previous chapter.

## 10.1 Obtain Data Online for Training a Neural Net

### 10.1.1 Problem

We want to find photographs online for training a cat recognition neural net.

### 10.1.2 Solution

Use the online database ImageNet to search for images of cats.

### 10.1.3 How It Works

ImageNet, http://www.image-net.org, is an image database organized according to the WordNet hierarchy. Each meaningful concept in WordNet is called a "synonym set." There are more than 100,000 sets and 14 million images in ImageNet. For example, type in "Siamese cat." Click on the link. You will see 445 images. You'll notice that there are a wide variety of shots from many angles and a wide range of distances.

```
Synset: Siamese cat, Siamese
Definition: a slender short-haired blue-eyed breed of cat having a
   pale coat with dark ears, paws, face, and tail tip.
Popularity percentile: 57%
Depth in WordNet: 8
```

This is a great resource! However, we are going to instead use pictures of our own cats for our test to avoid copyright issues. The database of photos on ImageNet may prove to be an excellent resource for you to use in training your own neural nets. However, you should review the ImageNet license agreement to determine whether your application can use these images without restrictions.

## 10.2 Generating Training Images of Cats

### 10.2.1 Problem

We want grayscale photographs for training a cat recognition neural net.

### 10.2.2 Solution

Take photographs using a digital camera. Crop them to a standard size manually, then process them using native MATLAB functions to create grayscale images.

### 10.2.3 How It Works

We first take pictures of several cats. We'll use them to train the net. The photos are taken using an iPhone 6. We limit the photos to facial shots of the cats. We then frame the shots so that they are reasonably consistent in size and minimize the background. We then convert them to grayscale.

We use the function `ImageArray` to read in the images. It takes a path to a folder containing the images to be processed. A lot of the code has nothing to do with image processing, just with dealing with unix files in the folder that are not images. `ScaleImage` is in the file reading loop to scale them. We flip them upside down so that they are the right side up from our viewpoint. We then average the color values to make grayscale. This reduces an n by n by 3 array to n by n. The rest of the code displays the images packed into a frame. Finally, we scale all the pixel values down by 256 so that each value is from 0 to 1. The body of `ImageArray` is shown in the listing below.

```
%% IMAGEARRAY Read an array of images from a directory
function [s, sName] = ImageArray( folderPath, scale )

c = cd;
cd(folderPath)

d = dir;
n = length(d);
j = 0;

s     = cell(n-2,1);
sName = cell(1,length(n));
for k = 1:n
  name = d(k).name;
  if( ~strcmp(name,'.') && ~strcmp(name,'..') )
    j        = j + 1;
    sName{j} = name;
    t        = ScaleImage(flipud(imread(name)),scale);
    s{j}     = (t(:,:,1)+ t(:,:,2) + t(:,:,3))/3;
  end
end

del   = size(s{1},1);
lX    = 3*del;

% Draw the images
NewFigure(folderPath);
colormap(gray);
n = length(s);
x = 0;
y = 0;
for k = 1:n
```

```
  image('xdata',[x;x+del],'ydata',[y;y+del],'cdata', s{k} );
  hold on
  x = x + del;
  if ( x == lX )
    x = 0;
    y = y + del;
  end
end
axis off
axis image

for k = 1:length(s)
  s{k} = double(s{k})/256;
end

cd(c)
```

The function has a built-in demo with our local folder of cat images. The images are scaled down by a factor of $2^4$, or 16, so that they are displayed as 64x64 pixel images.

```
%%% ImageArray>Demo
% Generate an array of cat images

c0 = cd;
p = mfilename('fullpath');
cd(fileparts(p));
ImageArray( fullfile('..','Cats'), 4 );
cd(c0);
```

The full set of images in the `Cats` folder, as loaded and scaled in the demo, is shown in Figure 10.3.

***Figure 10.3:*** 64x64 pixel grayscale cat images.

`ImageArray` averages the three colors to convert the color images to grayscale. It flips them upside down, since the image coordinates are opposite to that of MATLAB. We used the GraphicConverter™ application to crop the images around the cat face and make them all 1024x1024 pixels. One of the challenges of image matching is to do this process automatically. Also, typically training uses thousands of images. We will be using just a few to see if our neural net can determine if the test image is a cat, or even one we have used in training! `ImageArray` scales the image using the function `ScaleImage`, shown below.

```
%% SCALEIMAGE Scale an image by powers of 2.
function s2 = ScaleImage( s1, q )

% Demo
if( nargin < 1 )
  Demo
  return
end

n = 2^q;

[mR,~,mD] = size(s1);

m = mR/n;

s2 = zeros(m,m,mD,'uint8');

for i = 1:mD
  for j = 1:m
    r = (j-1)*n+1:j*n;
    for k = 1:m
      c         = (k-1)*n+1:k*n;
      s2(j,k,i) = mean(mean(s1(r,c,i)));
    end
  end
end
```

Notice that it creates the new image array as `uint8`. Figure 10.4 shows the results of scaling a full color image.

***Figure 10.4:*** Image scaled from 1024x1024 to 256x256.

## 10.3 Matrix Convolution

### 10.3.1 Problem

We want to implement convolution as a technique to emphasize key features in images, to make learning more effective. This will then be used in the next recipe to create a convolving layer for the neural net.

### 10.3.2 Solution

Implement convolution using MATLAB matrix operations.

### 10.3.3 How It Works

We create an n-by-n mask that we apply to an m-by-m, where m is greater than n. We start in the upper left corner of the matrix, as shown in Figure 10.5. We multiply the mask times the corresponding elements in the input matrix and do a double sum. That is the first element of the convolved output. We then move it column by column until the highest column of the mask is aligned with the highest column of the input matrix. We then return it to the first column and increment the row. We continue until we have traversed the entire input matrix and our mask is aligned with the maximum row and maximum column.

The mask represents a feature. In effect, we are seeing if the feature appears in different areas of the image. We can have multiple masks. There is one bias and one weight for each element of the mask for each feature. In this case, instead of 16 sets of weights and biases, we only have 4. For large images, the savings can be substantial. In this case, the convolution

***Figure 10.5:*** Convolution process showing the mask at the beginning and end of the process.

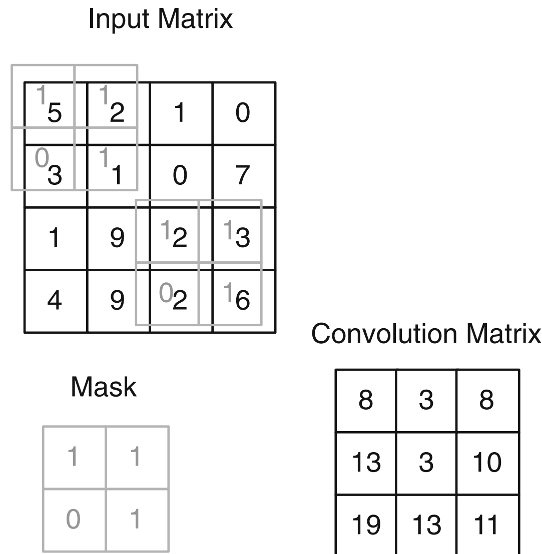


works on the image itself. Convolutions can also be applied to the output of other convolutional layers or pooling layers, as shown in Figure 10.2.

Convolution is implemented in `Convolve.m`. The mask is input `a` and the matrix to be convolved is input `b`.

```
function c = Convolve( a, b )

% Demo
if( nargin < 1 )
  Demo
  return
end

[nA,mA]  = size(a);
[nB,mB]  = size(b);
nC       = nB - nA + 1;
mC       = mB - mA + 1;
c        = zeros(nC,mC);
for j = 1:mC
  jR = j:j+nA-1;
  for k = 1:nC
    kR = k:k+mA-1;
    c(j,k) = sum(sum(a.*b(jR,kR)));
  end
end
```

The demo, which convolves a 3x3 mask with a 6x6 matrix, produces the following 4x4 matrix output.

```
>> Convolve
a =

     1     0     1
     0     1     0
     1     0     1
b =
     1     1     1     0     0     0
     0     1     1     1     0     1
     0     0     1     1     1     0
     0     0     1     1     0     1
     0     1     1     0     0     1
     0     1     1     0     0     1
ans =
     4     3     4     1
     2     4     3     5
     2     3     4     2
     3     3     2     3
```

## 10.4 Convolution Layer

### 10.4.1 Problem

We want to implement a convolution connected layer. This will apply a mask to an input image.

### 10.4.2 Solution

Use code from `Convolve` to implement the layer. It slides the mask across the image and the number of outputs is reduced.

### 10.4.3 How It Works

The "convolution" neural net scans the input with the mask. Each input to the mask passes through an activation function that is identical for a given mask. `ConvolutionLayer` has its own built-in neuron function shown in the listing.

```
%% CONVOLUTIONLAYER Convolution layer for a neural net
function y = ConvolutionLayer( x, d )

% Demo
if( nargin < 1 )
  if( nargout > 0 )
    y = DefaultDataStructure;
  else
    Demo;
  end
```

```
  return
end

a        = d.mask;
aFun     = str2func(d.aFun);
[nA,mA]  = size(a);
[nB,mB]  = size(x);
nC       = nB - nA + 1;
mC       = mB - mA + 1;
y        = zeros(nC,mC);
scale    = nA*mA;
for j = 1:mC
  jR = j:j+nA-1;
  for k = 1:nC
    kR = k:k+mA-1;
    y(j,k)  = sum(sum(a.*Neuron(x(jR,kR),d, aFun)));
  end
end

y = y/scale;

%%% ConvolutionLayer>Neuron
function y = Neuron( x, d, afun )
% Neuron function
y = afun(x.*d.w + d.b);
```

Figure 10.6 shows the inputs and outputs from the demo (not shown in the listing). The `tanh` activation function is used in this demo. The weights and biases are random. The convolution of the mask, which is all ones, is just the sum of all the points that it multiplies. The output is scaled by the number of elements in the mask.

## 10.5 Pooling to Outputs of a Layer

### 10.5.1 Problem

We want to pool the outputs of the convolution layer to reduce the number of points we need to process in further layers. This uses the `Convolve` function created in the previous recipe.

### 10.5.2 Solution

Implement a new function to take the output of the convolution function.

***Figure 10.6:*** Inputs and outputs for the convolution layer.

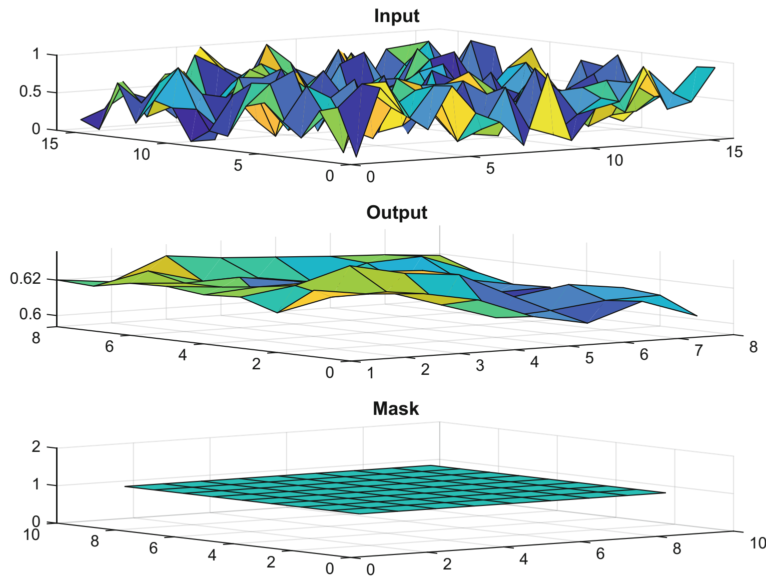


## 10.5.3 How It Works

Pooling layers take a subset of the outputs of the convolutional layers and pass that on. They do not have any weights. Pooling layers can use the maximum value of the pool or take the median or mean value. Our pooling function has all three as options. The pooling function divides the input into n-by-n subregions and returns an n-by-n matrix.

Pooling is implemented in `Pool.m`. Notice we use `str2func` instead of a switch statement. `a` is the matrix to be pooled, `n` is the number of pools, and `type` is the name of the pooling function.

```
function b = Pool( a, n, type )

% Demo
if( nargin < 1 )
  Demo
  return
end

if( nargin <3 )
  type = 'mean';
end

n = n/2;
p = str2func(type);

nA = size(a,1);
```

```
nPP = nA/n;

b = size(n,n);
for j = 1:n
  r = (j-1)*nPP +1:j*nPP;
  for k = 1:n
    c = (k-1)*nPP +1:k*nPP;
    b(j,k) = p(p(a(r,c)));
  end
end
```

These two demos create four pools from a 4x4 matrix. Each number in the output matrix is a pool of one quarter of the input matrix. It uses the default `'mean'` pool method.

```
>> Pool([1:4;3:6;6:9;7:10],4)
ans =
    2.5000    4.5000
    7.0000    9.0000
>> Pool([1:4;3:6;6:9;7:10],4,'max')
ans =
     4     6
     8    10
```

Pool is a neural layer whose activation function is effectively the argument passed to `Pool`.

# 10.6 Fully Connected Layer

## 10.6.1 Problem

We want to implement a fully connected layer.

## 10.6.2 Solution

Use `FullyConnectedNN` to implement the network.

## 10.6.3 How It Works

The “fully connected” neural net layer is the traditional neural net where every input is connected to every output, as shown in Figure 10.7. We implement the fully connected network with `n` inputs and `m` outputs. Each path to an output can have a different weight and bias. `FullyConnectedNN` can handle any number of inputs or outputs. The listing below shows the data structure function as well as the function body.

```
% FullyConnectedNN>Demo

function y = FullyConnectedNN( x, d )

% Demo
if( nargin < 1 )
```

***Figure 10.7:*** Fully connected neural net. This shows only one output.

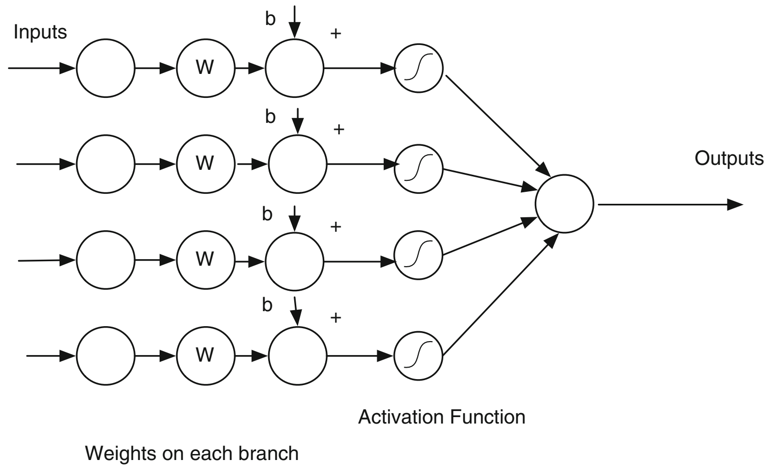


```
  if( nargout > 0 )
    y = DefaultDataStructure;
  else
    Demo;
  end
  return
end

y = zeros(d.m,size(x,2));

aFun = str2func(d.aFun);

n = size(x,1);
for k = 1:d.m
  for j = 1:n
    y(k,:) = y(k,:) + aFun(d.w(j,k)*x(j,:) + d.b(j,k));
  end
end

function d = DefaultDataStructure
%%% FullyConnectedNN>DefaultDataStructure
% Default Data Structure
```

Figure 10.8 shows the outputs from the built-in function demo. The `tanh` activation function is used in this demo. The weights and biases are random. The change in shape from input to output is the result of the activation function.

***Figure 10.8:*** The two outputs from the `FullyConnectedNN` demo function are shown versus the two inputs.

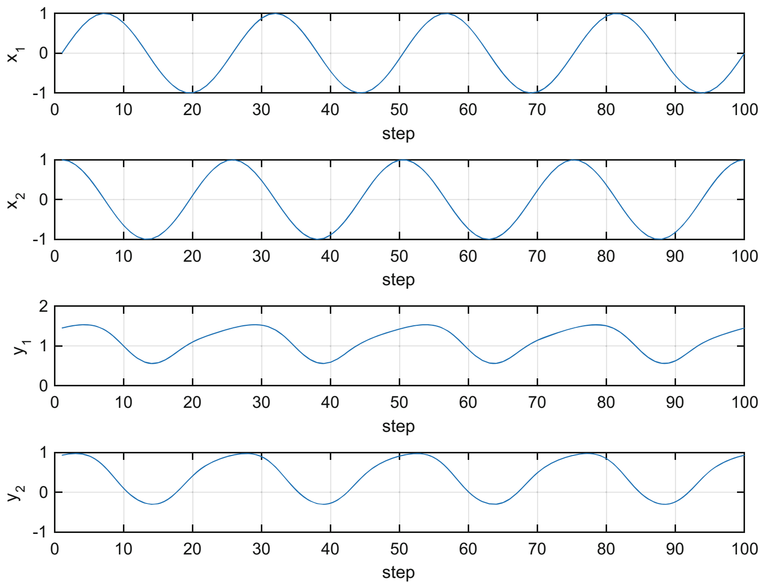


# 10.7 Determining the Probability

## 10.7.1 Problem

We want to calculate a probability that an output is what we expect from neural net outputs.

## 10.7.2 Solution

Implement the Softmax function. Given a set of inputs, it calculates a set of positive values that add up to 1. This will be used for the output nodes of our network.

## 10.7.3 How It Works

The *softmax* function is a generalization of the logistic function. The equation is:

$$p_j = \frac{e^{q_j}}{\sum_{k=1}^{N} e^{q_k}} \tag{10.1}$$

where $q$ is a vector of inputs, $N$ is the number of inputs, and $p$ are the output values that total 1.

The function is implemented in `Softmax.m`.

```
function [p, pMax, kMax] = Softmax( q )

q = reshape(q,[],1);
n = length(q);
p = zeros(1,n);
```

```
den = sum(exp(q));

for k = 1:n
  p(k) = exp(q(k))/den;
end

[pMax,kMax] = max(p);
```

The built-in demo passes in a short list of outputs.

```
function Demo
%% Softmax>Demo
q = [1,2,3,4,1,2,3];
[p, pMax, kMax] = Softmax( q )
sum(p)
```

The results of the demo are:

```
>> Softmax
p =
    0.0236    0.0643    0.1747    0.4748    0.0236    0.0643
       0.1747
pMax =
    0.4748
kMax =
     4
ans =
    1.0000
```

The last number is the sum of p, which should be (and is) 1.

## 10.8 Test the Neural Network

### 10.8.1 Problem

We want to integrate convolution, pooling, a fully connected layer, and Softmax so that our network outputs a probability.

### 10.8.2 Solution

The solution is to write a convolutional neural net. We integrate the convolution, pooling, fully connected net and Softmax functions. We then test it with randomly generated weights.

### 10.8.3 How It Works

Figure 10.9 shows the image processing neural network. It has one convolutional layer, one pooling layer, a fully connected layer and the final layer is the Softmax.

***Figure 10.9:*** Neural net for image processing.

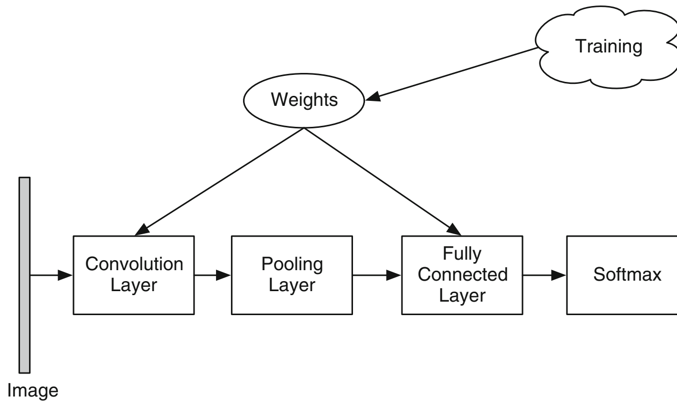


`ConvolutionNN` implements the network. It uses the functions `ConvolutionLayer`, `Pool`, `FullyConnectedNN` and `Softmax` that we have implemented in the prior recipes. The code in `ConvolutionNN`, which implements the network, is shown below, in the subfunction `NeuralNet`. It can generate plots if requested using `mesh`.

```
function r = NeuralNet( d, t, ~ )
%%% ConvolutionalNN>NeuralNet
% Execute the neural net. Plot if there are three inputs.

% Convolve the image
yCL   = ConvolutionLayer( t, d.cL );

% Pool outputs
yPool = Pool( yCL, d.pool.n, d.pool.type );

% Apply a fully connected layer
yFC   = FullyConnectedNN( yPool, d.fCNN );
[~,r] = Softmax( yFC );

% Plot if requested
if( nargin > 2 )
  NewFigure('ConvolutionNN');
  subplot(3,1,1);
  mesh(yCL);
  title('Convolution Layer')
  subplot(3,1,2);
        mesh(yPool);
  title('Pool Layer')
  subplot(3,1,3);
        mesh(yFC);
  title('Fully Connected Layer')
end
```

ConvolutionNN has additional subfunctions for defining the data structure and training and testing the network.

We begin by testing the neural net initialized with random weights, using TestNN. This is a script that loads the cat images using ImageArray, initializes a convolutional network with random weights, and then runs it with a selected test image.

```
>> TestNN
Image IMG_3886.png has a 13.1% chance of being a cat
```

As expected, an untrained neural net does not identify a cat! Figure 10.10 shows the output of the various stages of network processing.

***Figure 10.10:*** Stages in convolutional neural net processing.

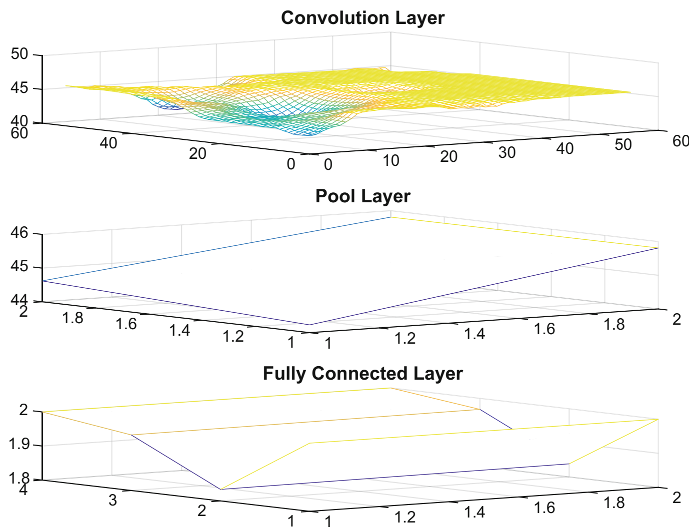


## 10.9 Recognizing a Number

### 10.9.1 Problem

We want to determine if an image is that of the number 3.

### 10.9.2 Solution

We train the neural network with a series of images of the number 3. We then use one picture from the training set and a separate picture and compute the probabilities that they are the number 3.

### 10.9.3 How It Works

We first run the script Digit3TrainingData to generate a training set. This is a simplified version of the training image generation script in Chapter 5, DigitTrainingData. It only produces one digit, in this case the number 3. input has all 256 bits of an image in a single column. output has the output number 1 for all images. We cycle among three fonts 'times','helvetica','courier' for variety. This will make the training more effective when the neural net sees different fonts. Unlike the script in Chapter 4, we store the images as 16x16 pixel images. We also save the three arrays, 'input', 'trainSets', 'testSets' in a .mat file directly using save.

```
%% Generate net training data for the digit 3
digits     = 3;
nImagesPer = 20;

% Prepare data
nDigits    = length(digits);
nImages    = nDigits*nImagesPer;
input      = cell(1,nImages);
output     = zeros(1,nImages);
fonts      = {'times','helvetica','courier'};

% Loop
kImage = 1;
for j = 1:nDigits
  fprintf('Digit %d\n', digits(j));
  for k = 1:nImagesPer
    kFont  = ceil(rand*length(fonts));
    pixels = CreateDigitImage( digits(j), fonts{kFont} );

    % Scale the pixels to a range 0 to 1
    input{kImage} = double(pixels)/255;
    kImage        = kImage + 1;
  end
  sets = randperm(10);
end

% Use 75% of the images for training and save the rest for testing
trainSets = sort(randperm(nImages,floor(0.75*nImages)));
testSets  = setdiff(1:nImages,trainSets);

save('digit3.mat', 'input', 'trainSets', 'testSets');
```

We then run the script TrainNNNumber to see if the input image is the number 3. This script loads in the data from the .mat file into the workspace so that input, trainSets, and

`testSets` are available directly. We get the default data structure from `ConvolutionalNN` and modify the settings for the optimization for `fminsearch`.

```
%% Train a neural net on a single digit
% Trains the net from the images in the loaded mat file.

% Switch to use one image or all for training purposes
useOneImage = false;

% This is needed to make runs consistent
rng('default')

% Load the image data
load('digit3');

% Training
if useOneImage
  % Use only one image for training
        trainSets       = 2;
  testSets  = setdiff(1:length(input),trainSets);
end
fprintf(1,'Training Image(s) [')
fprintf(1,'%1d ',trainSets);
d     = ConvolutionalNN;
d.opt = optimset('TolX',1e-5,'MaxFunEvals',400000,'maxiter',200000);
d     = ConvolutionalNN( 'train', d, input(trainSets) );
fprintf(1,']\nFunction value (should be zero) %12.4f\n',d.fVal);

% Test the net using a test image
for k = 1:length(testSets)
  [d, r] = ConvolutionalNN( 'test', d, input{testSets(k)} );
  fprintf(1,'Test image %d has a %4.1f%% chance of being a 3\n',
     testSets(k),100*r);
end

% Test the net using a test image
[d, r] = ConvolutionalNN( 'test', d, input{trainSets(1)} );

fprintf(1,'Training image %2d has a %4.1f%% chance of being a 3\n',
   trainSets(1),100*r);
```

We set `rng('default')`, since `fminsearch` uses random numbers at times. This makes each run the same. We run the script twice. The first time we use one number for training using the Boolean switch at the top. The second time we use the full training set, like in Chapter 9, setting the Boolean to false. We set `tolX = 1e-5`. This is the tolerance on the weights, which we are trying to solve. Making it smaller doesn't improve anything. If you make it really large, like 1, it will degrade the learning. The number of iterations needs to be greater than 10,000. Again, if you make it too small it won't converge. For one training

image, the script returns that the probability of image 2 or 19 being the number 3 is now 80.3% (presumably numbers with the same font). Other test images range from 35.6% to 47.4%.

```
>> TrainNNNumber
Training Image(s) [2 ]
Function value (should be zero) 0.1969
Test image 1 has a 35.6% chance of being a 3
Test image 6 has a 37.1% chance of being a 3
Test image 11 has a 47.4% chance of being a 3
Test image 18 has a 47.4% chance of being a 3
Test image 19 has a 80.3% chance of being a 3
Training image 2 has a 80.3% chance of being a 3
>> TrainNNNumber
Training Image(s) [2 3 4 5 7 8 9 10 12 13 14 15 16 17 20 ]
Function value (should be zero) 0.5734
Test image 1 has a 42.7% chance of being a 3
Test image 6 has a 42.7% chance of being a 3
Test image 11 has a 42.7% chance of being a 3
Test image 18 has a 42.7% chance of being a 3
Test image 19 has a 42.7% chance of being a 3
Training image 2 has a 42.7% chance of being a 3
```

When we use a lot of images for training representing the various fonts, the probabilities become consistent, though not as high as we would like. Although `fminsearch` does find reasonable weights we could not say that this network is very accurate.

## 10.10 Recognizing an Image

### 10.10.1 Problem

We want to determine if an image is that of a cat.

### 10.10.2 Solution

We train the neural network with a series of cat images. We then use one picture from the training set and a separate picture reserved for testing and compute the probabilities that they are cats.

### 10.10.3 How It Works

We run the script `TrainNN` to see if the input image is a cat. It trains the net from the images in the Cats folder. Many thousands of function evaluations are required for meaningful training, but allowing just a few function evaluations shows that the function is working.

```
%% Train a neural net on the Cats images
p  = mfilename('fullpath');
c0 = cd;
cd(fileparts(p));
```

```
folderPath = fullfile('..','Cats');
[s, name]  = ImageArray( folderPath, 4 );
d          = ConvolutionalNN;

% Use all but the last for training
s = s(1:end-1);

% This may take awhile
% Use at least 10000 iterations to see a higher change of being a cat
   !
disp('Start␣training...')
d.opt.Display = 'iter';
d.opt.MaxFunEvals = 500;
d =     ConvolutionalNN( 'train', d, s );

% Test the net using the last image that was not used in training
[d, r] = ConvolutionalNN( 'test', d, s{end} );

fprintf(1,'Image␣%s␣has␣a␣%4.1f%%␣chance␣of␣being␣a␣cat\n',name{end
   },100*r);

% Test the net using the first image
[d, r] = ConvolutionalNN( 'test', d, s{1} );

fprintf(1,'Image␣%s␣has␣a␣%4.1f%%␣chance␣of␣being␣a␣cat\n',name
   {1},100*r);
```

The script returns that the probability of either image being a cat is now 38.8%. This is an improvement considering we only trained it with one image. It took a couple of hours to process.

```
>> TrainNN

Exiting: Maximum number of function evaluations has been exceeded
         - increase MaxFunEvals option.
         Current function value: 0.612029

Image IMG_3886.png has a 38.8% chance of being a cat
Image IMG_0191.png has a 38.8% chance of being a cat
```

`fminsearch` uses a direct search method (Nelder–Mead simplex), and it is very sensitive to initial conditions.

In fact, using this search method poses a fundamental performance barrier for this neural net training, especially for deep learning, where the combinatorics of different weight combos are so big. Better (and faster) results with a global optimization method are likely.

The training code from `ConvolutionNN` is shown below. It uses MATLAB `fminsearch`. `fminsearch` tweaks the gains and biases until it gets a good fit between all the images input and the training image.

```
function d = Training( d, t )
%%% ConvolutionalNN>Training

d             = Indices( d );
x0            = DToX( d );
[x,d.fVal]      = fminsearch( @RHS, x0, d.opt, d, t );
d             = XToD( x, d );
```

We can improve the results with:

- Adjust `fminsearch` parameters.
- More images.
- More features (masks).
- Change the connections in the fully connected layer.
- Adding the ability of `ConvolutionalNN` to handle RGB images directly, rather than converting them to grayscale.
- Use a different search method such as a genetic algorithm.

## 10.11 Summary

This chapter has demonstrated the steps for implementing a convolutional neural network using MATLAB. Convolutional neural nets were used to process pictures of numbers and cats for learning. When trained, the neural net was asked to identify other pictures to determine if they were pictures of a cat or a number. Table 10.1 lists the functions and scripts included in the companion code.

***Table 10.1:*** Chapter Code Listing

| File | Description |
|---|---|
| Activation | Generate activation functions. |
| ConvolutionalNN | Implement a convolutional neural net. |
| ConvolutionLayer | Implement a convolutional layer. |
| Convolve | Convolve a 2D array using a mask. |
| Digit3TrainingData | Create training data for a single digit. |
| FullyConnectedNN | Implement a fully connected neural network. |
| ImageArray | Read in images in a folder and convert to grayscale. |
| Pool | Pool a 2D array. |
| ScaleImage | Scale and image. |
| Softmax | Implement the Softmax function. |
| TrainNN | Train the convolutional neural net with cat images. |
| TrainNNNumber | Train the convolutional neural net on digit images. |
| TestNN | Test the convolutional neural net on a cat image. |
| TrainingData.mat | Data from TestNN. |

# CHAPTER 11

# Neural Aircraft Control

Longitudinal control is the control of an aircraft that needs to work at the altitude and speed changes. In this chapter, we will implement a neural net to produce the critical parameters for a nonlinear aircraft control system. This is an example of online learning and applies techniques from multiple previous chapters.

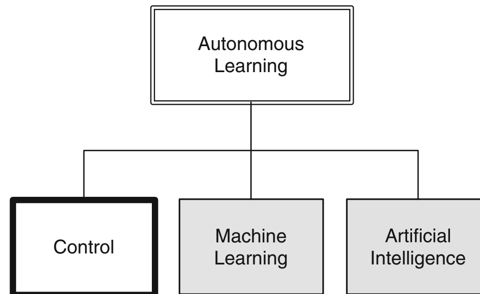


The longitudinal dynamics of an aircraft are also known as the pitch dynamics. The dynamics are entirely in the plane of symmetry of the aircraft. The plane of symmetry is defined as a plane that cuts the aircraft in half vertically. Most airplanes are symmetric about this plane. These dynamics include the forward and vertical motion of the aircraft and the pitching of the aircraft about the axis perpendicular to the plane of symmetry. Figure 11.1 shows an aircraft in flight. $\alpha$ is the angle-of-attack, the angle between the wing and the velocity vector. We assume that the wind direction is opposite that of the velocity vector, that is, the aircraft produces all of its wind. Drag is along the wind direction and lift is perpendicular to drag. The pitch moment is around the center of mass. The model we will derive uses a small set of parameters, yet reproduces the longitudinal dynamics reasonably well. It is also easy for you to modify the model to simulate any aircraft of interest.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_11

***Figure 11.1:*** Diagram of an aircraft in flight showing all the important quantities for longitudinal dynamics simulation.

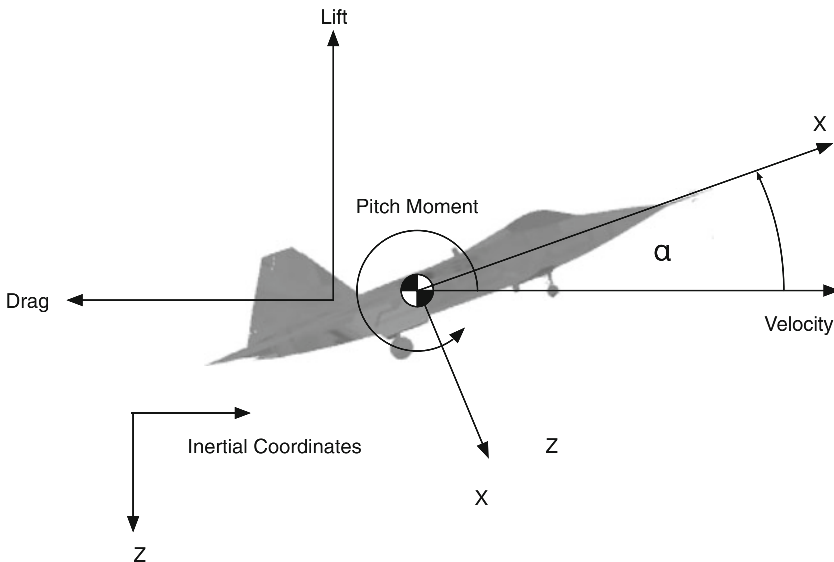


## 11.1 Longitudinal Motion

The next few recipes will involve the longitudinal control of an aircraft with a neural net to provide learning. We will:

1. Model the aircraft dynamics
2. Find an equilibrium solution about which we will control the aircraft
3. Learn how to write a sigma-pi neural net
4. Implement the PID control
5. Implement the neural net
6. Simulate the system

In this recipe, we will model the longitudinal dynamics of an aircraft for use in learning control. We will derive a simple longitudinal dynamics model with a "small" number of parameters. Our control will use nonlinear dynamics inversion with a proportional integral differential (PID) controller to control the pitch dynamics [16, 17]. Learning will be done using a sigma-pi neural network.

We will use the learning approach developed at NASA Dryden Research Center [30]. The baseline controller is a dynamic inversion type controller with a PID control law. A neural net [15] provides learning while the aircraft is operating. The neural network is a sigma-pi type

network meaning that the network sums the products of the inputs with their associated weights. The weights of the neural network are determined by a training algorithm that uses:

1. Commanded aircraft rates from the reference model
2. PID errors
3. Adaptive control rates fed back from the neural network

### 11.1.1 Problem

We want to model the longitudinal dynamics of an aircraft.

### 11.1.2 Solution

The solution is to write the right-hand side function for the aircraft longitudinal dynamics differential equations.

### 11.1.3 How It Works

We summarized the symbols for the dynamical model in Table 11.1 Our aerodynamic model is very simple. The lift and drag are:

$$L = pSC_L \tag{11.1}$$

$$D = pSC_D \tag{11.2}$$

where $S$ is the wetted area, the area that interacts with the airflow and is the area that is counted in computing the aerodynamic forces, and $p$ is the dynamic pressure, the pressure on the aircraft caused by its velocity:

$$p = \frac{1}{2}\rho v^2 \tag{11.3}$$

where $\rho$ is the atmospheric density and $v$ is the magnitude of the velocity. Atmospheric density is a function of altitude. For low-speed flight this is mostly the wings. Most books use $q$ for dynamic pressure. We use $q$ for pitch angular rate (also a convention), so we use $p$ for pressure here to avoid confusion.

The lift coefficient, $C_L$ is:

$$C_L = C_{L_\alpha}\alpha \tag{11.4}$$

and the drag coefficient, $C_D$ is:

$$C_D = C_{D_0} + kC_L^2 \tag{11.5}$$

The drag equation is called the drag polar. Increasing the angle of attack increases the aircraft lift, but also increases the aircraft drag. The coefficient $k$ is:

$$k = \frac{1}{\pi\epsilon_0 AR} \tag{11.6}$$

***Table 11.1:*** Aircraft Dynamics Symbols

| Symbol | Description | Units |
|---|---|---|
| $g$ | Acceleration of gravity at sea-level | 9.806 m/s$^2$ |
| $h$ | Altitude | m |
| $k$ | Coefficient of lift induced drag | |
| $m$ | Mass | kg |
| $p$ | Dynamic pressure | N/m$^2$ |
| $q$ | Pitch angular rate | rad/s |
| $u$ | x-velocity | m/s |
| $w$ | z-velocity | m/s |
| $C_L$ | Lift coefficient | |
| $C_D$ | Drag coefficient | |
| $D$ | Drag | N |
| $I_y$ | Pitch moment of inertia | kg-m$^2$ |
| $L$ | Lift | N |
| $M$ | Pitch moment (torque) | Nm |
| $M_e$ | Pitch moment due to elevator | Nm |
| $r_e$ | Elevator moment arm | m |
| $S$ | Wetted area of wings (the area that contributes to lift and drag) | m$^2$ |
| $S_e$ | Wetted area of elevator | m$^2$ |
| $T$ | Thrust | N |
| $X$ | X force in the aircraft frame | N |
| $Z$ | Z force in the aircraft frame | N |
| $\alpha$ | Angle of attack | rad |
| $\gamma$ | Flight path angle | rad |
| $\rho$ | Air density | kg/m$^3$ |
| $\theta$ | Pitch angle | rad |

where $\epsilon_0$ is the Oswald efficiency factor, which is typically between 0.75 and 0.85. $AR$ is the wing aspect ratio. The aspect ratio is the ratio of the span of the wing to its chord. For complex shapes, it is approximately given by the formula:

$$AR = \frac{b^2}{S} \tag{11.7}$$

where $b$ is the span and $S$ is the wing area. Span is measured from wingtip to wingtip. Gliders have very high aspect ratios and delta-wing aircraft have low aspect ratios.

The aerodynamic coefficients are nondimensional coefficients that when multiplied by the wetted area of the aircraft, and the dynamic pressure, produce the aerodynamic forces.

The dynamical equations, the differential equations of motion, are [5]:

$$m(\dot{u} + qw) = X - mg\sin\theta + T\cos\epsilon \tag{11.8}$$
$$m(\dot{w} - qu) = Z + mg\cos\theta - T\sin\epsilon \tag{11.9}$$
$$I_y\dot{q} = M \tag{11.10}$$
$$\dot{\theta} = q \tag{11.11}$$

$m$ is the mass, $u$ is the x-velocity, $w$ is the z-velocity, $q$ is the pitch angular rate, $\theta$ is the pitch angle, $T$ is the engine thrust, $\epsilon$ is the angle between the thrust vector and the x-axis, $I_y$ is the pitch inertia, $X$ is the x-force, $Z$ is the z-force, and $M$ is the torque about the pitch axis. The coupling between x and z velocities is due to writing the force equations in the rotating frame. The pitch equation is about the center of mass. These are a function of $u$, $w$, $q$, and altitude, $h$, which is found from:

$$\dot{h} = u\sin\theta - w\cos\theta \tag{11.12}$$

The angle of attack, $\alpha$ is the angle between the $u$ and $w$ velocities and is:

$$\tan\alpha = \frac{w}{u} \tag{11.13}$$

The flight path angle $\gamma$ is the angle between the vector velocity direction and the horizontal. It is related to $\theta$ and $\alpha$ by the relationship:

$$\gamma = \theta - \alpha \tag{11.14}$$

This does not appear in the equations, but it is useful to compute when studying aircraft motion. The forces are:

$$X = L\sin\alpha - D\cos\alpha \tag{11.15}$$
$$Z = -L\cos\alpha - D\sin\alpha \tag{11.16}$$

The moment, or torque, is assumed because of the offset of the center-of-pressure and center of mass, which is assumed to be along the $x$-axis:

$$M = (c_p - c)Z \tag{11.17}$$

where $c_p$ is the location of the center of pressure. The moment due to the elevator is:

$$M_e = qr_eS_e\sin(\delta) \tag{11.18}$$

$S_e$ is the wetted area of the elevator and $r_E$ is the distance from the center of mass to the elevator. The dynamical model is in `RHSAircraft`. The atmospheric density model is an exponential model and is included as a subfunction in this function. `RHSAircraft` returns the default data structure if no inputs are given.

```
function [xDot, lift, drag, pD] = RHSAircraft( ~, x, d )

if( nargin < 1 )
  xDot = DataStructure;
  return
end

g     = 9.806; % Acceleration of gravity (m/s^2)

u     = x(1); % Forward velocity
w     = x(2); % Up velocity
q     = x(3); % Pitch angular rate
theta = x(4); % Pitch angle
h     = x(5); % Altitude

rho   = AtmDensity( h ); % Density in kg/m^3

alpha = atan(w/u);
cA    = cos(alpha);
sA    = sin(alpha);

v     = sqrt(u^2 + w^2);
pD    = 0.5*rho*v^2; % Dynamic pressure

cL    = d.cLAlpha*alpha;
cD    = d.cD0 + d.k*cL^2;

drag  = pD*d.s*cD;
lift  = pD*d.s*cL;

x     =  lift*sA - drag*cA;
z     = -lift*cA - drag*sA;
m     =  d.c*z + pD*d.sE*d.rE*sin(d.delta);

sT    = sin(theta);
cT    = cos(theta);

tEng  = d.thrust*d.throttle;
cE    = cos(d.epsilon);
sE    = sin(d.epsilon);

uDot  = (x + tEng*cE)/d.mass - q*w - g*sT + d.externalAccel(1);
wDot  = (z - tEng*sE)/d.mass + q*u + g*cT + d.externalAccel(2);
qDot  = m/d.inertia                       + d.externalAccel(3);
hDot  = u*sT - w*cT;

xDot  = [uDot;wDot;qDot;q;hDot];
```

We will use a model of the F-16 aircraft for our simulation. The F-16 is a single engine supersonic multi-role combat aircraft used by many countries. The F-16 is shown in Figure 11.2.

***Figure 11.2:*** F-16 model.

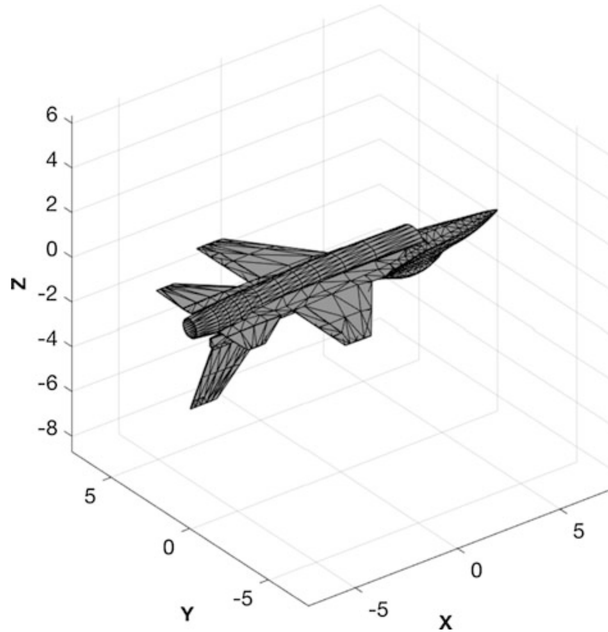


The inertia matrix is found by taking this model, distributing the mass amongst all the vertices, and computing the inertia from the formulas

$$m_k = \frac{m}{N} \tag{11.19}$$

$$c = \sum_k m_k r_k \tag{11.20}$$

$$I = \sum_k m_k (r_k - c)^2 \tag{11.21}$$

where $N$ is the number of nodes and $r_k$ is the vector from the origin (which is arbitrary) to node $k$.

```
inr =

   1.0e+05 *

    0.3672    0.0002   -0.0604
    0.0002    1.4778    0.0000
   -0.0604    0.0000    1.7295
```

***Table 11.2:*** F-16 Data

| **Symbol** | **Field** | **Value** | **Description** | **Units** |
|---|---|---|---|---|
| $C_{L_\alpha}$ | `cLAlpha` | 6.28 | Lift coefficient | |
| $C_{D_0}$ | `cD0` | 0.0175 | Zero lift drag coefficient | |
| $k$ | `k` | 0.1288 | Lift coupling coefficient | |
| $\epsilon$ | `epsilon` | 0 | Thrust angle from the x-axis | rad |
| $T$ | `thrust` | 76.3e3 | Engine thrust | N |
| $S$ | `s` | 27.87 | Wing area | m$^2$ |
| $m$ | `mass` | 12,000 | Aircraft mass | kg |
| $I_y$ | `inertia` | 1.7295e5 | z-axis inertia | kg-m$^2$ |
| $c - c_p$ | `c` | 1 | Offset of center-of-mass from the center-of-pressure | m |
| $S_e$ | `sE` | 3.5 | Elevator area | m$^2$ |
| $r_e$ | (rE) | 4.0 | Elevator moment arm | m |

The F-16 data are given in Table 11.2.

There are many limitations to this model. First of all, the thrust is applied immediately with 100% accuracy. The thrust is also not a function of airspeed or altitude. Real engines take some time to achieve the commanded thrust and the thrust levels change with airspeed and altitude. In the model, the elevator also responds instantaneously. Elevators are driven by motors, usually hydraulic, but sometimes pure electric, and they take time to reach a commanded angle. In our model, the aerodynamics are very simple. In reality, lift and drag are complex functions of airspeed and angle of attack and are usually modeled with large tables of coefficients. We also model the pitching moment by a moment arm. Usually, the torque is modeled by a table. No aerodynamic damping is modeled although this appears in most complete aerodynamic models for aircraft. You can easily add these features by creating functions:

```
C_L = CL(v,h,alpha,delta)
C_D = CD(v,h,alpha,delta)
C_M = CL(v,h,vdot,alpha,delta)
```

## 11.2 Numerically Finding Equilibrium

### 11.2.1 Problem

We want to determine the equilibrium state for the aircraft. This is the orientation at which all forces and torques balance.

### 11.2.2 Solution

The solution is to compute the Jacobian for the dynamics. The Jacobian is a matrix of all first-order partial derivatives of a vector valued function, in this case the dynamics of the aircraft.

### 11.2.3 How It Works

We want to start every simulation from an equilibrium state. This is done using the function `EquilibriumState`. It uses `fminsearch` to minimize:

$$\dot{u}^2 + \dot{w}^2 \tag{11.22}$$

given the flight speed, altitude, and flight path angle. It then computes the elevator angle needed to zero the pitch angular acceleration. It has a built-in demo for equilibrium level flight at 10 km.

```
function [x, thrust, delta, cost] = EquilibriumState( gamma, v, h, d
    )

%% Code
if( nargin < 1 )
  Demo;
  return
end

% [Forward velocity, vertical velocity, pitch rate pitch angle and
    altitude
x               = [v;0;0;0;h];
[~,~,drag]      = RHSAircraft( 0, x, d );
y0              = [0;drag];
cost(1)         = CostFun( y0, d, gamma, v, h );
y               = fminsearch( @CostFun, y0, [], d, gamma, v, h );
w               = y(1);
thrust          = y(2);
u               = sqrt(v^2-w^2);
alpha           = atan(w/u);
theta           = gamma + alpha;
cost(2)         = CostFun( y, d, gamma, v, h );
x               = [u;w;0;theta;h];
d.thrust        = thrust;
d.delta         = 0;
[xDot,~,~,p]    = RHSAircraft( 0, x, d );
```

`CostFun` is the cost functional given below.

```
function cost = CostFun ( y, d, gamma, v, h )
%% EquilibriumState>CostFun
% Cost function for fminsearch. The cost is the square of the
    velocity
% derivatives (the first two terms of xDot from RHSAircraft).
%
% See also RHSAircraft.
```

```
w           = y(1);
d.thrust          = y(2);
d.delta     = 0;
u           = sqrt(v^2-w^2);
alpha       = atan(w/u);
theta       = gamma + alpha;
x           = [u;w;0;theta;h];
xDot        = RHSAircraft( 0, x, d );
cost        = xDot(1:2)'*xDot(1:2);
```

The vector of values is the first input. Our first guess is that thrust equals drag. The vertical velocity and thrust are solved for by `fminsearch`. `fminsearch` searches over thrust and vertical velocity to find an equilibrium state.

The results of the demo are:

```
>> EquilibriumState
Velocity               250.00 m/s
Altitude             10000.00 m
Flight path angle        0.00 deg
Z speed                 13.84 m/s
Thrust               11148.95 N
Angle of attack          3.17 deg
Elevator               -11.22 deg
Initial cost         9.62e+01
Final cost           1.17e-17
```

The initial and final costs show how successful `fminsearch` was in achieving the objective of minimizing the $w$ and $u$ accelerations.

## 11.3 Numerical Simulation of the Aircraft

### 11.3.1 Problem

We want to simulate the aircraft.

### 11.3.2 Solution

The solution is to create a script that calls the right-hand side of the dynamical equations, `RHSAircraft` in a loop, and plot the results.

### 11.3.3 How It Works

The simulation script is shown below. It computes the equilibrium state, then simulates the dynamics in a loop by calling `RungeKutta`. It applies a disturbance to the aircraft. It then uses `PlotSet` to plot the results.

```
%% Initialize
nSim    = 2000;     % Number of time steps
dT      = 0.1;      % Time step (sec)
dRHS    = RHSAircraft;  % Get the default data structure
h       = 10000;
gamma   = 0.0;
v       = 250;
nPulse  = 10;
[x,  dRHS.thrust, dRHS.delta, cost] = EquilibriumState( gamma, v, h,
   dRHS );
fprintf(1,'Finding Equilibrium: Starting Cost %12.4e Final Cost %12.4
   e\n',cost);

accel = [0.0;0.1;0.0];

%% Simulation
xPlot = zeros(length(x)+2,nSim);
for k = 1:nSim
  % Plot storage
  [~,L,D]      = RHSAircraft( 0, x, dRHS );
  xPlot(:,k)   = [x;L;D];
  % Propagate (numerically integrate) the state equations
  if( k > nPulse )
    dRHS.externalAccel = [0;0;0];
  else
    dRHS.externalAccel = accel;
  end
  x            = RungeKutta( @RHSAircraft, 0, x, dT, dRHS );
  if( x(5) <= 0 )
    break;
  end
end
```

The applied external acceleration puts the aircraft into a slight climb with some noticeable oscillations.

```
>> AircraftSimOpenLoop
Velocity             250.00 m/s
Altitude           10000.00 m
Flight path angle      0.57 deg
Z speed               13.83 m/s
Thrust             12321.13 N
Angle of attack        3.17 deg
Elevator              11.22 deg
Initial cost       9.62e+01
Final cost         5.66e-17
Finding Equilibrium: Starting Cost   9.6158e+01 Final Cost   5.6645e
    -17
```

The simulation results are shown in Figure 11.3. The aircraft climbs steadily. Two oscillations are seen. A high frequency one primarily associated with pitch and a low frequency one with the velocity of the aircraft.

***Figure 11.3:*** Open loop response to a pulse for the F-16 in a shallow climb.

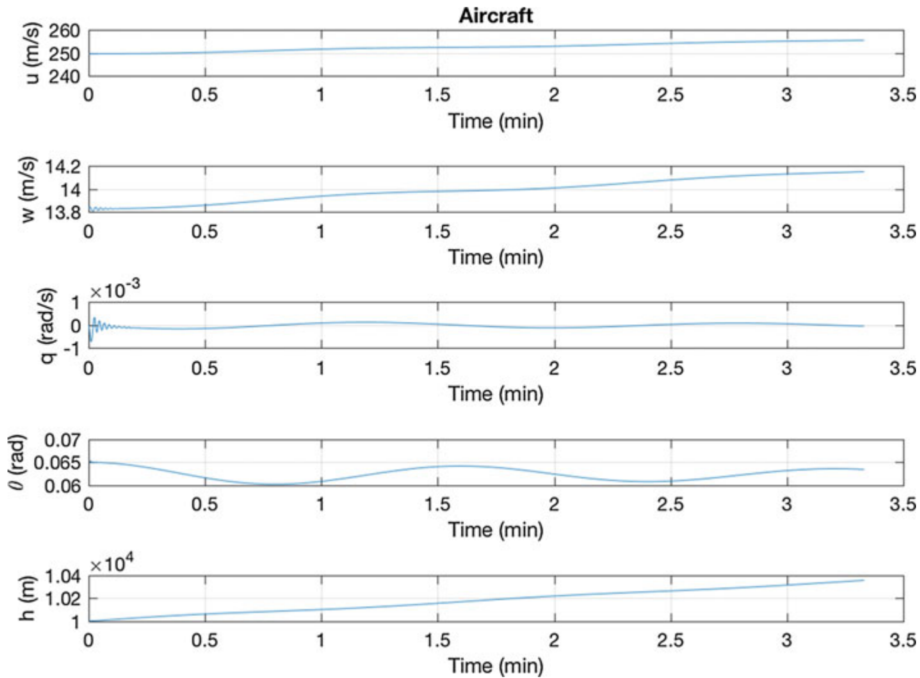


## 11.4 Activation Function

### 11.4.1 Problem

We are going to implement a neural net so that our aircraft control system can learn. We need an activation function to scale and limit measurements.

### 11.4.2 Solution

Use a sigmoid function as our activation function.

### 11.4.3 How It Works

The neural net uses the following sigmoid function

$$g(x) = \frac{1 - e^{-kx}}{1 + e^{-kx}} \tag{11.23}$$

The sigmoid function with $k = 1$ is plotted in the following script.

```
s = (1-exp(-x))./(1+exp(-x));

PlotSet( x, s, 'x label', 'x', 'y label', 's',...
  'plot title', 'Sigmoid', 'figure title', 'Sigmoid' );
```

Results are shown in Figure 11.4.

***Figure 11.4:*** Sigmoid function. At large values of $x$, the sigmoid function returns $\pm 1$.

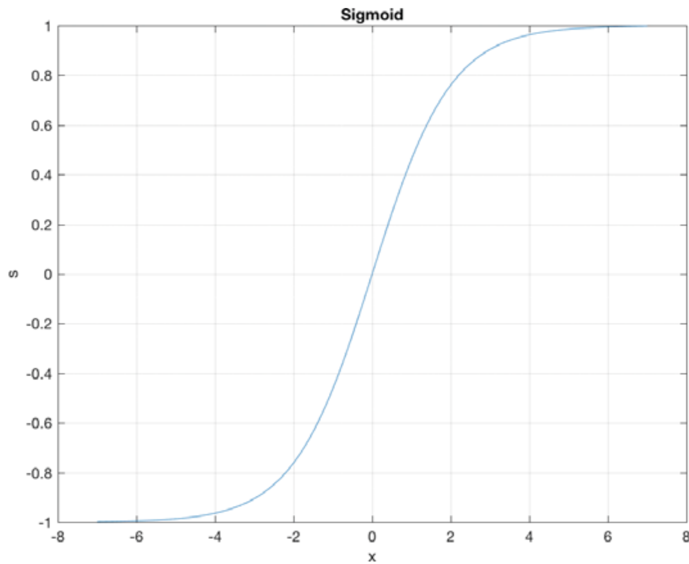


# 11.5 Neural Net for Learning Control

## 11.5.1 Problem

We want to use a neural net to add learning to the aircraft control system.

## 11.5.2 Solution

Use a sigma-pi neural net function. A sigma-pi neural net sums the inputs and products of the inputs to produce a model.

## 11.5.3 How It Works

The adaptive neural network for the pitch axis has seven inputs. The output of the neural network is a pitch angular acceleration that augments the control signal coming from the dynamic inversion controller. The control system is shown in Figure 11.5. The left-most box produces the reference model given the pilot input. The output of the reference model is a vector of the

desired states that are differenced with the true states and fed to the PID controller and the neural network. The output of the PID is differenced with the output of the neural network. This is fed into the model inversion block that drives the aircraft dynamics.

***Figure 11.5:*** Aircraft control system. It combines a PID controller with dynamic inversion to handle nonlinearities. A neural net provides learning.

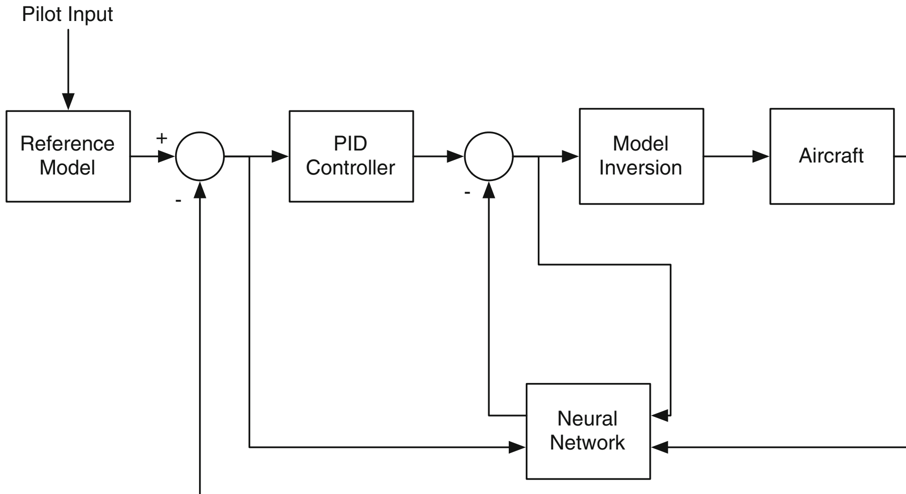


The sigma-pi neural net is shown in Figure 11.6 for a two-input system.

***Figure 11.6:*** Sigma-pi neural net. $\Pi$ stands for product and $\Sigma$ stands for sum.

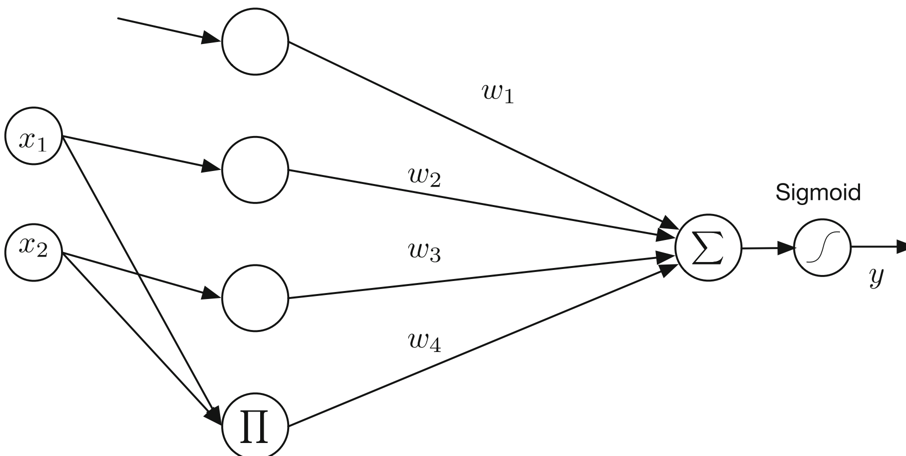


The output is:

$$y = w_1 c + w_2 x_1 + w_3 x_2 + w_4 x_1 x_2 \tag{11.24}$$

The weights are selected to represent a nonlinear function. For example, suppose we want to represent the dynamic pressure:

$$y = \frac{1}{2}\rho v^2 \tag{11.25}$$

We let $x_1 = \rho$ and $x_2 = v^2$. Set $w_4 = \frac{1}{2}$ and all other weights to zero. Suppose we didn't know the constant $\frac{1}{2}$. We would like our neural net to determine the weight through measurements. Learning for a neural net means determining the weights so that our net replicates the function it is modeling. Define the vector $z$, which is the result of the product operations. In our two-input case this would be:

$$z = \begin{bmatrix} c \\ x_1 \\ x_2 \\ x_1 x_2 \end{bmatrix} \tag{11.26}$$

$c$ is a constant. The output is:

$$y = w^T z \tag{11.27}$$

We could assemble multiple inputs and outputs:

$$\begin{bmatrix} y_1 & y_2 & \cdots \end{bmatrix} = w^T \begin{bmatrix} z_1 & z_2 & \cdots \end{bmatrix} \tag{11.28}$$

where $z_k$ is a column array. We can solve for the weights $w$ using least squares given the outputs, $y$, and inputs, $x$. Define the vector of $y$ to be $Y$ and the matrix of $z$ to be $Z$. The solution for $w$ is:

$$Y = Z^T w \tag{11.29}$$

The least squares solution is:

$$w = \left(ZZ^T\right)^{-1} ZY^T \tag{11.30}$$

This gives the best fit to $w$ for the measurements $Y$ and inputs $Z$. Suppose we take another measurement. We would then repeat this with bigger matrices. As a side note, you would really compute this using an inverse. There are better numerical methods for doing least squares. MATLAB has the `pinv` function. For example:

```
>> z = rand(4,4);
>> w = rand(4,1);
>> y = w'*z;
>> wL = inv(z*z')*z*y'
wL =
     0.8308
     0.5853
     0.5497
     0.9172
```

```
>> w
w =
    0.8308
    0.5853
    0.5497
    0.9172

>> pinv(z')*y'
ans =
    0.8308
    0.5853
    0.5497
    0.9172
```

As you can see, they all agree! This is a good way to initially train your neural net. Collect as many measurements as you have values of $z$ and compute the weights. Your net is then ready to go.

The recursive approach is to initialize the recursive trainer with $n$ values of $z$ and $y$.

$$p = \left(ZZ^T\right)^{-1} \tag{11.31}$$

$$w = pZY \tag{11.32}$$

The recursive learning algorithm is:

$$p = p - \frac{pzz^Tp}{1+z^Tpz} \tag{11.33}$$

$$k = pz \tag{11.34}$$

$$w = w + k\left(y - z^Tw\right) \tag{11.35}$$

`RecursiveLearning` demonstrates recursive learning or training. It starts with an initial estimate based on a four-element training set. It then recursively learns based on new data.

```
wN   = w + 0.1*randn(4,1); % True weights are a little different
n    = 300;
zA   = randn(4,n); % Random inputs
y    = wN'*zA; % 100 new measurements

% Batch training
p    = inv(Z*Z'); % Initial value
w    = p*Z*Y; % Initial value

%% Recursive learning
dW = zeros(4,n);
for j = 1:n
  z       = zA(:,j);
  p       = p - p*(z*z')*p/(1+z'*p*z);
  w       = w + p*z*(y(j) - z'*w);
  dW(:,j) = w - wN; % Store for plotting
end
```

```
%% Plot the results
yL = cell(1,4);
for j = 1:4
  yL{j} = sprintf('\\Delta W_%d',j);
end

PlotSet(1:n,dW,'x label','Sample','y label',yL,...
        'plot title','Recursive Training',...
        'figure title','Recursive Training');
```

Figure 11.7 shows the results. After an initial transient, the learning converges. Every time you run this you will get different answers because we initialize with random values.

***Figure 11.7:*** Recursive training or learning. After an initial transient the weights converge quickly.

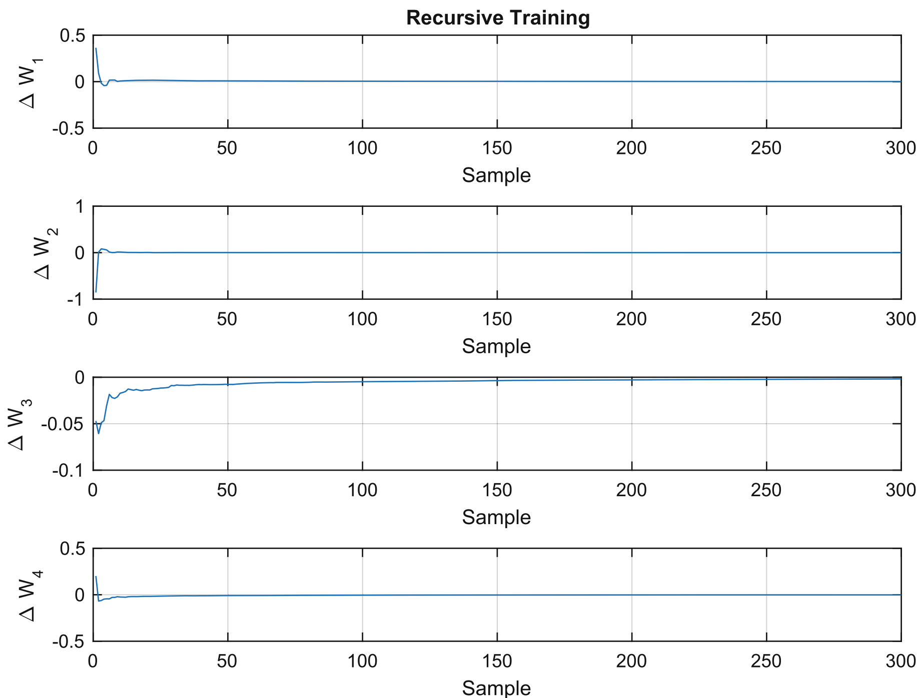


You will notice that the recursive learning algorithm is identical in form to the Kalman Filter given in Section 4.1.3. Our learning algorithm was derived from batch least squares, which is an alternative derivation for the Kalman Filter.

## 11.6 Enumeration of All Sets of Inputs

### 11.6.1 Problem

One issue with a sigma-pi neural network is the number of possible nodes. For design purposes, we need a function to enumerate all possible sets of combinations of inputs. This is to determine the limitation of the complexity of a sigma-pi neural network.

### 11.6.2 Solution

Write a combination function that computes the number of sets.

### 11.6.3 How It Works

In our sigma-pi network, we hand-coded the products of the inputs. For more general code, we want to enumerate all combinations of inputs. If we have $n$ inputs and want to take them $k$ at a time, the number of sets is:

$$\frac{n!}{(n-k)!k!} \tag{11.36}$$

The code to enumerate all sets is in the function `Combinations`.

```
function c = Combinations( r, k )

%% Demo
if( nargin < 1 )
  Combinations(1:4,3)
  return
end

%% Special cases
if( k == 1 )
  c = r';
  return
elseif( k == length(r) )
  c = r;
  return
end

%% Recursion
rJ      = r(2:end);
c   = [];
if( length(rJ) > 1 )
  for j = 2:length(r)-k+1
    rJ            = r(j:end);
    nC            = NumberOfCombinations(length(rJ),k-1);
    cJ            = zeros(nC,k);
    cJ(:,2:end)   = Combinations(rJ,k-1);
    cJ(:,1)       = r(j-1);
```

```
    if( ~isempty(c) )
      c = [c;cJ];
    else
      c = cJ;
    end
  end
else
  c = rJ;
end
c = [c;r(end-k+1:end)];
```

This handles two special cases on input and then calls itself recursively for all other cases. Here are some examples:

```
>> Combinations(1:4,3)
ans =
     1     2     3
     1     2     4
     1     3     4
     2     3     4
 >> Combinations(1:4,2)
ans =
     1     2
     1     3
     1     4
     2     3
     2     4
     3     4
```

You can see that if we have four inputs and want all possible combinations, we end up with 14 in total! This indicates a practical limit to a sigma-pi neural network, as the number of weights will grow fast as the number of inputs increases.

## 11.7 Write a Sigma-Pi Neural Net Function

### 11.7.1 Problem

We need a sigma-pi net function for general problems.

### 11.7.2 Solution

Use a sigma-pi function.

### 11.7.3 How It Works

The following code shows how we implement the sigma-pi neural net. `SigmaPiNeuralNet` has `action` as its first input. You use this to access the functionality of the function. Actions are:

1. "initialize" – initialize the function
2. "set constant" – set the constant term
3. "batch learning" – perform batch learning
4. "recursive learning" – perform recursive learning
5. "output" – generate outputs without training

You usually go in order when running the function. Setting the constant is not needed if the default of one is fine.

The functionality is distributed among sub-functions called from the `switch` statement.

The demo shows an example of using the function to model dynamic pressure. Our inputs are the altitude and the square of the velocity. The neural net will try to fit:

$$y = w_1 c + w_2 h + w_3 v^2 + w_4 h v^2 \tag{11.37}$$

to

$$y = 0.6125 e^{-0.0817 h.^{1.15}} v^2 \tag{11.38}$$

We first get a default data structure. Then we initialize the filter with an empty `x`. We then get the initial weights by using batch learning. The number of columns of x should be at least twice the number of inputs. This gives a starting $p$ matrix and initial estimate of weights. We then perform recursive learning. It is important that the field `kSigmoid` is small enough so that valid inputs are in the linear region of the sigmoid function. Note that this can be an array so that you can use different scalings on different inputs.

```
function Demo
% Demonstrate a sigma-pi neural net for dynamic pressure
x       = zeros(2,1);

d       = SigmaPiNeuralNet;
[~, d]  = SigmaPiNeuralNet( 'initialize', x, d );

h       = linspace(10,10000);
v       = linspace(10,400);
v2      = v.^2;
q       = 0.5*AtmDensity(h).*v2;

n       = 5;
x       = [h(1:n);v2(1:n)];
```

```
d.y      = q(1:n)';
[y, d]   = SigmaPiNeuralNet( 'batch learning', x, d );

fprintf(1,'Batch Results\n#         Truth   Neural Net\n');
for k = 1:length(y)
  fprintf(1,'%d: %12.2f %12.2f\n',k,q(k),y(k));
end

n = length(h);
y = zeros(1,n);
x = [h;v2];
for k = 1:n
  d.y = q(k);
  [y(k), d]  = SigmaPiNeuralNet( 'recursive learning', x(:,k), d );
end
```

The batch results are as follows for five examples of dynamic pressures at low altitude. As you can see the truth model and neural net outputs are quite close:

```
>> SigmaPiNeuralNet
Batch Results
#          Truth   Neural Net
1:         61.22         61.17
2:        118.24        118.42
3:        193.12        192.88
4:        285.38        285.52
5:        394.51        394.48
```

The recursive learning results are shown in Figure 11.8. The results are pretty good over a wide range of altitudes. You could then just use the "update" action during aircraft operation.

## 11.8 Implement PID Control

### 11.8.1 Problem

We want a PID controller to control the aircraft.

### 11.8.2 Solution

Write a function to implement PID control. The inputs will be the pitch angle error.

***Figure 11.8:*** Recursive training for the dynamic pressure example.

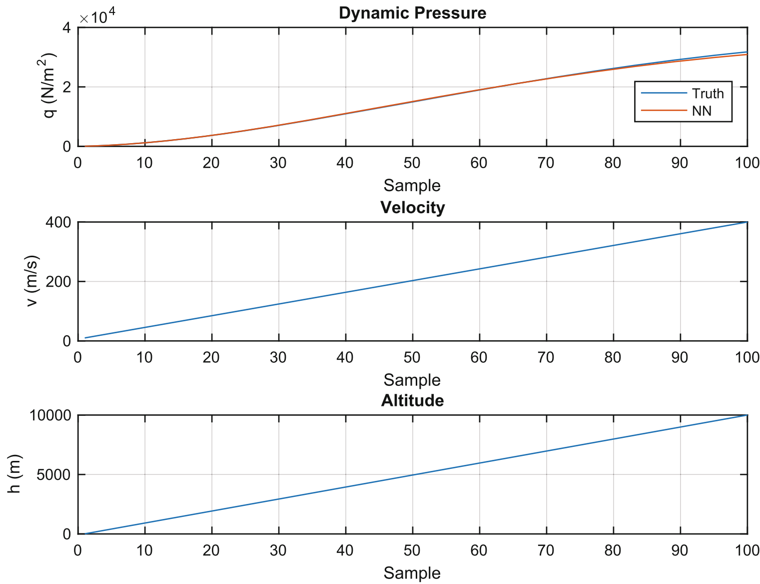


## 11.8.3 How It Works

Assume we have a double integrator driven by a constant input:

$$\ddot{x} = u \tag{11.39}$$

where $u = u_d + u_c$.
The result is:

$$x = \frac{1}{2}ut^2 + x(0) + \dot{x}(0)t \tag{11.40}$$

The simplest control is to add a feedback controller

$$u_c = -K\left(\tau_d \dot{x} + x\right) \tag{11.41}$$

where $K$ is the forward gain and $\tau$ is the damping time constant. Our dynamical equation is now:

$$\ddot{x} + K\left(\tau_d \dot{x} + x\right) = u_d \tag{11.42}$$

The damping term will cause the transients to die out. When that happens, the second derivative and first derivatives of $x$ are zero and we end up with an offset:

$$x = \frac{u}{K} \tag{11.43}$$

This is generally not desirable. You could increase $K$ until the offset were small, but that would mean your actuator would need to produce higher forces or torques. What we have at the moment is a proportional derivative (PD) controller. Let's add another term to the controller:

$$u_c = -K\left(\tau_d \dot{x} + x + \frac{1}{\tau_i}\int x\right) \tag{11.44}$$

This is now a proportional integral derivative (PID) controller. There is now a gain proportional to the integral of $x$. We add the new controller, and then take another derivative to get:

$$\dddot{x} + K\left(\tau_d \ddot{x} + \dot{x} + \frac{1}{\tau_i}x\right) = \dot{u}_d \tag{11.45}$$

Now in steady state:

$$x = \frac{\tau_i}{K}\dot{u}_d \tag{11.46}$$

If $u$ is constant, the offset is zero. Define $s$ as the derivative operator.

$$s = \frac{d}{dt} \tag{11.47}$$

Then:

$$s^3 x(s) + K\left(\tau_d s^2 x(s) + s x(s) + \frac{1}{\tau_i}x(s)\right) = s u_d(s) \tag{11.48}$$

Note that:

$$\frac{u_c(s)}{x(s)} = K\left(1 + \tau_d s + \frac{1}{\tau_i s}\right) \tag{11.49}$$

where $\tau_d$ is the rate time constant, which is how long the system will take to damp and $\tau_i$ is how fast the system will integrate out a steady disturbance.

where $s = j\omega$ and $j = \sqrt{-1}$. The closed loop transfer function is:

$$\frac{x(s)}{u_d(s)} = \frac{s}{s^3 + K\tau_d s^2 + Ks + K/\tau_i} \tag{11.50}$$

The desired closed loop transfer function is:

$$\frac{x(s)}{u_d(s)} = \frac{s}{(s+\gamma)(s^2 + 2\zeta\sigma s + \sigma^2)} \tag{11.51}$$

or

$$\frac{x(s)}{u_d(s)} = \frac{s}{s^3 + (\gamma + 2\zeta\sigma)s^2 + \sigma(\sigma + 2\zeta\gamma)s + \gamma\sigma^2} \tag{11.52}$$

The parameters are:

$$K = \sigma(\sigma + 2\zeta\gamma) \tag{11.53}$$

$$\tau_i = \frac{\sigma + 2\zeta\gamma}{\gamma\sigma} \tag{11.54}$$

$$\tau_d = \frac{\gamma + 2\zeta\sigma}{\sigma(\sigma + 2\zeta\gamma)} \tag{11.55}$$

This is a design for a PID. However, it is not possible to write this in the desired state space form:

$$\dot{x} = Ax + Au \tag{11.56}$$

$$y = Cx + Du \tag{11.57}$$

because it has a pure differentiator. We need to add a filter to the rate term so that it looks like:

$$\frac{s}{\tau_r s + 1} \tag{11.58}$$

instead of $s$. We aren't going to derive the constants and will leave it as an exercise for the reader. The code for the PID is in `PID`.

```
function [a, b, c, d] = PID(  zeta, omega, tauInt, omegaR, tSamp )

% Demo
if( nargin < 1 )
  Demo;
  return
end

% Input processing
if( nargin < 4 )
  omegaR = [];
end

% Default roll-off
if( isempty(omegaR) )
  omegaR = 5*omega;
end

% Compute the PID gains
omegaI   = 2*pi/tauInt;

c2   = omegaI*omegaR;
c1   = omegaI+omegaR;
b1   = 2*zeta*omega;
b2   = omega^2;
g    = c1 + b1;
```

```
kI   = c2*b2/g;
kP   = (c1*b2 + b1.*c2  - kI)/g;
kR   = (c1*b1 + c2 + b2 - kP)/g;

% Compute the state space model
a    =  [0 0;0 -g];
b    =  [1;g];
c    =  [kI -kR*g];
d    =  kP + kR*g;

% Convert to discrete time
if( nargin > 4 )
  [a,b] = CToDZOH(a,b,tSamp);
end
```

It is interesting to evaluate the effect of the integrator. This is shown in Figure 11.9. The code is the demo in PID. Instead of numerically integrating the differential equations we convert them into sampled time and propagate them. This is handy for linear equations. The double integrator equations are in the form:

$$x_{k+1} = ax_k + bu_k \quad (11.59)$$

$$y = cx_k + du_k \quad (11.60)$$

This is the same form as the PID controller.

```
% The double integrator plant
dT              = 0.1; % s
aP              = [0 1;0 0];
bP              = [0;1];
[aP, bP]        = CToDZOH( aP, bP, dT );

% Design the controller
[a, b, c, d]    = PID(  1, 0.1, 100, 0.5, dT );

% Run the simulation
n    = 2000;
p    = zeros(2,n);
x    = [0;0];
xC   = [0;0];

for k = 1:n
  % PID Controller
  y        = x(1);
  xC       = a*xC + b*y;
  uC       = c*xC + d*y;
  p(:,k)   = [y;uC];
  x        = aP*x + bP*(1-uC); % Unit step response
end
```

***Figure 11.9:*** Proportional integral derivative control given a unit input.

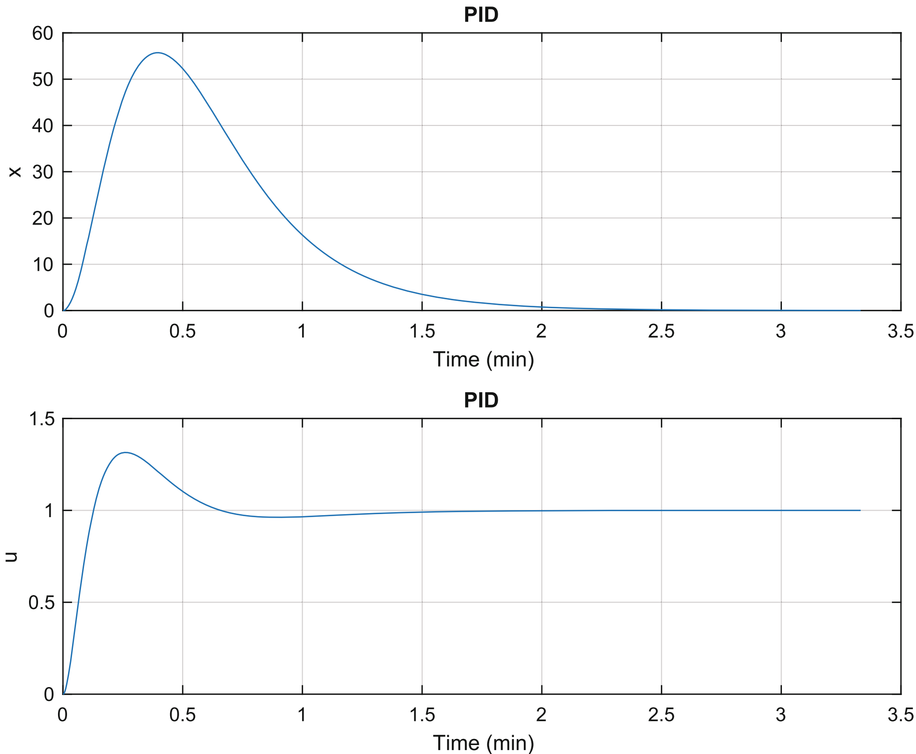


It takes about 2 minutes to drive $x$ to zero, which is close to the 100 seconds specified for the integrator.

## 11.9 PID Control of Pitch

### 11.9.1 Problem

We want to control the pitch angle of an aircraft with a PID controller.

### 11.9.2 Solution

Write a script to implement the controller with the PID controller and pitch dynamic inversion compensation.

### 11.9.3 How It Works

The PID controller changes the elevator angle to produce a pitch acceleration to rotate the aircraft. The elevator is the moveable horizontal surface that is usually on the tail wing of an aircraft. In addition, additional elevator movement is needed to compensate for changes in the accelerations due to lift and drag as the aircraft changes its pitch orientation. This is done using the pitch dynamic inversion function. This returns the pitch acceleration, which must be compensated for when applying the pitch control.

```
function qDot = PitchDynamicInversion( x, d )

if( nargin < 1 )
  qDot = DataStructure;
  return
end

u     = x(1);
w     = x(2);
h     = x(5);

rho   = AtmDensity( h );

alpha = atan(w/u);
cA    = cos(alpha);
sA    = sin(alpha);

v     = sqrt(u^2 + w^2);
pD    = 0.5*rho*v^2; % Dynamic pressure

cL    = d.cLAlpha*alpha;
cD    = d.cD0 + d.k*cL^2;

drag  = pD*d.s*cD;
lift  = pD*d.s*cL;

z     = -lift*cA - drag*sA;
m     = d.c*z;
qDot  = m/d.inertia;
```

The simulation incorporating the controls is `AircraftSim` below. There is a flag to turn on control and another to turn on the learning control. We command a 0.2-radian pitch angle using the PID control. The results are shown in Figure 11.10, Figure 11.11 and Figure 11.12.

The maneuver increases the drag and we don't adjust the throttle to compensate. This will cause the airspeed to drop. In implementing the controller we neglected to consider coupling between states, but this can be added easily.

***Figure 11.10:*** Aircraft pitch angle change. The aircraft oscillates because of the pitch dynamics.

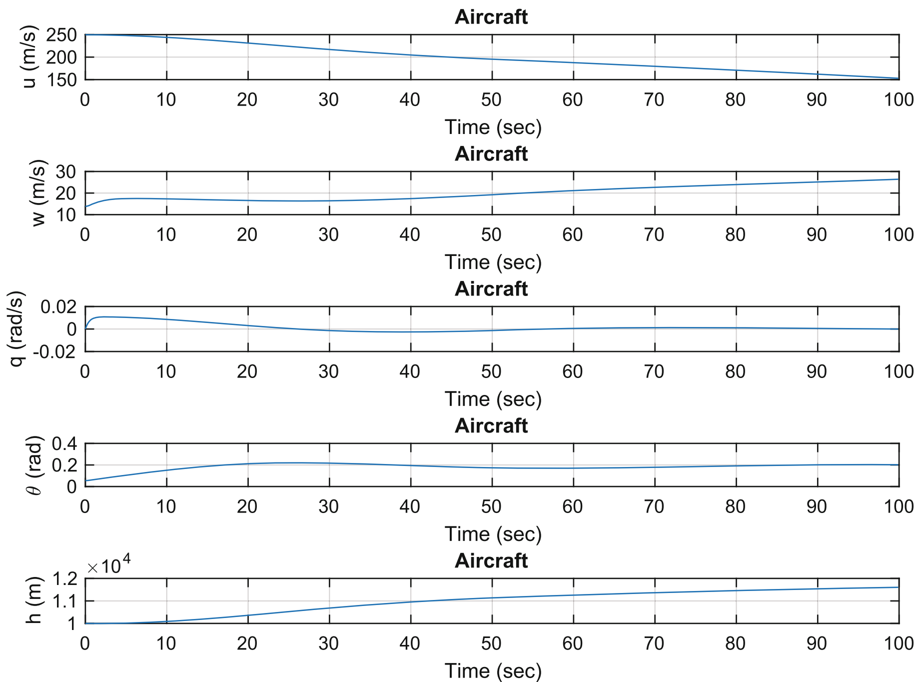


## 11.10 Neural Net for Pitch Dynamics

### 11.10.1 Problem

We want a nonlinear inversion controller with a PID controller and the sigma-pi neural net.

### 11.10.2 Solution

Train the neural net with a script that takes the angle and velocity squared input and computes the pitch acceleration error.

### 11.10.3 How It Works

The `PitchNeuralNetTraining` script computes the pitch acceleration for a slightly different set of parameters. It then processes the delta-acceleration. The script passes a range of pitch angles to the function and learns the acceleration. We use the velocity squared as an input because the dynamic pressure is proportional to the velocity squared. The base acceleration (in `dRHSL`) is for our "a-priori" model. `dRHS` is the measured values. We assume that these are obtained during flight testing.

```
% This is from flight testing
dRHS    = RHSAircraft;    % Get the default data structure has F-16 data
h   = 10000;
```

***Figure 11.11:*** Aircraft pitch angle change. Notice the changes in lift and drag with angle.

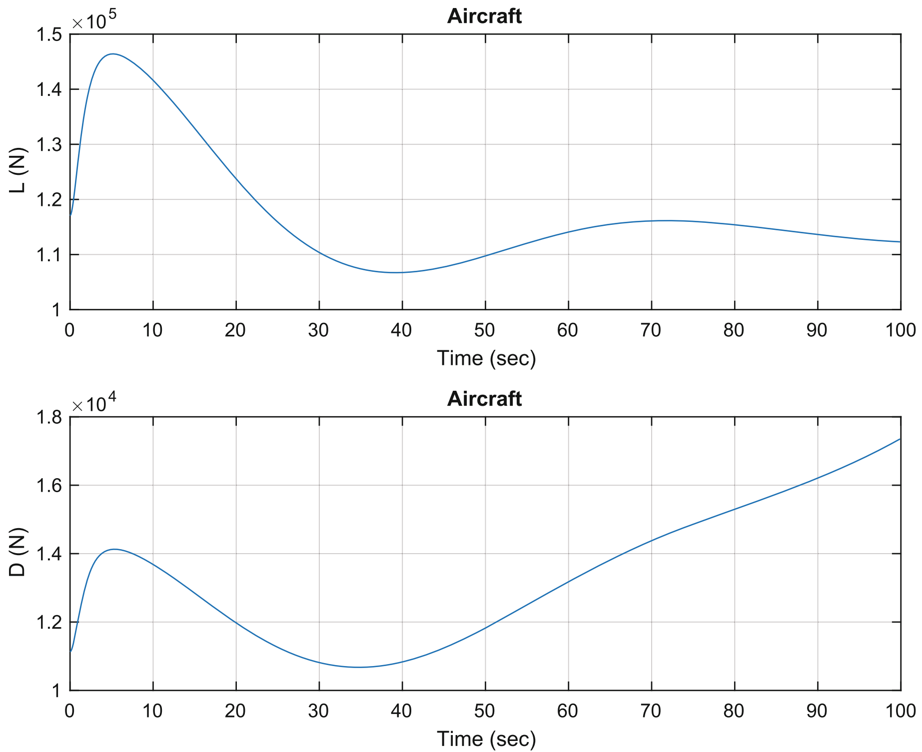


***Figure 11.12:*** Aircraft pitch angle change. The PID acceleration is much lower than the pitch inversion acceleration.

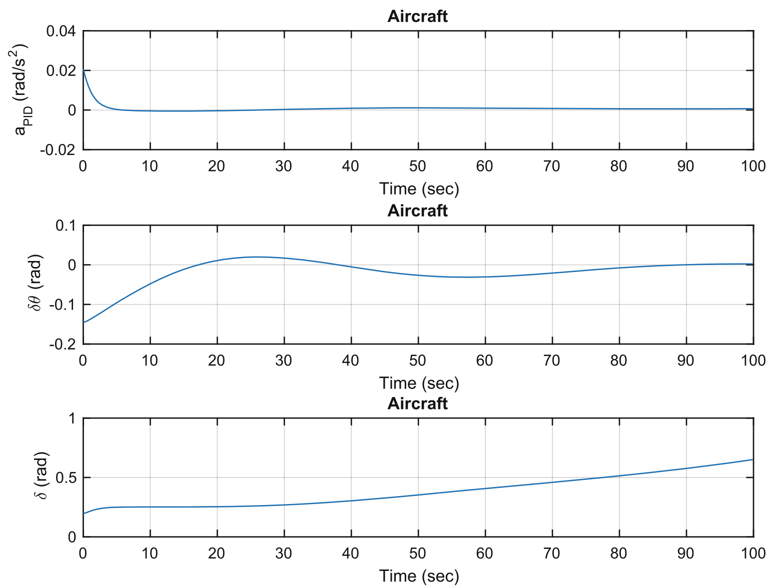

```
gamma          = 0.0;
v              = 250;

% Get the equilibrium state
[x,  dRHS.thrust, deltaEq, cost] = EquilibriumState( gamma, v, h,
   dRHS );

% Angle of attack
alpha          = atan(x(2)/x(1));
cA             = cos(alpha);
sA             = sin(alpha);

% Create the assumed properties
dRHSL          = dRHS;
dRHSL.cD0      = 2.2*dRHS.cD0;
dRHSL.k        = 1.0*dRHSL.k;

% 2 inputs
xNN      = zeros(2,1);
d        = SigmaPiNeuralNet;
[~, d]   = SigmaPiNeuralNet( 'initialize', xNN, d );

theta    = linspace(0,pi/8);
v        = linspace(300,200);
n        = length(theta);
aT       = zeros(1,n);
aM       = zeros(1,n);

for k    = 1:n
  x(4)   = theta(k);
  x(1)   = cA*v(k);
  x(2)   = sA*v(k);
  aT(k)  = PitchDynamicInversion( x, dRHSL );
  aM(k)  = PitchDynamicInversion( x, dRHS  );
end

% The delta pitch acceleration
dA         = aM - aT;

% Inputs to the neural net
v2         = v.^2;
xNN        = [theta;v2];

% Outputs for training
d.y        = dA';
[aNN, d]   = SigmaPiNeuralNet( 'batch learning', xNN, d );
```

```
% Save the data for the aircraft simulation
thisPath = fileparts(mfilename('fullpath'));
save( fullfile(thisPath,'DRHSL'),'dRHSL' );
save( fullfile(thisPath,'DNN'), 'd'  );

for j = 1:size(xNN,2)
  aNN(j,:) = SigmaPiNeuralNet( 'output', xNN(:,j), d );
end

% Plot the results
```

The script first finds the equilibrium state using `EquilibriumState`. It then sets up the sigma-pi neural net using `SigmaPiNeuralNet`. `PitchDynamicInversion` is called twice, once to get the model aircraft acceleration `aM` (the way we want the aircraft to behave) and once to get the true acceleration `aT`. The delta acceleration, `dA` is used to train the neural net. The neural net produces `aNN`. The resulting weights are saved in a .mat file for use in `AircraftSim`. The simulation uses `dRHS`, but our pitch acceleration model uses `dRHSL`. The latter is saved in another .mat file.

```
>> PitchNeuralNetTraining
Velocity              250.00 m/s
Altitude            10000.00 m
Flight path angle       0.00 deg
Z speed                13.84 m/s
Thrust              11148.95 N
Angle of attack         3.17 deg
Elevator               11.22 deg
Initial cost        9.62e+01
Final cost          1.17e-17
```

As can be seen, the neural net reproduces the model very well. The script also outputs `DNN.mat`, which contains the trained neural net data.

## 11.11 Nonlinear Simulation

### 11.11.1 Problem

We want to demonstrate our learning control system for controlling the longitudinal dynamics of an aircraft.

***Figure 11.13:*** Neural net fitted to the delta acceleration.

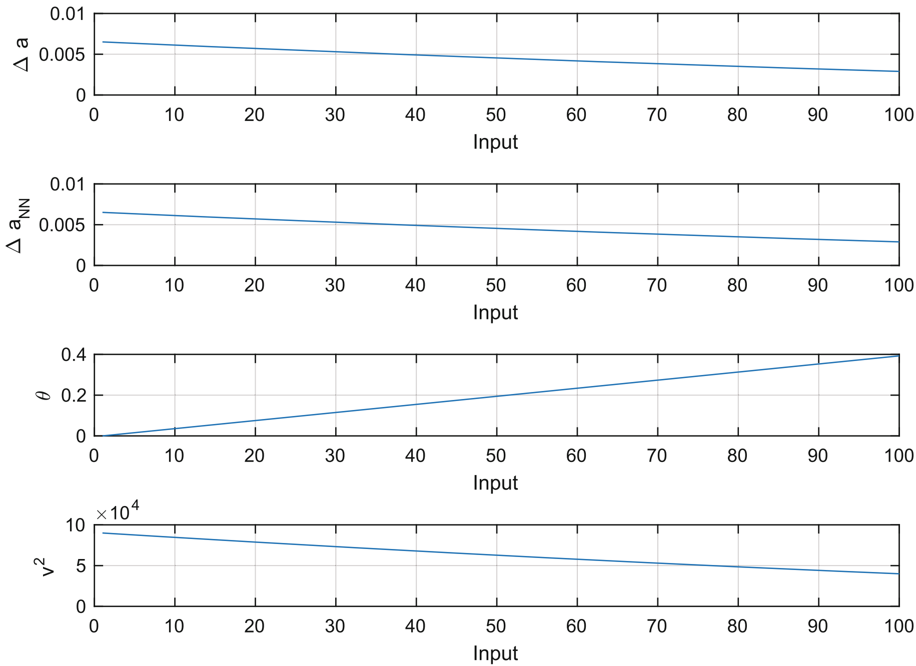


### 11.11.2 Solution

Enable the control functions to the simulation script described in `AircraftSimOpenLoop`.

### 11.11.3 How It Works

After training, the neural net in the previous recipe we set `addLearning` to true. The weights are read in. We command a 0.2-radian pitch angle using the PID learning control. The results are shown in Figure 11.14, Figure 11.15, and Figure 11.16. The figures show without learning control on the left and with learning control on the right.

Learning control helps the performance of the controller. However, the weights are fixed throughout the simulation. Learning occurs prior to the controller becoming active. The control system is still sensitive to parameter changes since the learning part of the control was computed for a pre-determined trajectory. Our weights were determined only as a function of pitch angle and velocity squared. Additional inputs would improve the performance. There are many opportunities for you to try to expand and improve the learning system.

***Figure 11.14:*** Aircraft pitch angle change. Lift and drag variations are shown.

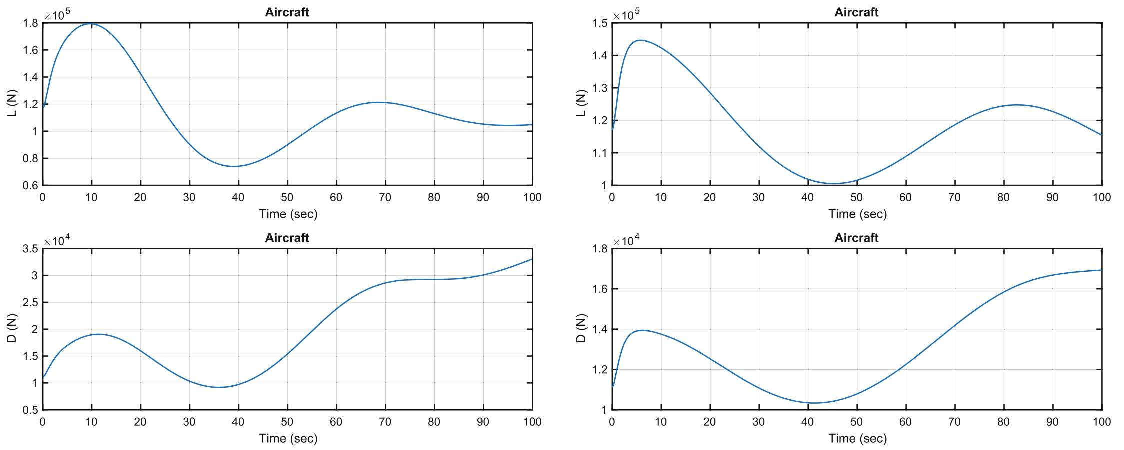


***Figure 11.15:*** Aircraft pitch angle change. Without learning control, the elevator saturates.

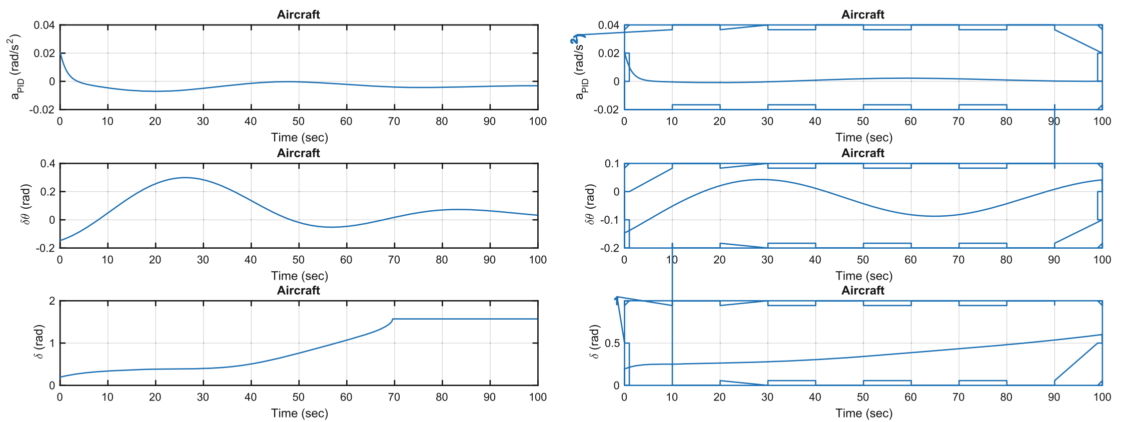

***Figure 11.16:*** Aircraft pitch angle change. The PID acceleration is much lower than the pitch dynamic inversion acceleration.

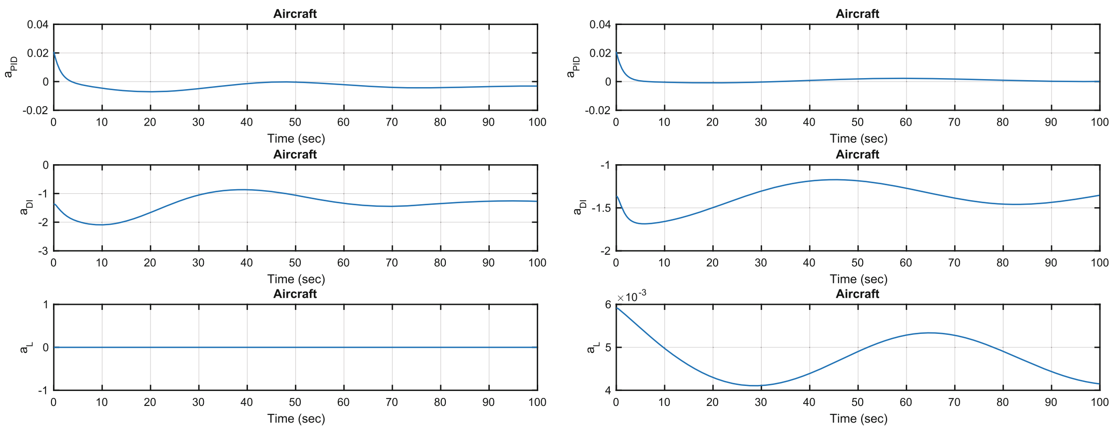


## 11.12 Summary

This chapter has demonstrated adaptive or learning control for an aircraft. You learned about model tuning, model reference adaptive control, adaptive control, and gain scheduling. You also learned how to use a neural net as part of an aircraft control system. Table 11.3 lists the functions and scripts included in the companion code.

***Table 11.3:*** Chapter Code Listing

| **File** | **Description** |
|---|---|
| AircraftSim | Simulation of the longitudinal dynamics of an aircraft. |
| AtmDensity | Atmospheric density using a modified exponential model. |
| EquilibriumState | Finds the equilibrium state for an aircraft. |
| PID | Implements a PID controller. |
| PitchDynamicInversion | Pitch angular acceleration. |
| PitchNeuralNetTraining | Train the pitch acceleration neural net. |
| QCR | Generates a full state feedback controller. |
| RecursiveLearning | Demonstrates recursive neural net training or learning. |
| RHSAircraft | Right-hand side for aircraft longitudinal dynamics. |
| SigmaPiNeuralNet | Implements a sigma-pi neural net. |
| Sigmoid | Plots a sigmoid function. |

CHAPTER 12

# Multiple Hypothesis Testing

## 12.1 Overview

Tracking is the process of determining the position of other objects as their position changes with time. Air traffic control radar systems are used to track aircraft. Aircraft in flight must track all nearby objects to avoid collisions and to determine if they are threats. Automobiles with radar cruise control use their radar to track cars in front of them so that the car can maintain safe spacing and avoid a collision.

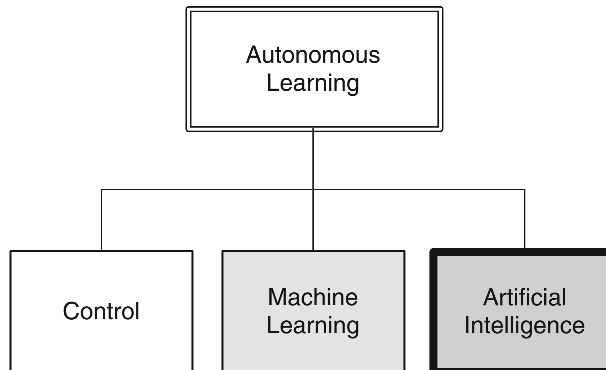


When you are driving, you maintain situation awareness by identifying nearby cars and figuring out what they are going to do next. Your brain processes data from your eyes to characterize a car. You track objects by their appearance, since, in general, the cars around you all look different. Of course, at night you only have tail lights so the process is harder. You can often guess what each car is going to do, but sometimes you guess wrong and that can lead to collisions.

Radar systems just see blobs. Cameras should be able to do what your eyes and brain do, but that requires a lot of processing. As noted, at night it is hard to reliably identify a car. As the blobs are measured by radar we want to collect all blobs, as they vary in position and speed, and attach them to a particular car's track. This way we can reliably predict where it will go next. This leads to the topic of this chapter, track-oriented multiple hypothesis testing (MHT).

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_12

***Table 12.1:*** Multiple Hypothesis Testing Terms

| **Term** | **Definition** |
|---|---|
| Clutter | Transient objects of no interest to the tracking system. |
| Cluster | A collection of tracks that are linked by common observations. |
| Error Ellipsoid | An ellipsoidal volume around an estimated position. |
| Family | A set of tracks with a common root node. At most, one track per family can be included in a hypothesis. A family can at most represent one target. |
| Gate | A region around an existing track position. Measurements within the gate are associated with the track. |
| Hypothesis | A set of tracks that do not share any common observations. |
| N-Scan Pruning | Using the track scores from the last N scans of data to prune tracks. The count starts from a root node. When the tracks are pruned, a new root node is established. |
| Observation | A measurement that indicates the presence of an object. The observation may be of a target or be spurious. |
| Pruning | Removal of low-score tracks. |
| Root Node | An established track to which observations can be attached and which may spawn additional tracks. |
| Scan | A set of data taken simultaneously. |
| Target | An object being tracked. |
| Trajectory | The path of a target. |
| Track | A trajectory that is propagated. |
| Track Branch | A track in a family that represents a different data association hypothesis. Only one branch can be correct. |
| Track Score | The log-likelihood ratio for a track. |

Track-oriented MHT is a powerful technique for assigning measurements to tracks of objects when the number of objects is unknown or changing. It is absolutely essential for accurate tracking of multiple objects. MHT terms are defined in Table 12.1.

Hypotheses are sets of tracks with consistent data, that is, where no measurements are assigned to more than one track. The track-oriented approach recomputes the hypotheses using the newly updated tracks after each scan of data is received. Rather than maintaining, and expanding, hypotheses from scan to scan, the track-oriented approach discards the hypotheses formed on scan $k-1$. The tracks that survive pruning are propagated to the next scan $k$ where new tracks are formed, using the new observations, and reformed into hypotheses. Except for the necessity to delete some tracks based upon low probability, no information is lost because the track scores that are maintained contain all the relevant statistical data.

The software in this chapter uses a powerful track-pruning algorithm that does the pruning in one step. Because of its speed, ad-hoc pruning methods are not required, leading to more robust and reliable results. The track management software is, as a consequence, quite simple.

The MHT Module requires the GNU Linear Programming Kit (GLPK; http://www.gnu.org/software/glpk/) and specifically, the MATLAB mex wrapper GLPKMEX (http://glpkmex.sourceforge.net). Both are distributed under the GNU license. Both the GLPK library and the GLPKMEX program are operating-system-dependent and must be compiled from the source code on your computer. Once GLPK is installed, the mex must be generated from MATLAB from the GLPKMEX source code.

The command that is executed from MATLAB to create the mex should look like:

```
mex -v -I/usr/local/include glpkcc.cpp /usr/local/lib/libglpk.a
```

where the “v” specifies verbose printout and you should replace `/usr/local` with your operating-system dependent path to your installation of GLPK. The resulting mex file (Mac) is:

```
glpkcc.mexmaci64
```

The MHT software was tested with GLPK version 4.47 and GLPKMEX version 2.11.

## 12.2 Theory

### 12.2.1 Introduction

Figure 12.1 shows the general tracking problem in the context of automobile tracking. Two scans of data are shown. When the first scan is done, there are two tracks. The uncertainty ellipsoids are shown and they are based on all previous information. In the $k-1$ scan (a scan is a set of measurements taken at the same time), three measurements are observed. Each scan has multiple measurements, the measurements in each new scan are numbered beginning with 1, and the measurement numbers are not meant to imply any correlation across subsequent scans. One and 3 are within the ellipsoids of the two tracks, but 2 is in both. It may be a measurement of either of the tracks or a spurious measurement. In scan $k$, four measurements are taken. Only measurement 4 is in one of the uncertainty ellipsoids. Three may be interpreted as spurious, but it is actually because of a new track from a third vehicle that it separates from the blue track. Measurement 1 is outside of the red ellipsoid, but is actually a good measurement of the red track and (if correctly interpreted) indicates that the model is erroneous. 4 is a good measurement of the blue track and indicates that the model is valid. Measurement 2 of scan k is outside both uncertainty ellipsoids. The illustration shows how the tracking system should behave, but without the tracks it would be difficult to interpret the measurements. As shown a measurement can be:

1. Valid
2. Spurious
3. A new track

“Spurious” means that the measurement is not associated with any tracked object and isn’t a new track. We can’t determine the nature of any measurement without going through the MHT process.

***Figure 12.1:*** Tracking problem.

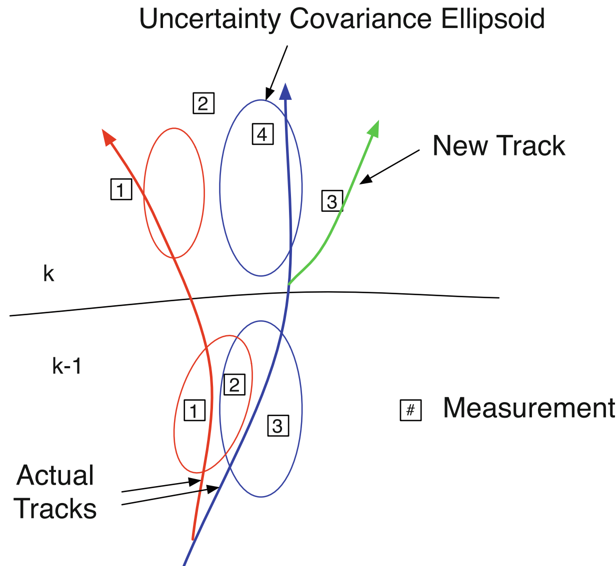


We define a contact as an observation where the signal-to-noise ratio is above a certain threshold. The observation then constitutes a measurement. Low signal-to-noise ratio observations can happen in both optical and radar systems. Thresholding reduces the number of observations that need to be associated with tracks, but may lose valid data. An alternative is to treat all observations as valid, but adjust the measurement error accordingly.

Valid measurements must then be assigned to tracks. An ideal tracking system would be able to categorize each measurement accurately and then assign them to the correct track. The system must also be able to identify new tracks and remove tracks that no longer exist. A tracking system may have to deal with hundreds of objects (perhaps after a collision or because of debris in the road).

A sophisticated system should be able to work with multiple objects as groups or clusters if the objects are more or less moving in the same direction. This reduces the number of states a system must handle. If a system handles groups, then it must be able to handle groups spawning from groups.

If we were confident that we were only tracking one vehicle, all of the data might be incorporated into the state estimate. An alternative is to incorporate only the data within the covariance ellipsoids and treat the remainders as outliers. If the latter strategy were taken, it would be sensible to remember that data in case future measurements were also “outliers,” in which case the filter might go back and incorporate different sets of outliers into the solution. This could easily happen if the model were invalid, for example, if the vehicle, which had been cruising at a constant speed, suddenly began maneuvering and the filter model did not allow for maneuvers.

The multiple model filters helps with the erroneous model problem and should be used any time a vehicle might change mode. It does not tell us how many vehicles we are tracking,

however. With multiple models, each model would have its own error ellipsoids and the measurements would fit one better than the other, assuming that one of the models was a reasonable model for the tracked vehicle in its current mode.

### 12.2.2 Example

Referring to Figure 12.1 in the first scan, we have three measurements. 1 and 3 are associated with existing tracks and are used to update those tracks. 2 could be associated with either. It might be a spurious measurement or it could be a new track, so the algorithm forms a new hypothesis. In scan 2 measurement 4 is associated with the blue track. 1, 2, and 3 are not within the error ellipsoids of either track. Since the figure shows the true track, we can see that 1 is associated with the red track. Both 1 and 2 are just outside the error ellipsoid for the red track. Measurement 2 in scan 2 might be consistent with measurement 2 in scan 1 and could result in a new track. Measurement 3 in scan 2 is a new track, but we likely don't have enough information to create a track until we have more scans of data.

### 12.2.3 Algorithm

In classical multiple target tracking, [24], the problem is divided into two steps, association and estimation. Step 1 associates contacts with targets and step 2 estimates each target's state. Complications arise when there is more than one reasonable way to associate contacts with targets. The multiple hypothesis testing (MHT) approach is to form alternative hypotheses to explain the source of the observations. Each hypothesis assigns observations to targets or false alarms.

There are two basic approaches to MHT [3]. The first, following Reid [21], operates within a structure in which hypotheses are continually maintained and updated as observation data are received. In the second, the track-oriented approach to MHT, tracks are initiated, updated, and scored before being formed into hypotheses. The scoring process consists of comparing the likelihood that the track represents a true target versus the likelihood that it is a collation of false alarms. Thus, unlikely tracks can be deleted before the next stage, in which tracks are formed into a hypothesis. It is a good thing to discard the old hypotheses and start from scratch each time because this approach maintains the important track data while preventing an explosion of an impractically large number of hypotheses.

The track-oriented approach recomputes the hypotheses using the newly updated tracks after each scan of data is received. Rather than maintaining, and expanding, hypotheses from scan to scan, the track-oriented approach discards the hypotheses formed on scan k-1. The tracks that survive pruning are predicted to the next scan k where new tracks are formed, using the new observations, and reformed into hypotheses. Except for the necessity to delete some tracks based upon low probability or N-scan pruning, no information is lost because the track scores that are maintained contain all the relevant statistical data.

Track scoring is done using log-likelihood ratios. $LR$ is the likelihood ratio, $LLR$ the log-likelihood ratio, and $L$ is the likelihood.

$$L(K) = \log[\mathrm{LR}(K)] = \sum_{k=1}^{K} [\mathrm{LLR}_K(k) + \mathrm{LLR}_S(k)] + \log[L_0] \tag{12.1}$$

where the subscript $K$ denotes kinematic (position) and the subscript $S$ denotes signal (measurement). It is assumed that the two are statistically independent.

$$L_0 = \frac{P_0(H_1)}{P_0(H_0)} \tag{12.2}$$

where $H_1$ and $H_0$ are the true target and false alarm hypotheses. $\log$ is a natural logarithm. The likelihood ratio for the kinematic data is the probability that the data are a result of the true target divided by the probability that the data are due to a false alarm:

$$\mathrm{LR}_K = \frac{p(D_K|H_1)}{p(D_K|H_0)} = \frac{e^{-d^2/2}/((2\pi)^{M/2}\sqrt{|S|}}{1/V_C} \tag{12.3}$$

where

1. $M$ in the denominator of the third formula is the measurement dimension
2. $V_C$ is the measurement volume
3. $S = HPT^T + R$ the measurement residual covariance matrix
4. $d^2 = y^T S^{-1} y$ is the normalized statistical distance for the measurement

The statistical distance is defined by the residual $y$, the difference between the measurement and the estimated measurement, and the covariance matrix $S$. The numerator is the multivariate Gaussian.

### 12.2.4 Measurement Assignment and Tracks

The following are the rules for each measurement:

1. Each measurement creates a new track.
2. Each measurement in each gate updates the existing track. If there is more than one measurement in a gate, the existing track is duplicated with the new measurement.
3. All existing tracks are updated with a “missed” measurement, creating a new track.

Figure 12.2 gives an example. We are starting with two tracks. There are two tracks and three measurements. All three measurements are in the gate for track 1, but only one is in the gate for track 2. Each measurement produces a new track. The three measurements produce three tracks based on track 1 and the one measurement produces one track based on track 2.

There are three types of tracks created from each scan, in general:

1. An existing track is updated with a new measurement, assuming it corresponds to that track.
2. An existing track is carried along with no update, assuming that no measurement was made for it in that scan.

3. A completely new track is generated for each measurement, assuming that the measurement represents a new object.

Each track also spawns a new track assuming that there was no measurement for the track. Thus in this case, three measurements and two tracks result in nine new tracks. Tracks 7–9 are initiated based only on the measurement, which may not be enough information to initiate the full state vector. If this is the case, there would be an infinite number of tracks associated with each measurement, not just one new track. If we have a radar measurement we have azimuth, elevation, range, and range rate. This gives all position states and one velocity state.

***Figure 12.2:*** Measurement and gates. M0 is an "absent" measurement. An absent measurement is one that should exist, but does not.

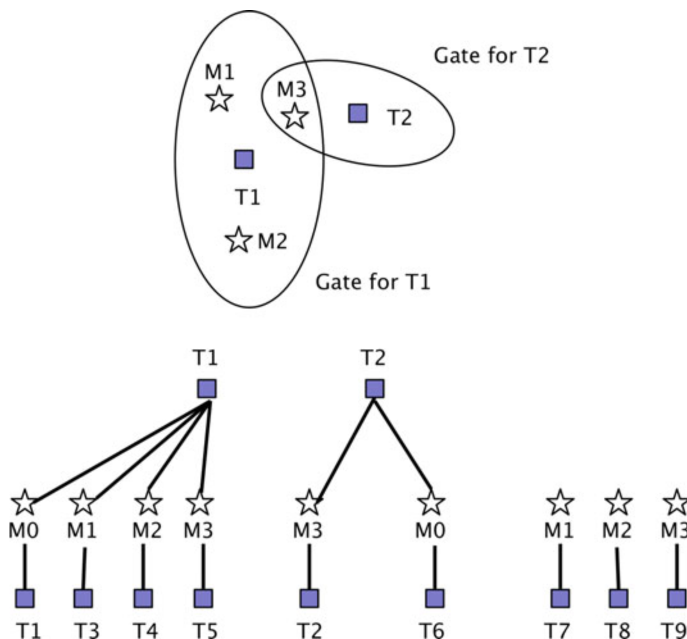


### 12.2.5 Hypothesis Formation

In MHT, a valid hypothesis is any compatible set of tracks. In order for two or more tracks to be compatible, they cannot describe the same object, and they cannot share the same measurement at any of the scans. The task in hypothesis formation is to find one or more combinations of tracks that: 1) are compatible, and 2) maximize some performance function.

Before discussing the method of hypothesis formation, it is useful to first consider track formation and how tracks are associated with unique objects. New tracks may be formed in one of two ways:

1. The new track is based on some existing track, with the addition of a new measurement.
2. The new track is NOT based on any existing tracks; it is based solely on a single new measurement.

Recall that each track is formed as a sequence of measurements across multiple scans. In addition to the raw measurement history, every track also contains a history of state and covariance data that is computed from a Kalman Filter. Kalman Filters are explored in Chapter 8. When a new measurement is appended to an existing track, we are spawning a new track that includes all of the original track's measurements, plus this new measurement. Therefore, the new track is describing the same object as the original track.

A new measurement can also be used to generate a completely new track that is independent of past measurements. When this is done, we are effectively saying that the measurement does not describe any of the objects that are already being tracked. It therefore must correspond to a new/different object.

In this way, each track is given an object ID to distinguish which object it describes. Within the context of track-tree diagrams, all of the tracks inside the same track-tree have the same object ID. For example, if at some point there are 10 separate track-trees, this means that 10 separate objects are being tracked in the MHT system. When a valid hypothesis is formed, it may turn out that only a few of these objects have compatible tracks.

The hypothesis formation step is formulated as a mixed integer linear program (MILP) and solved using GLPK. Each track is given an aggregate score that reflects the component scores attained from each measurement. The MILP formulation is constructed to select a set of tracks that add up to give the highest score, such that:

1. No two tracks have the same object ID

2. No two tracks have the same measurement index for any scan

In addition, we extend the formulation with an option to solve for multiple hypotheses, rather than just one. The algorithm will return the "M best" hypotheses, in descending order of score. This enables tracks to be preserved from alternate hypotheses that may be very close in score to the best.

### 12.2.6 Track Pruning

The N-scan track pruning is carried out every step using the last $n$ scans of data. We employ a pruning method in which the following tracks are preserved:

- Tracks with the "N" highest scores

- Tracks that are included in the "M best" hypotheses

- Tracks that have both 1) the object ID and 2) the first "P" measurements found in the "M best" hypotheses.

We use the results of hypothesis formation to guide track pruning. The parameters N, M, and P can be tuned to improve performance. The objective with pruning is to reduce the number of tracks as much as possible, while not removing any tracks that should be part of the actual true hypothesis.

The second item listed above is to preserve all tracks included in the “M best” hypotheses. Each of these is a full path through a track-tree, which is clear. The third item listed above is similar, but less constrained. Consider one of the tracks in the “M best” hypotheses. We will preserve this full track. In addition, we will preserve all tracks that stem from scan “P” of this track.

Figure 12.3 provides an example of which tracks in a track-tree might be preserved. The diagram shows 17 different tracks over five scans. The green track represents one of the tracks found in the set of “M best” hypotheses, from the hypothesis formation step. This track would be preserved. The orange tracks all stem from the node in this track at scan 2. These would be preserved if we set $P = 2$ from the description above.

***Figure 12.3:*** Track pruning example. This shows multiple scans (simultaneous measurements) and how they might be used to remove tracks that do not fit all of the data.

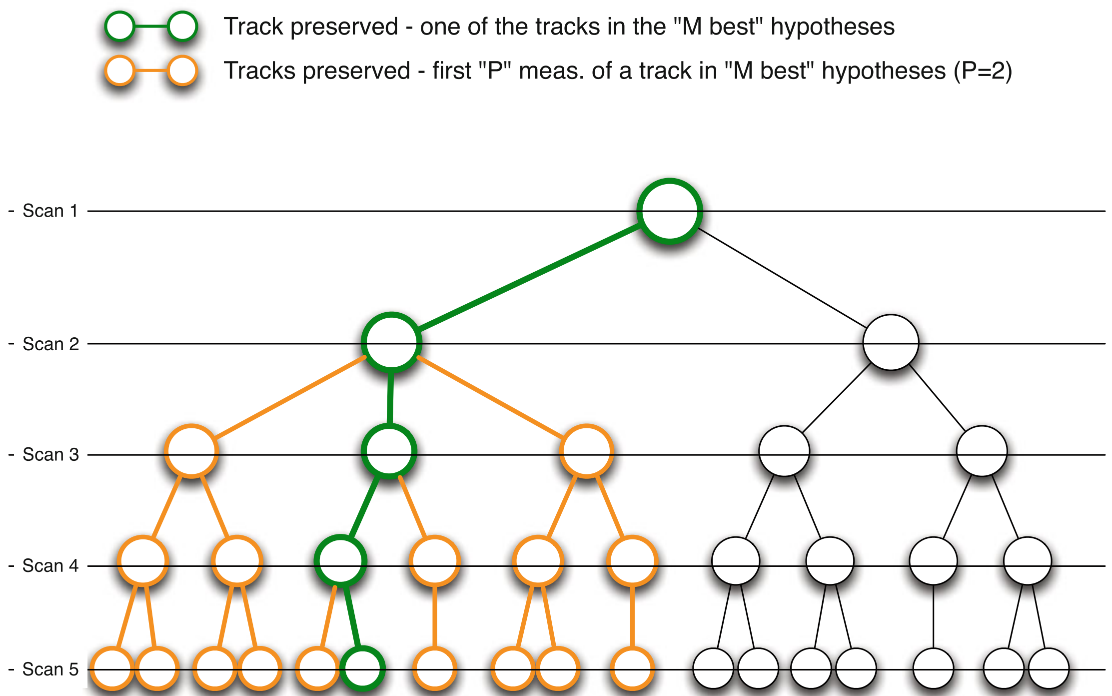

## 12.3 Billiard Ball Kalman Filter

### 12.3.1 Problem

You want to estimate the trajectory of multiple billiard balls. In the billiard ball example, we assume that we have multiple balls moving at once. Let's say we have a video camera placed above the table, and we have software that can measure the position of each ball for each video frame. That software cannot, however, determine the identity of any ball. This is where MHT comes in. We use MHT to develop a set of tracks for the moving balls.

### 12.3.2 Solution

The solution is to create a linear Kalman Filter.

### 12.3.3 How It Works

The core estimation algorithm for the MHT system is the Kalman Filter. The Kalman Filter consists of a simulation of the dynamics and an algorithm to incorporate the measurements. For the examples in this chapter we use a fixed gain Kalman Filter. The model is:

$$x_{k+1} = ax_k + bu_k \tag{12.4}$$

$$y_k = cx_k \tag{12.5}$$

$x_k$ is the state, a column vector that includes position and velocity. $y_k$ is the measurement vector. $u_k$ is the input, the accelerations on the billiard balls. $c$ relates the state to the measurement, $y$. If the only measurement were position then:

$$c = \begin{bmatrix} 1 & 0 \end{bmatrix} \tag{12.6}$$

This is a discrete time equation. Since the second column is zero, it is only measuring position. Let's assume we have no input accelerations. Also assume that the time step is $\tau$. Then our equations become:

$$\begin{bmatrix} s \\ v \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & \tau \\ 0 & 1 \end{bmatrix} \begin{bmatrix} s \\ v \end{bmatrix}_k \tag{12.7}$$

$$y_k = \begin{bmatrix} 1 & 0 \end{bmatrix} \begin{bmatrix} s \\ v \end{bmatrix}_k \tag{12.8}$$

where $s$ is position and $v$ is velocity, $y_k = s$. This says that the new position is the old position plus velocity times time. Our measurement is just position. If there are no external accelerations, the velocity is constant. If we can't measure acceleration directly then this is our model. Our filter will estimate velocity given changes in position.

A track, in this case, is a sequence of $s$. MHT assigns measurements, $y$ to the track. If we know that we have only one object and that our sensor is measuring the track accurately, and doesn't have any false measurements or possibility of missing measurements, we can use the Kalman Filter directly.

The KFBilliardsDemo simulates billiard balls. It includes two functions to represent the dynamics. The first is RHSBilliards, which is the right-hand-side of the billiards ball dynamics, which were just given above. This computes the position and velocity given external accelerations. The function BilliardCollision applies conservation of momentum whenever a ball hits a bumper. Balls can't collide with other balls. The first part of the script is the simulation that generates a measurement vector for all of the balls. The second part of the script initializes one Kalman Filter per ball. This script perfectly assigns measurements to each track. The function KFPredict is the prediction step, i.e., the simulation of the ball motion. It uses the linear model described above. KFUpdate incorporates the measurements. MHTDistance is just for information purposes. The initial positions and velocity vectors of the balls are random. The script fixes the seed for the random number generator to make every run the same, which is handy for debugging. If you comment out this code each run will be different.

Here, we initialize the ball positions.

```
% The number of balls and the random initial position and velocity
d       = struct('nBalls',3,'xLim',[-1 1], 'yLim', [-1 1]);
sigP    = 0.4; % 1 sigma noise for the position
sigV    = 1; % 1 sigma noise for the velocity
sigMeas = 0.00000001; % 1 sigma noise for the measurement

% Set the initial state for  2 sets of position and velocity
x  = zeros(4*d.nBalls,1);
rN = rand(4*d.nBalls,1);

for k = 1:d.nBalls
  j        = 4*k-3;
  x(j  ,1) = sigP*(rN(j  ) - 0.5);
  x(j+1,1) = sigV*(rN(j+1) - 0.5);
  x(j+2,1) = sigP*(rN(j+2) - 0.5);
  x(j+3,1) = sigV*(rN(j+3) - 0.5);
end
```

We then simulate them. Their motion is in a straight line unless they collide with a bumper.

```
% Sensor measurements
nM  = 2*d.nBalls;
y   = zeros(nM,n);
iY  = zeros(nM,1);

for k = 1:d.nBalls
  j = 2*k-1;
  iY(j  )  = 4*k-3;
  iY(j+1)  = 4*k-1;
end

for k = 1:n
```

```
  % Collisions
  x = BilliardCollision( x, d );

  % Plotting
  xP(:,k)       = x;

  % Integrate using a 4th Order Runge-Kutta integrator
  x = RungeKutta(@RHSBilliards, 0, x, dT, d );

  % Measurements with Gaussian random noise
  y(:,k) = x(iY) + sigMeas*randn(nM,1);

end
```

We then process the measurements through the Kalman Filter. `KFPredict` predicts the next position of the balls and `KFUpdate` incorporates measurements. The prediction step does not know about collisions.

```
%% Implement the Kalman Filter

% Covariances
r0      = sigMeas^2*[1;1];      % Measurement covariance
q0      = [1;60;1;60];      % The baseline plant covariance diagonal
p0      = [0.1;1;0.1;1];        % Initial state covariance matrix
   diagonal

% Plant model
a       = [1 dT;0 1];
b       = [dT^2/2;dT];
zA      = zeros(2,2);
zB      = zeros(2,1);

% Create the Kalman Filter data structures. a is for two balls.
for k = 1:d.nBalls
  kf(k) = KFInitialize( 'kf', 'm', x0(4*k-3:4*k), 'x', x0(4*k-3:4*k)
     ,...
                        'a', [a zA;zA a], 'b', [b zB;zB b],'u'
                           ,[0;0],...
                        'h', [1 0 0 0;0 0 1 0], 'p', diag(p0), ...
                        'q', diag(q0),'r', diag(r0) );
 end

% Size arrays for plotting
pUKF = zeros(4*d.nBalls,n);
xUKF = zeros(4*d.nBalls,n);
t    = 0;
```

```
for k = 1:n
  % Run the filters
  for j = 1:d.nBalls

    % Store for plotting
    i            = 4*j-3:4*j;
    pUKF(i,k)    = diag(kf(j).p);
    xUKF(i,k)    = kf(j).m;

    % State update
    kf(j).t    = t;
    kf(j)      = KFPredict( kf(j) );

    % Incorporate the measurements
    i          = 2*j-1:2*j;
    kf(j).y    = y(i,k);
    kf(j)      = KFUpdate( kf(j) );
  end

  t = t + dT;

end
```

The results of the Kalman Filter demo are shown in Figure 12.4, Figure 12.5 and Figure 12.6. The covariances and states for all balls are plotted, but we only show one here. The covariances always follow the same trend with time. As the filter accumulates measurements it adjusts the covariances based on the ratio between the model covariance, i.e., how accurate the model is assumed to be, and the measurement covariances. The covariances are not related to the actual measurements at all. The Kalman Filter errors are shown in Figure 12.6. They are large whenever the ball hits a bumper, since the model does not include collisions with the bumpers. They rapidly decrease because our measurements have little noise.

The following code, excerpted from the above demo, is specialized drawing code to show the billiards on the table. It calls `plot` for each ball. Colors are taken from the array `c` and are blue, green, red, cyan, magenta, yellow, and black. You can run this from the command line once you have computed `xP` and `yP`, which are the $x$ and $y$ positions of the balls. The code uses the legend handles to associate the balls with the tracks in the plot in the legend. It manually sets the limits (`gca` as a handle to the current axes).

```
% Plot the simulation results
NewFigure( 'Billiard␣Balls' )
c  = 'bgrcmyk';
kX = 1;
kY = 3;
s  = cell(1,d.nBalls);
l = [];
for k = 1:d.nBalls
  plot(xP(kX,1),xP(kY,1),['o',c(k)])
```

***Figure 12.4:*** The four balls on the billiards table.

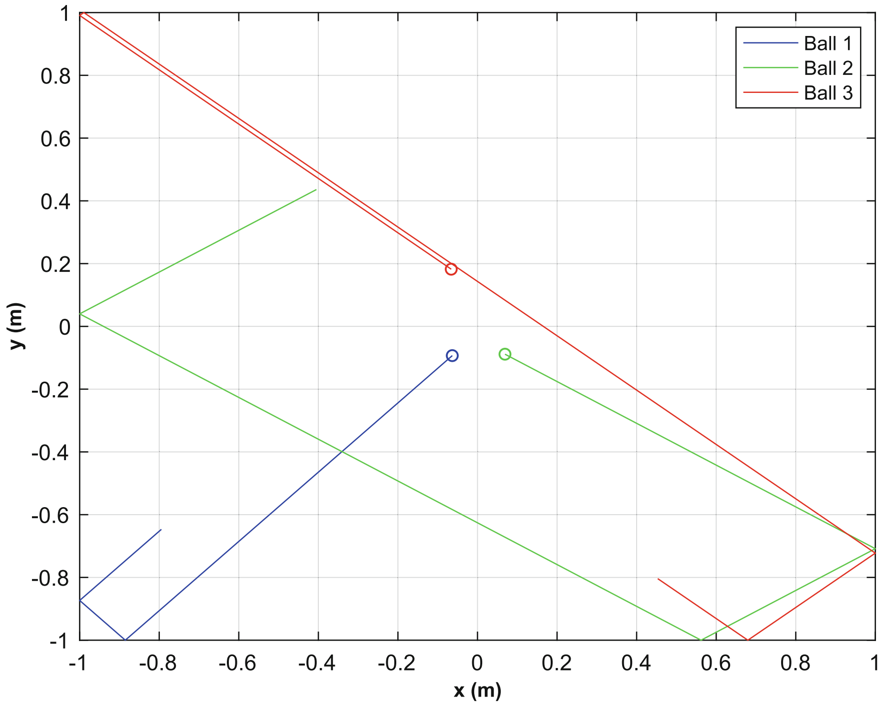


***Figure 12.5:*** The filter covariances.

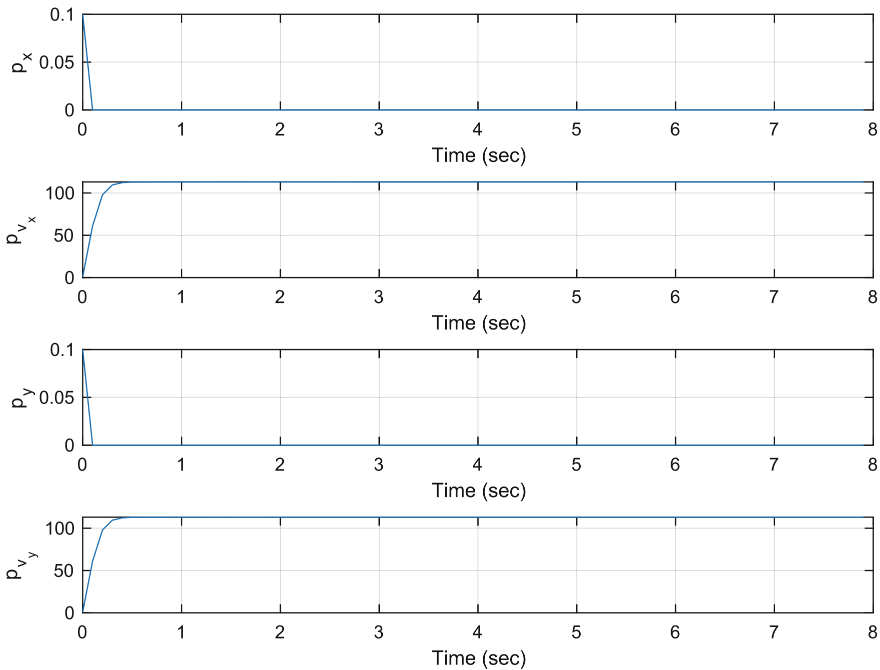

***Figure 12.6:*** The filter errors.

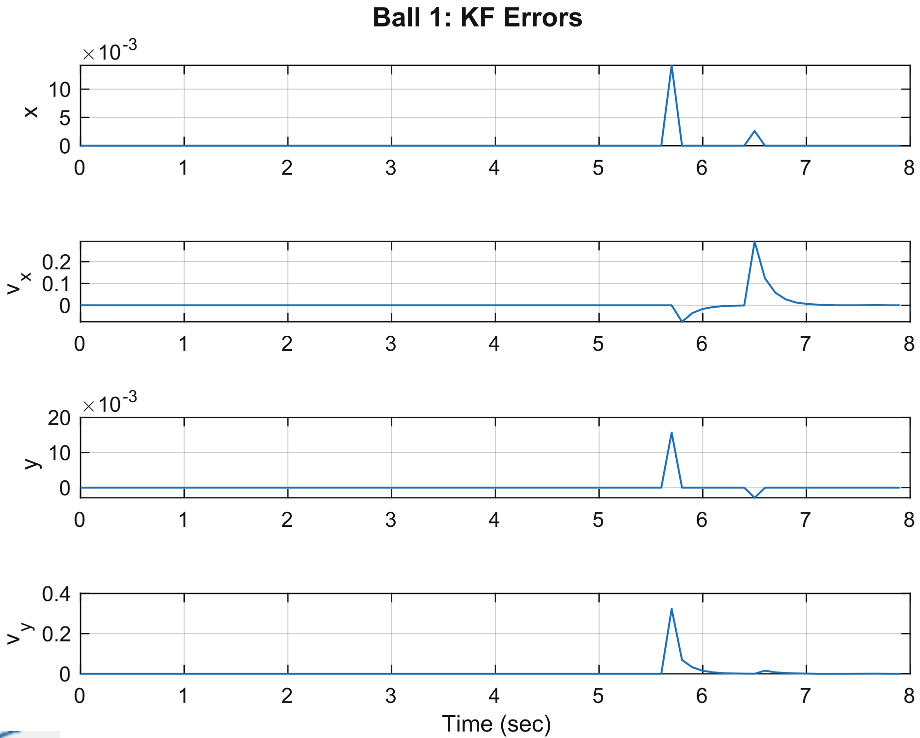

```
    hold on
    l(k)    = plot(xP(kX,:),xP(kY,:),c(k));
    kX      = kX + 4;
    kY      = kY + 4;
    s{k}    = sprintf('Ball␣%d',k);
end

xlabel('x␣(m)');
ylabel('y␣(m)');
set(gca,'ylim',d.yLim,'xlim',d.xLim);
legend(l,s)
grid on
```

You can change the covariances, `sigP,sigV,sigMeas` in the script and see how it impacts the errors and the covariances.

# 12.4 Billiard Ball MHT

## 12.4.1 Problem

You want to estimate the trajectory of multiple billiard balls.

## 12.4.2 Solution

The solution is to create an MHT system with a linear Kalman Filter. This example involves billiard balls bouncing off of the bumpers of a billiard table. The model does not include the bumper collisions.

## 12.4.3 How It Works

The following code adds the MHT functionality. It first runs the demo, just like in the example above, and then tries to sort the measurements into tracks. It only has two balls. When you run the demo you will see the graphical user interface (GUI), Figure 12.7, and the Tree, Figure 12.8, change as the simulation progresses. We only include the MHT code in the following listing.

```
% Create the track data data structure
mhtData = MHTInitialize('probability␣false␣alarm', 0.001,...
                        'probability␣of␣signal␣if␣target␣present',
                           0.999,...
                        'probability␣of␣signal␣if␣target␣absent',
                           0.001,...
                        'probability␣of␣detection', 1, ...
                        'measurement␣volume', 1.0, ...
                        'number␣of␣scans', 3, ...
                        'gate', 0.2,...
                        'm␣best', 2,...
                        'number␣of␣tracks', 1,...
                        'scan␣to␣track␣function',
                           @ScanToTrackBilliards,...
                        'scan␣to␣track␣data',struct('r',diag(r0),'p',
                           diag(p0)),...
                        'distance␣function',@MHTDistance,...
                        'hypothesis␣scan␣last', 0,...
                        'filter␣data',kf(1),...
                        'prune␣tracks', 1,...
                        'remove␣duplicate␣tracks␣across␣all␣trees'
                           ,1,...
                        'average␣score␣history␣weight',0.01,...
                        'filter␣type','kf');

% Create the tracks
for k = 1:d.nBalls
        trk(k) = MHTInitializeTrk( kf(k) );
end
```

```
% Size arrays
b = MHTTrkToB( trk );

%% Initialize MHT GUI
MHTGUI;
MLog('init')
MLog('name','Billiards␣Demo')
TOMHTTreeAnimation( 'initialize', trk );
TOMHTTreeAnimation( 'update', trk );

t = 0;

for k = 1:n

  % Get the measurements - zScan.data
  z = reshape( y(:,k), 2, d.nBalls );
  zScan = AddScan( z(:,1) );
  for j = 2:size(z,2)
    zScan = AddScan( z(:,j),[],zScan);
  end

  % Manage the tracks and generate hypotheses
  [b, trk, sol, hyp] = MHTTrackMgmt( b, trk, zScan, mhtData, k, t );

  % Update MHTGUI display
  if( ~isempty(zScan) && graphicsOn )
    if (treeAnimationOn)
      TOMHTTreeAnimation( 'update', trk );
    end
    MHTGUI(trk,sol,'hide');
    drawnow
  end

  t = t + dT;
end

% Show the final GUI
if (~treeAnimationOn)
  TOMHTTreeAnimation( 'update', trk );
end
if (~graphicsOn)
  MHTGUI(trk,sol,'hide');
end
MHTGUI;
```

The parameter pairs in `MHTInitialize` are described in Table 12.2.

Figure 12.7 shows the MHT GUI. This shows the GUI at the end of the simulation. The table shows scans on the x-axis and tracks on the y-axis (vertical). Each track is numbered as

***Table 12.2:*** Multiple Hypothesis Testing Parameters

| **Term** | **Definition** |
|---|---|
| `'probability false alarm'` | The probability that a measurement is spurious |
| `'probability of signal if target present'` | The probability of getting a signal if the target is present |
| `'probability of signal if target absent'` | The probability of getting a signal if the target is absent |
| `'probability of detection'` | Probability of detection of a target |
| `'measurement volume'` | Scales the likelihood ratio. |
| `'number of scans'` | The number of scans to consider in hypothesis formulation |
| `'gate'` | The size of the gate |
| `'m best'` | Number of hypotheses to consider |
| `'number of tracks'` | Number of tracks to maintain |
| `'scan to track function'` | Pointer to the scan to track function. This is custom for each application. |
| `'scan to track data'` | Data for the scan to track function |
| `'distance function'` | Pointer for the MHT distance function. Different definitions are possible. |
| `'hypothesis scan last'` | The last scan used in a hypothesis |
| `'prune tracks'` | Prune tracks if true |
| `'filter type'` | Type of Kalman Filter |
| `'filter data'` | Data for the Kalman Filter |
| `'remove duplicate tracks across all trees'` | If true removed duplicate tracks from all trees |
| `'average score history weight'` | A number to multiply the average score history |
| `'create track''` | If entered it will create a track instead of using an existing track |

xxx.yyy, where xxx is the track and yyy is the tag. Every track is assigned a new tag number. For example, 95.542 is track 95 and tag 542 means it is the 542nd track generated. The numbers in the table show the measurements associated with the track and the scan. TRK 3.21 and TRK 3.57 are duplicates. In both cases one measurement per scan is associated with the TRK. Their scores are the same because they are consistent. We can only pick one or the other for our hypothesis. TRK 95.542 doesn't get a measurement from scan 77, but for the rest of the scans it gets measurement 2. Scans 77 through 80 are active. A scan is a set of four position measurements. The summary shows there are seven active tracks, but we know (although the software does not necessarily) that there are only four balls in play. The number of scans are the ones currently in use to determine valid tracks. There are two active hypotheses.

Figure 12.8 shows the decision tree. You can see that with scan 80, two new tracks are created. This means that MHT thinks that there could be as many as four tracks. However, at

***Figure 12.7:*** The multiple hypothesis testing (MHT) graphic user interface (GUI).

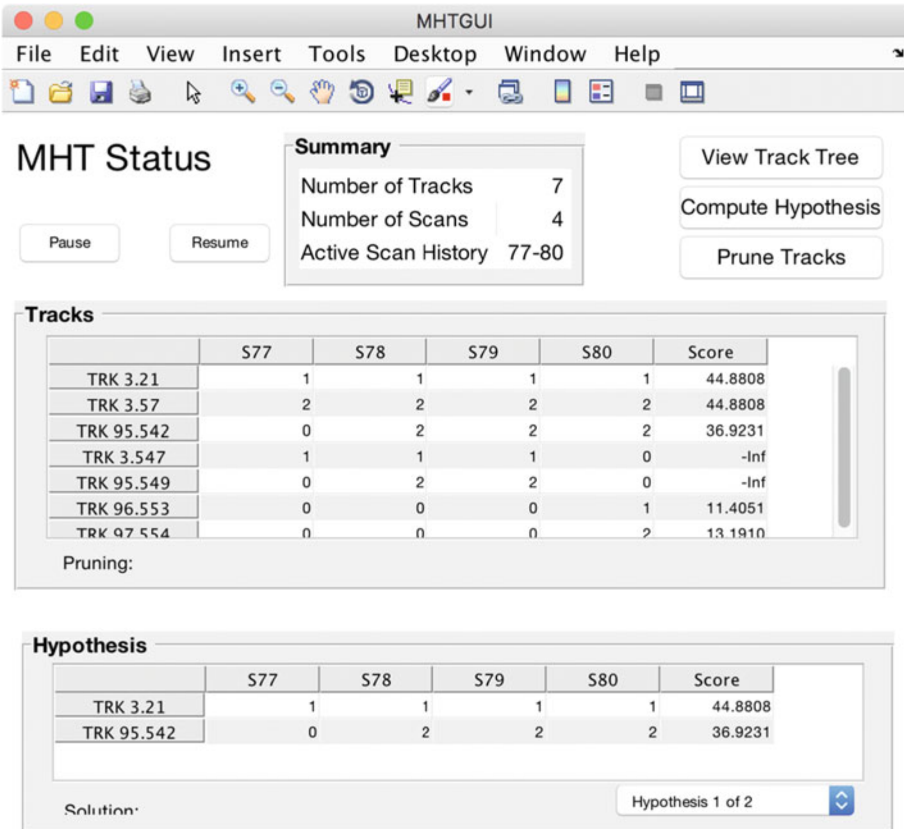


this point only two tracks, 3 and 95, have multiple measurements associated with them.

Figure 12.9 shows the information window. This shows the MHT algorithm's thinking. It gives the decisions made with each scan.

The demo shows that the MHT algorithm correctly associates measurements with tracks.

***Figure 12.8:*** The MHT tree. The blue bars gives the score assigned to each track. Longer is better. The numbers in the framed black boxes are the track numbers.

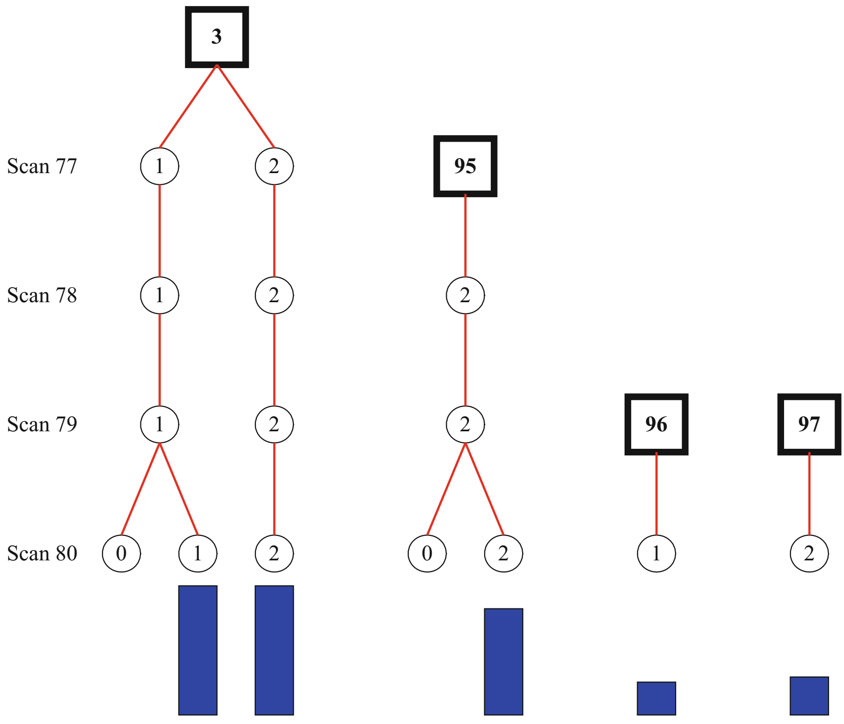


***Figure 12.9:*** The MHT information window. It tells you what the MHT algorithm is thinking.

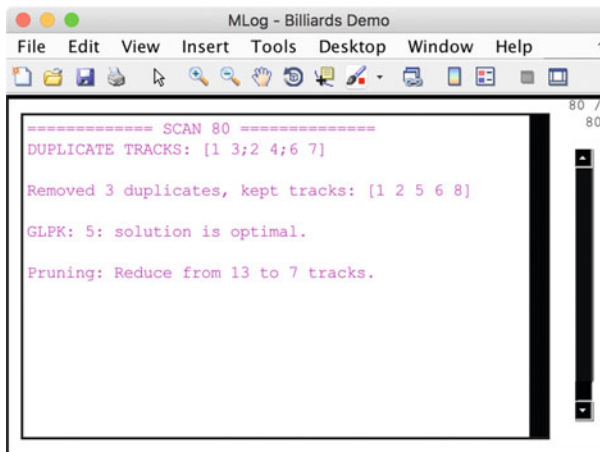

## 12.5 One-Dimensional Motion

### 12.5.1 Problem

You want to estimate the position of an object moving in one direction with unknown accelerations.

### 12.5.2 Solution

The solution is to create a linear Kalman Filter with an acceleration state.

### 12.5.3 How It Works

In this demo, we have a model of objects that includes an unknown acceleration state.

$$\begin{bmatrix} s \\ v \\ a \end{bmatrix}_{k+1} = \begin{bmatrix} 1 & \tau & \frac{1}{2}\tau^2 \\ 0 & 1 & \tau \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} s \\ v \\ a \end{bmatrix}_k \tag{12.9}$$

$$y_k = \begin{bmatrix} 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} s \\ v \\ a \end{bmatrix}_k \tag{12.10}$$

where $s$ is position, $v$ is velocity and $a$ is acceleration. $y_k = s$. $\tau$ is the time step. The input to the acceleration state is the time rate of change of acceleration.

The function DoubleIntegratorWithAccel creates the matrices shown above:

```
>> [a, b]  = DoubleIntegratorWithAccel( 0.5 )

a =
    1.0000    0.5000    0.1250
         0    1.0000    0.5000
         0         0    1.0000


b =
      0
      0
      1
```

with $\tau$ = 0.5 s.

We will set up the simulation so that one object has no acceleration, but starts in front of the other. The other will overtake the first. We want to see if MHT can sort out the trajectories. Passing would happen all the time with autonomous driving.

The following code implements the Kalman Filters for two vehicles. The simulation runs first to generate the measurements. The Kalman Filter runs next. Note that the plot array is updated after the filter update. This keeps it in sync with the simulation.

```
%% Run the Kalman Filter
% The covariances
r       = r(1,1);
q       = diag([0.5*aRand*dT^2;aRand*dT;aRand].^2 + q0);

% Create the Kalman Filter data structures
d1     = KFInitialize( 'kf', 'm', [0;0;0],  'x', [0;0;0], 'a', a, 'b',
    b, 'u',0,...
                      'h', h(1,1:3), 'p', diag(p0), 'q', q, 'r', r );
d2     = d1;
d1.m    = x(1:3,1) + sqrt(p0).*rand(3,1);
d2.m    = x(4:6,1) + sqrt(p0).*rand(3,1);
xE     = zeros(6,n);

for k = 1:n
  d1      = KFPredict( d1 );
  d1.y    = z(1,k);
  d1      = KFUpdate( d1 );

  d2      = KFPredict( d2 );
  d2.y    = z(2,k);
  d2      = KFUpdate( d2 );

  xE(:,k) = [d1.m;d2.m];
end
```

We use `PlotSet` with the argument `'plot set'` to group inputs and the argument `'legend'` to put legends on each plot. `'plot set'` takes a cell array of `1` $\times$ `n` arrays and `'legend'` takes a cell array of cell arrays as inputs. We don't need to numerically integrate the equations of motion because the state equations have already done that. You can always propagate a linear model in this fashion. We set the model noise matrix using `aRand`, but don't actually input any random accelerations. As written, our model is perfect, which is never true in a real system, hence the need for model uncertainty.

Figure 12.10 shows the states and the errors. The filters track all three states for both objects pretty well. The acceleration and velocity estimates converge with 10 s. It does a good job of estimating the fixed disturbance acceleration despite only having a position, $s$, measurement.

***Figure 12.10:*** The object states and filter errors.

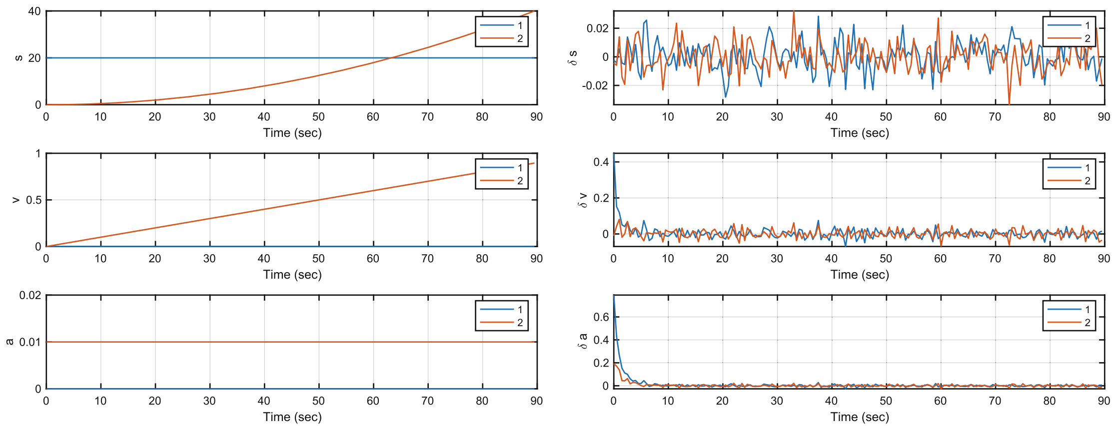


## 12.6 One-Dimensional Motion with Track Association

The next problem is one in which we need to associate measurements with a track.

### 12.6.1 Problem

You want to estimate the position of an object moving in one direction with measurements that need to be associated with a track.

### 12.6.2 Solution

The solution is to create an MHT system with the Kalman Filter as the state estimator.

### 12.6.3 How It Works

The MHT code is shown below. We append the MHT software to the script shown above. The Kalman Filters are embedded in the MHT software. We first run the simulation and gather the measurements and then process them in the MHT code.

```
% Initialize the MHT parameters
[mhtData, trk] = MHTInitialize( 'probability false alarm', 0.001,...
                                'probability of signal if target 
                                   present', 0.999,...
                                'probability of signal if target 
                                   absent', 0.001,...
                                'probability of detection', 1, ...
                                'measurement volume', 1.0, ...
                                'number of scans', 3, ...
                                'gate', 0.2,...
                                'm best', 2,...
                                'number of tracks', 1,...
                                'scan to track function',
                                   @ScanToTrack1D,...
```

```
                                'scan␣to␣track␣data',struct('v',0)
                                   ,...
                                'distance␣function',@MHTDistance,...
                                'hypothesis␣scan␣last', 0,...
                                'prune␣tracks', true,...
                                'filter␣type','kf',...
                                'filter␣data', f,...
                                'remove␣duplicate␣tracks␣across␣all␣
                                   trees',true,...
                                'average␣score␣history␣weight'
                                   ,0.01,...
                                'create␣track', '');

% Size arrays
m               = zeros(3,n);
p               = zeros(3,n);
scan            = cell(1,n);
b               = MHTTrkToB( trk );

TOMHTTreeAnimation( 'initialize', trk );
TOMHTTreeAnimation( 'update', trk );

% Initialize the MHT GUI
MHTGUI;
MLog('init')
MLog('name','MHT␣1D␣Demo')

t = 0;

for k = 1:n

  % Get the measurements
  zScan = AddScan( z(1,k) );
  zScan = AddScan( z(2,k), [], zScan );

  % Manage the tracks
  [b, trk, sol, hyp] = MHTTrackMgmt( b, trk, zScan, mhtData, k, t );

  % Update MHTGUI display
  MHTGUI(trk,sol,'update');

  % A guess for the initial velocity of any new track
  for j = 1:length(trk)
      mhtData.fScanToTrackData.v = mhtData.fScanToTrackData.v + trk(j
         ).m(1);
  end
```

```
  mhtData.fScanToTrackData.v = mhtData.fScanToTrackData.v/length(trk)
      ;

  % Animate the tree
  TOMHTTreeAnimation( 'update', trk );
  drawnow;
  t = t + dT;
end
```

Figure 12.11 shows the states and the errors. The MHT-hypothesized tracks are a good fit to the data.

***Figure 12.11:*** The MHT object states and estimated states. The colors are switched between plots.

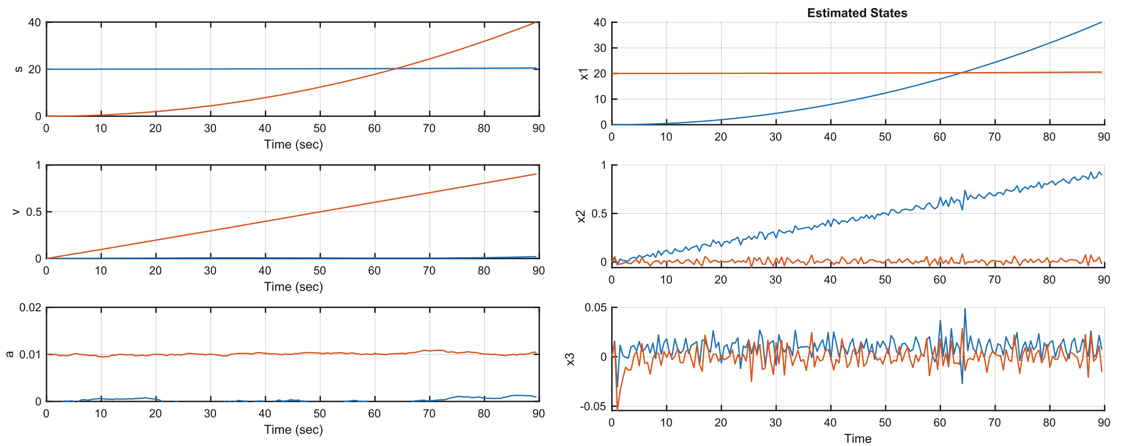


Figure 12.12 shows the MHT GUI and the tree. Track 1 contains only measurements from object 2. Track 2 contains only measurements from object 1. Three hundred and fifty-four and 360 are spurious tracks. Three hundred and fifty-four has a measurement 1 for scan 177, but none for the following scan. Three hundred and sixty was created on scan 180 and has just one measurement. One and 2 have the same score. The results show that the MHT software has successfully sorted out the measurements and assigned them correctly. At this point, the end of the sim, four scans are active.

## 12.7 Summary

This chapter has demonstrated the fundamentals of multiple hypothesis testing. Table 12.3 lists the functions and scripts included in the companion code.

*Figure 12.12:* The GUI and MHT tree. The tree shows the MHT decision process.

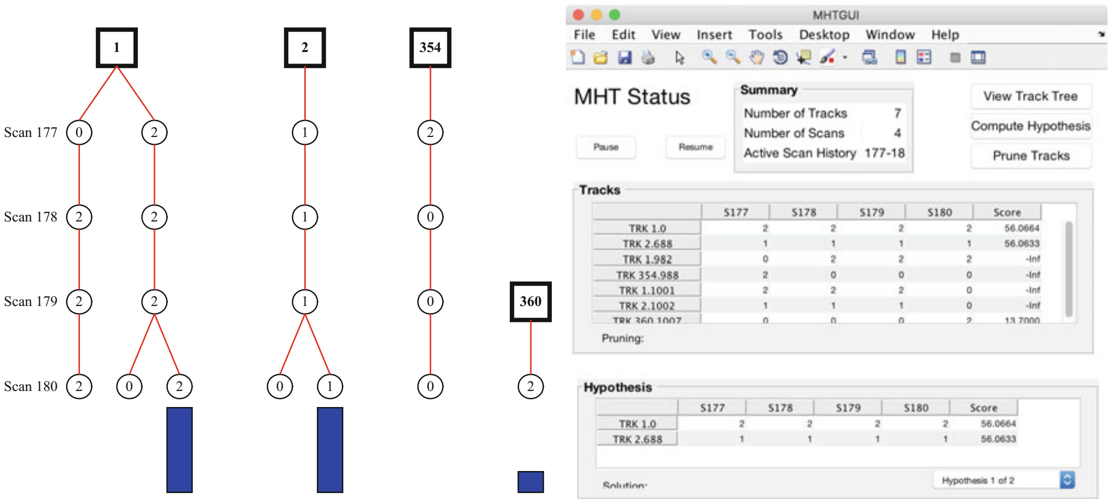


*Table 12.3:* Chapter Code Listing

| File | Description |
|---|---|
| AddScan | Add a scan to the data. |
| CheckForDuplicateTracks | Look through the recorded tracks for duplicates. |
| MHTDistanceUKF | Compute the MHT distance. |
| MHTGUI.fig | Saved layout data for the MHT GUI. |
| MHTGUI | GUI for the MHT software. |
| MHTHypothesisDisplay | Display hypotheses in a GUI. |
| MHTInitialize | Initialize the MHT algorithm. |
| MHTInitializeTrk | Initialize a track. |
| MHTLLRUpdate | Update the log likelihood ratio. |
| MHTMatrixSortRows | Sort rows in the MHT. |
| MHTMatrixTreeConvert | Convert to and from a tree format for the MHT data. |
| MHTTrackMerging | Merge MHT tracks. |
| MHTTrackMgmt | Manage MHT tracks. |
| MHTTrackScore | Compute the total score for the track. |
| MHTTrackScoreKinematic | Compute the kinematic portion of the track score. |
| MHTTrackScoreSignal | Compute the signal portion of the track score. |
| MHTTreeDiagram | Draw an MHT tree diagram. |
| MHTTrkToB | Convert tracks to a B matrix. |
| PlotTracks | Plot object tracks. |
| Residual | Compute the residual. |
| TOMHTTreeAnimation | Track-oriented MHT tree diagram animation. |
| TOMHTAssignment | Assign a scan to a track. |
| TOMHTPruneTracks | Prune the tracks. |

CHAPTER 13

# Autonomous Driving with Multiple Hypothesis Testing

In this chapter, we will apply the multiple hypothesis testing (MHT) techniques from the previous chapter to the interesting problem of autonomous driving. Consider a primary car that is driving along a highway at variable speeds. It carries a radar that measures azimuth, range, and range rate. Cars pass the primary car, some of which change lanes from behind the car and cut in front. The multiple hypothesis system tracks all cars around the primary car. At the start of the simulation there are no cars in the radar field of view. One car passes and cuts in front of the radar car. The other two just pass in their lanes. You want to accurately track all cars that your radar can see.

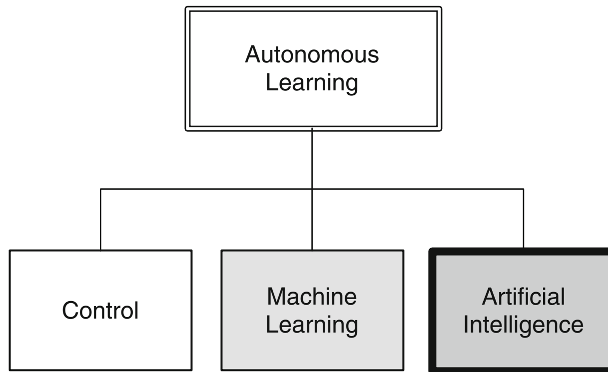


There are two elements to this problem. One is to model the motion of the tracked automobiles using measurements to improve your estimate of each automobile's location and velocity. The second is to systematically assign measurements to different tracks. A track should represent a single car, but the radar is just returning measurements on echoes, it doesn't know anything about the source of the echoes.

You will solve the problem by first implementing a Kalman Filter to track one automobile. We need to write measurement and dynamics functions that will be passed to the Kalman filter, and we need a simulation to create the measurements. Then we will apply the MHT techniques developed in the previous chapter to this problem.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_13

We'll do the following things in this chapter.

1. Model the automobile dynamics
2. Model the radar system
3. Write the control algorithms
4. Implement visualization to let us see the maneuvers in 3D
5. Implement the Unscented Kalman Filter
6. Implement MHT

## 13.1 Automobile Dynamics

### 13.1.1 Problem

We need to model the car dynamics. We will limit this to a planar model in two dimensions. We are modeling the location of the car in $x/y$ and the angle of the wheels, which allows the car to change direction.

### 13.1.2 Solution

Write a right-hand side function that can be called `RungeKutta`.

### 13.1.3 How It Works

Much like with the radar, we will need two functions for the dynamics of the automobile. `RHSAutomobile` is used by the simulation. `RHSAutomobile` has the full dynamic model including the engine and steering model. Aerodynamic drag, rolling resistance and side force resistance (the car doesn't slide sideways without resistance) are modeled. `RHSAutomobile` handles multiple automobiles. An alternative would be to have a one-automobile function and call `RungeKutta` once for each automobile. The latter approach works in all cases, except when you want to model collisions. In many types of collisions two cars collide and then stick, effectively becoming a single car. A real tracking system would need to handle this situation. Each vehicle has six states. They are:

1. x position
2. y position
3. x velocity
4. y velocity
5. Angle about vertical
6. Angular rate about vertical

The velocity derivatives are driven by the forces and the angular rate derivative by the torques. The planar dynamics model is illustrated in Figure 13.1 [29]. Unlike the reference, we constrain the rear wheels to be fixed and the angles for the front wheels to be the same.

***Figure 13.1:*** Planar automobile dynamical model.

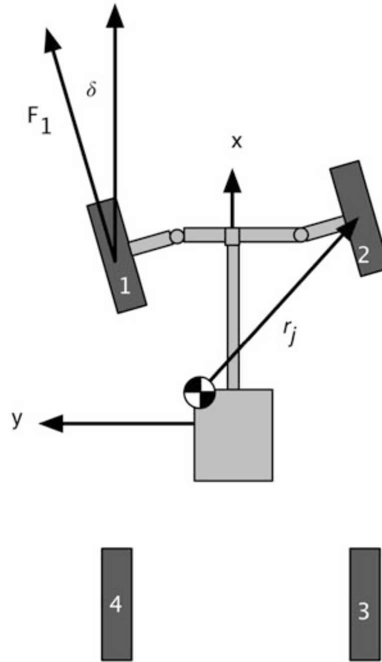


The dynamical equations are written in the rotating frame.

$$m(\dot{v}_x - 2\omega v_y) = \sum_{k=1}^{4} F_{k_x} - qC_{D_x} A_x u_x \tag{13.1}$$

$$m(\dot{v}_y + 2\omega v_x) = \sum_{k=1}^{4} F_{k_y} - qC_{D_y} A_y u_y \tag{13.2}$$

$$I\dot{\omega} = \sum_{k=1}^{4} r_k^{\times} F_k \tag{13.3}$$

where the dynamic pressure is:

$$q = \frac{1}{2}\rho\sqrt{v_x^2 + v_y^2} \tag{13.4}$$

and

$$v = \begin{bmatrix} v_x \\ v_y \end{bmatrix} \tag{13.5}$$

The unit vector is:

$$u = \frac{\begin{bmatrix} v_x \\ v_y \end{bmatrix}}{\sqrt{v_x^2 + v_y^2}} \tag{13.6}$$

The normal force is $mg$, where $g$ is the acceleration of gravity. The force at the tire contact point, where the tire touches the road, for tire $k$ is:

$$F_{t_k} = \begin{bmatrix} T/\rho - F_r \\ -F_c \end{bmatrix} \tag{13.7}$$

***Figure 13.2:*** Wheel force and torque.

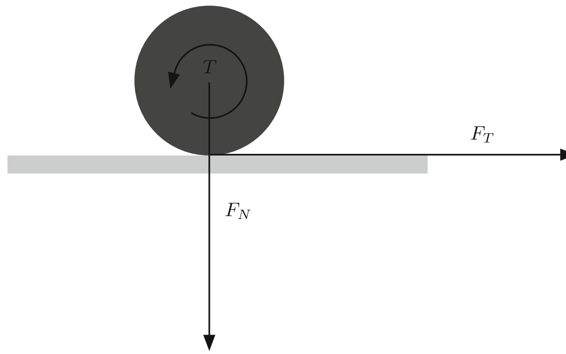


where $\rho$ is the radius of the tire and $F_r$ is the rolling friction and is:

$$F_r = f_0 + K_1 v_{t_x}^2 \tag{13.8}$$

where $v_{t_x}$ is the velocity in the tire frame in the rolling direction. For front wheel drive cars, the torque, $T$, is zero for the rear wheels. The contact friction is:

$$F_c = \mu_c mg \frac{v_{t_y}}{|v_t|} \tag{13.9}$$

This the force perpendicular to the normal rolling direction of the wheel, that is, into or out of the paper in Figure 13.2. The velocity term ensures that the friction force does not cause limit cycling. That is, when the $y$ velocity is zero, the force is zero. $\mu_c$ is a constant for the tires.

The transformation from tire to body frame is:

$$c = \begin{bmatrix} \cos\delta & -\sin\delta \\ \sin\delta & \cos\delta \end{bmatrix} \tag{13.10}$$

where $\delta$ is the steering angle so that:

$$F_k = cF_{t_k} \tag{13.11}$$

$$v_t = c^T \begin{bmatrix} v_x \\ v_y \end{bmatrix} \tag{13.12}$$

The kinematical equations that related yaw angle and yaw angular rate are

$$\dot{\theta} = \omega \tag{13.13}$$

and the inertial velocity $V$, the velocity needed to tell you where the car is going, is:

$$V = \begin{bmatrix} \cos\theta & -\sin\theta \\ \sin\theta & \cos\theta \end{bmatrix} v \tag{13.14}$$

We'll show you the dynamics simulation when we get to the graphics part of the chapter in Section 13.4

## 13.2 Modeling the Automobile Radar

### 13.2.1 Problem

The sensor utilized for this example will be the automobile radar. The radar measures azimuth, range, and range rate. We need two functions: one for the simulation and the second for use by the Unscented Kalman Filter.

### 13.2.2 Solution

Build a radar model in a MATLAB function. The function will use analytical derivations of range and range rate.

### 13.2.3 How It Works

The radar model is extremely simple. It assumes that the radar measures line-of-site range, range rate, and azimuth, the angle from the forward axis of the car. The model skips all the details of radar signal processing and outputs those three quantities. This type of simple model is always best when you start a project. Later on, you will need to add a very detailed model that has been verified against test data to demonstrate that your system works as expected.

The position and velocity of the radar are entered through the data structure. This does not model the signal-to-noise ratio of a radar. The power received by a radar goes as $\frac{1}{r^4}$. In this model, the signal goes to zero at the maximum range that is specified in the function. The range is found from the difference in position between the radar and the target. If $\delta$ is the difference, we write:

$$\delta = \begin{bmatrix} x - x_r \\ y - y_r \\ z - z_r \end{bmatrix} \tag{13.15}$$

Range is then:

$$\rho = \sqrt{\delta_x^2 + \delta_y^2 + \delta_z^2} \tag{13.16}$$

The delta velocity is:

$$\nu = \begin{bmatrix} v_x - v_{x_r} \\ v_y - v_{y_r} \\ v_z - v_{z_r} \end{bmatrix} \tag{13.17}$$

In both equations, the subscript $r$ denotes the radar. The range rate is:

$$\dot{\rho} = \frac{\nu^T \delta}{\rho} \tag{13.18}$$

The `AutoRadar` function handles multiple targets and can generate radar measurements for an entire trajectory. This is really convenient because you can give it your trajectory and see what it returns. This gives you a physical feel for the problem without running a simulation. It also allows you to be sure the sensor model is doing what you expect! This is important because all models have assumptions and limitations. It may be that the model really isn't suitable for your application. For example, this model is two-dimensional. If you are concerned about your system getting confused about a car driving across a bridge above your automobile, this model will not be useful in testing that scenario.

Notice that the function has a built-in demo and, if there are no outputs, will plot the results. Adding demos to your code is a nice way to make your functions more user friendly to other people using your code and even to you when you encounter the code again several months after writing the code! We put the demo in a subfunction because it is long. If the demo is one or two lines, a subfunction isn't necessary. Just before the demo function is the function defining the data structure.

The second function, `AutoRadarUKF` is the same core code, but designed to be compatible with the Unscented Kalman Filter. We could have used `AutoRadar`, but this is more convenient. The transformation matrix, `cITOC` (inertial to car transformation) is two-dimensional, since the simulation is in a flat world.

```
s        = sin(d.theta);
c        = cos(d.theta);
cIToC    = [c s;-s c];
dR       = cIToC*x(1:2);
dV       = cIToC*x(3:4);

rng      = sqrt(dR'*dR);
y        = [rng; dR'*dV/rng; atan(dR(2)/dR(1))];
```

The radar returns range, range rate, and the azimuth angle of the target. Even though we are using radar as our sensor, there is no reason why you couldn't use a camera, laser rangefinder or sonar instead. The limitation in the algorithms and software provided in this book is that it will only handle one sensor. You can get software from Princeton Satellite Systems that

expands this to multiple sensors. For example, cars carry radar, cameras, and lidar. You may want to integrate all of their measurements together. Figure 13.3 shows the internal radar demo. The target car is weaving in front of the radar. It is receding at a steady velocity, but the weave introduces a time-varying range rate.

***Figure 13.3:*** Built-in radar demo. The target is weaving in front of the radar.

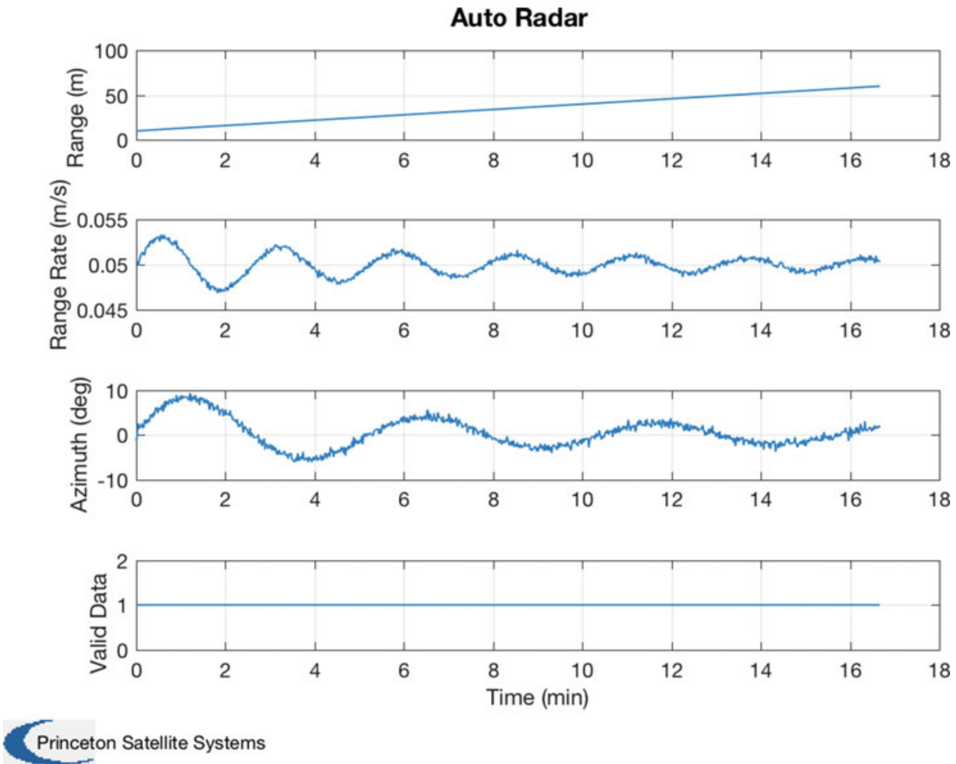


# 13.3 Automobile Autonomous Passing Control

## 13.3.1 Problem

To have something interesting for our radar to measure, we need our cars to perform some maneuvers. We will develop an algorithm for a car to change lanes.

## 13.3.2 Solution

The cars are driven by steering controllers that execute basic automobile maneuvers. Throttle (accelerator pedal) and steering angle can be controlled. Multiple maneuvers can be chained together. This provides a challenging test for the MHT system. The first function is for autonomous passing and the second performs the lane change.

### 13.3.3 How It Works

The `AutomobilePassing` implements passing control by pointing the wheels at the target. It generates a steering angle demand and torque demand. Demand is what we want the steering to do. In a real automobile, the hardware will attempt to meet the demand, but there will be a time lag before the wheel angle or motor torque meets the wheel angle or torque demand commanded by the controller. In many cases, you are passing the demand to another control system that will try and meet the demand. The algorithms are quite simple. They don't care if anyone gets in the way. They also don't have any control for avoiding another vehicle. The code assumes that the lane is empty. Don't try this with your car! The state is defined by the `passState` variable. Prior to passing, the `passState` is 0. During the passing, it is 1. When it returns to its original lane, the state is set to 0.

```
% Lead the target unless the passing car is in front
if( passee.x(1) + dX > passer.x(1) )
  xTarget = passee.x(1) + dX;
else
  xTarget = passer.x(1) + dX;
end

% This causes the passing car to cut in front of the car being passed
if( passer(1).passState == 0 )
  if( passer.x(1) > passee.x(1) + 2*dX )
    dY = 0;
    passer(1).passState = 1;
  end
else
  dY = 0;
end

% Control calculation
target          = [xTarget;passee.x(2) + dY];
theta           = passer.x(5);
dR              = target - passer.x(1:2);
angle           = atan2(dR(2),dR(1));
err             = angle - theta;
passer.delta    = gain(1)*(err + gain(3)*(err - passer.errOld));
passer.errOld   = err;
passer.torque   = gain(2)*(passee.x(3) + dV - passer.x(3));
```

The second function performs a lane change. It implements lane change control by pointing the wheels at the target. The function generates a steering angle demand and a torque demand. The default gains work reasonably well. You should always supply defaults that make sense.

```
% Default gains
if( nargin < 5 )
  gain = [0.05 80 120];
end
```

```
% Lead the target unless the passing car is in front
xTarget          = passer.x(1) + dX;

% Control calculation
target           = [xTarget;y];
theta            = passer.x(5);
dR               = target - passer.x(1:2);
angle            = atan2(dR(2),dR(1));
err              = angle - theta;
passer.delta     = gain(1)*(err + gain(3)*(err - passer.errOld));
passer.errOld    = err;
passer.torque    = gain(2)*(v - passer.x(3));
```

## 13.4 Automobile Animation

### 13.4.1 Problem

We want to visualize the cars as the maneuver.

### 13.4.2 How It Works

We create a function to read in `.obj` files. We then write a function to draw and animate the model.

### 13.4.3 Solution

The first part is to find an automobile model. A good resource is TurboSquid https://www.turbosquid.com. You will find thousands of models. We need `.obj` format and prefer a low polygon count. Ideally, we want models with triangles. In the case of the model found for this chapter, it had rectangles so we converted them to triangles using a Macintosh application, Cheetah3D https://www.cheetah3d.com. An OBJ model comes with an .obj file, an .mtl file (material file), and images for textures. We will only use the .obj file.

`LoadOBJ` loads the file and puts it into a data structure. The data structure uses the `g` field of the OBJ file to break the file into components. In this case, the components are the four tires and the rest of the car. The demo is just `LoadOBJ( 'MyCar.obj' )`. You do need the extension, `.obj`. The car is shown in Figure 13.4.

The image is generated with one call to `patch` per component.

The first part of `DrawComponents` initializes the model and updates it. We save, and return, pointers to the patches so that we only have to update the vectors with each call.

***Figure 13.4:*** Automobile 3D model.

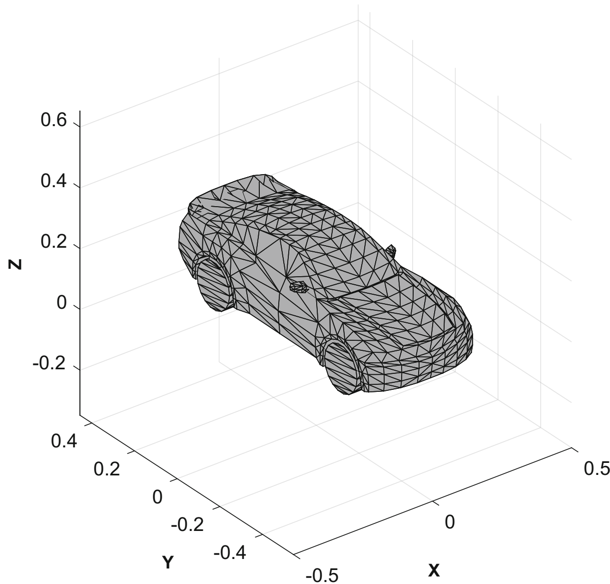


```
switch( lower(action) )
  case 'initialize'

    n = length(g.component);
    h = zeros(1,n);

    for k = 1:n
      h(k) = DrawMesh(g.component(k) );
    end

  case 'update'
    UpdateMesh(h,g.component,x);

  otherwise
    warning('%s␣not␣available',action);
end
```

The mesh is drawn with a call to `patch`. `patch` has many options that are worth exploring. We use the minimal set. We make the edges black to make the model easier to see. The Phong reflection model is an empirical lighting model. It includes diffuse and specular lighting.

```
function h = DrawMesh( m )

h = patch(  'Vertices', m.v, 'Faces',   m.f, 'FaceColor', m.color,...
            'EdgeColor',[0 0 0],'EdgeLighting', 'phong',...
            'FaceLighting', 'phong');
```

Updating is done by rotating the vertices around the z-axis and then adding the x and y positional offsets. The input array is [x;y;yaw]. We then set the new vertices. The function can handle an array of positions, velocities, and yaw angles.

```
function UpdateMesh( h, c, x )

for j = 1:size(x,2)
  for k = 1:length(c)
    cs      = cos(x(3,j));
    sn      = sin(x(3,j));
    b       = [cs -sn 0 ;sn cs 0;0 0 1];
    v       = (b*c(k).v')';
    v(:,1)  = v(:,1) + x(1,j);
    v(:,2)  = v(:,2) + x(2,j);
    set(h(k),'vertices',v);
  end
end
```

The graphics demo `AutomobileDemo` implements passing control. `AutomobileInitialize` reads in the OBJ file. The following code sets up the graphics window:

```
% Set up the figure
NewFig( 'Car␣Passing' )
axes('DataAspectRatio',[1 1 1],'PlotBoxAspectRatio',[1 1 1] );

h(1,:) = DrawComponents( 'initialize', d.car(1).g );
h(2,:) = DrawComponents( 'initialize', d.car(2).g );

XLabelS('X␣(m)')
YLabelS('Y␣(m)')
ZLabelS('Z␣(m)')

set(gca,'ylim',[-4 4],'zlim',[0 2]);

grid on
view(3)
rotate3d on
```

During each pass through the simulation loop, we update the graphics. We call `DrawComponents` once per car along with the stored `patch` handles for each car's components. We adjust the limits so that we maintain a tight focus on the two cars. We could have used the camera fields in the axes data structure for this too. We call `drawnow` after setting the new `xlim` for smooth animation.

```
% Draw the cars
pos1 = x([1 2]);
pos2 = x([7 8]);
DrawComponents( 'update', d.car(1).g, h(1,:), [pos1;pi/2 + x( 5)] );
DrawComponents( 'update', d.car(2).g, h(2,:), [pos2;pi/2 + x(11)] );

xlim = [min(x([1 7]))-10 max(x([1 7]))+10];
set(gca,'xlim',xlim);
drawnow
```

Figure 13.5 shows four points in the passing sequence.

***Figure 13.5:*** Automobile simulation snap shots showing passing.

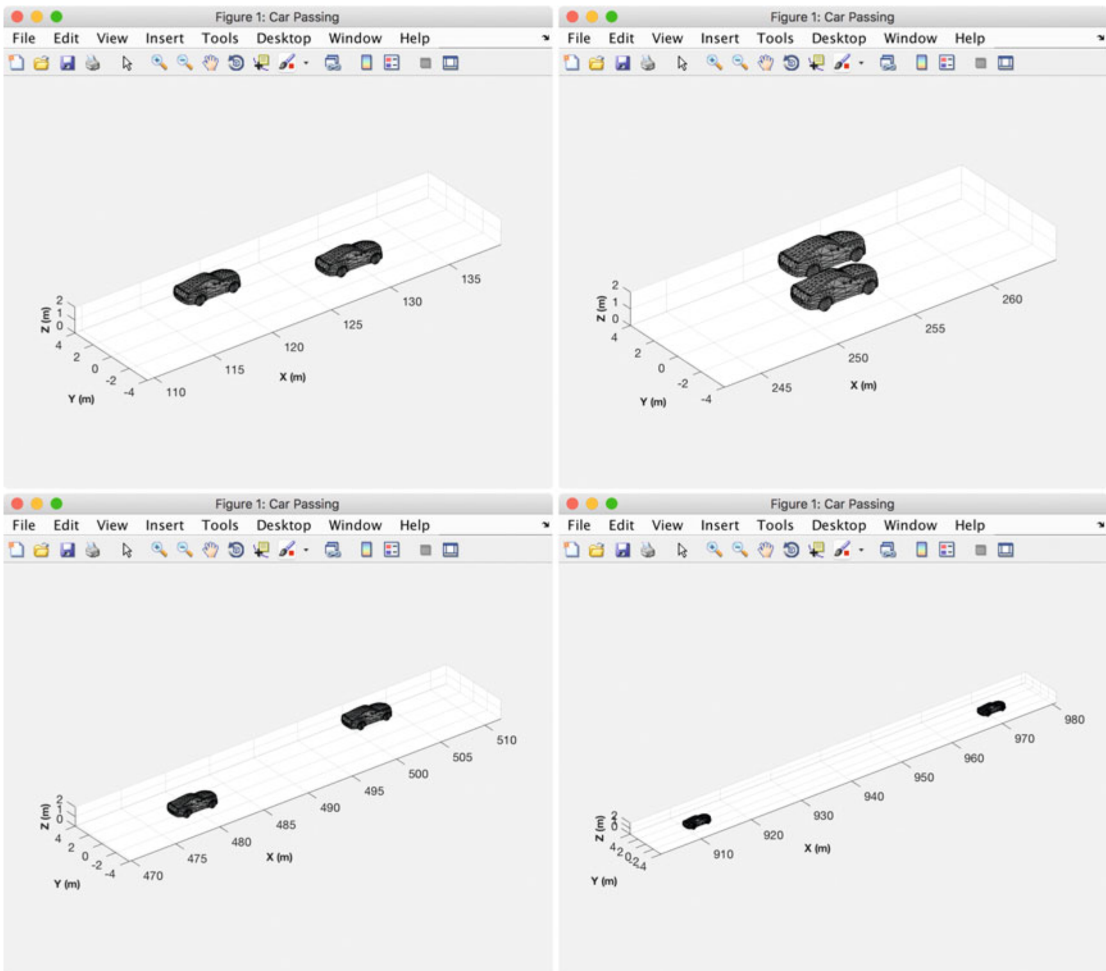

# 13.5 Automobile Simulation and the Kalman Filter

## 13.5.1 Problem

You want to track a car using radar measurements to track an automobile maneuvering around your car. Cars may appear and disappear at any time. The radar measurement needs to be turned into the position and velocity of the tracked car. In between radar measurements you want to make your best estimate of where the automobile will be at a given time.

## 13.5.2 Solution

The solution is to implement an Unscented Kalman Filter to take radar measurements and update a dynamical model of the tracked automobile.

## 13.5.3 How It Works

We first create the function `RHSAutomobileXY` with the Kalman Filter dynamical model. The Kalman Filter right-hand side is just the differential equations.

$$\dot{x} = v_x \tag{13.19}$$

$$\dot{y} = v_y \tag{13.20}$$

$$\dot{v}_x = 0 \tag{13.21}$$

$$\dot{v}_y = 0 \tag{13.22}$$

The dot means time derivative or rate of change with time. These are the state equations for the automobile. This model says that the position change with time is proportional to the velocity. It also says the velocity is constant. Information about velocity changes will come solely from the measurements. We also don't model the angle or angular rate. This is because we aren't getting information about it from the radar. However, you may want to try including it!

The `RHSAutomobileXY` function is shown below; it is of only one code! This is because it just models the dynamics of the point mass.

```
xDot = [x(3:4);0;0];
```

The demonstration simulation is the same simulation used to demonstrate the multiple hypothesis system tracking. This simulation just demonstrates the Kalman Filter. Since the Kalman Filter is the core of the package, it is important that it works well before adding the measurement assignment part.

`MHTDistanceUKF` finds the MHT distance for use in gating computations using the Unscented Kalman Filter (UKF). The MHT distance is the distance between the observation and predicted locations. The measurement function is of the form h(x,d), where d is the UKF data structure. `MHTDistanceUKF` uses sigma points. The code is similar to `UKFUpdate`. As the uncertainty gets smaller, the residual must be smaller to remain within the gate.

```
pS       = d.c*chol(d.p)';
nS       = length(d.m);
nSig     = 2*nS + 1;
mM       = repmat(d.m,1,nSig);
if( length(d.m) == 1 )
    mM = mM';
end

x        = mM + [zeros(nS,1) pS -pS];

[y, r]   = Measurement( x, d );
mu       = y*d.wM;
b        = y*d.w*y' + r;
del      = d.y - mu;
k        = del'*(b\del);

%% MHTDistanceUKF>Measurement
function [y, r] = Measurement( x, d )
%       Measurement from the sigma points

nSigma   = size(x,2);
lR       = length(d.r);
y        = zeros(lR,nSigma);
r        = d.r;
iR       = 1:lR;

for j = 1:nSigma
        f           = feval( d.hFun, x(:,j), d.hData );
        y(iR,j)     = f;
        r(iR,iR)    = d.r;
end
```

The simulation `UKFAutomobileDemo` uses a car data structure to contain all of the car information. A MATLAB function `AutomobileInitialize` takes parameter pairs and builds the data structure. This is a lot cleaner than assigning the individual fields in your script. It will return a default data structure if nothing is entered as an argument.

The first part of the demo, is the automobile simulation. It generates the measurements of the automobile positions to be used by the Kalman Filter. The second part of the demo processes the measurements in the UKF to generate the estimates of the automobile track. You could move the code that generates the simulated data into a separate file if you were reusing the simulation results repeatedly.

The results of the script are shown in Figure 13.6, Figure 13.7, and Figure 13.8.

***Figure 13.6:*** Automobile trajectories.

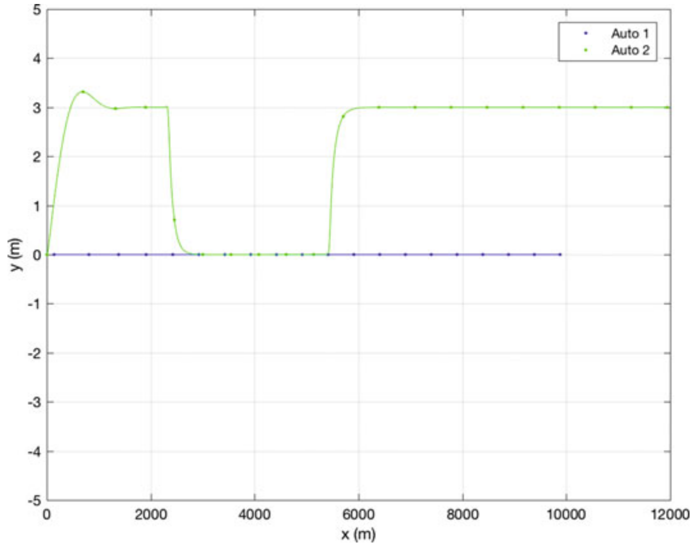


***Figure 13.7:*** The true states and Unscented Kalman Filter (UKF) estimated states.

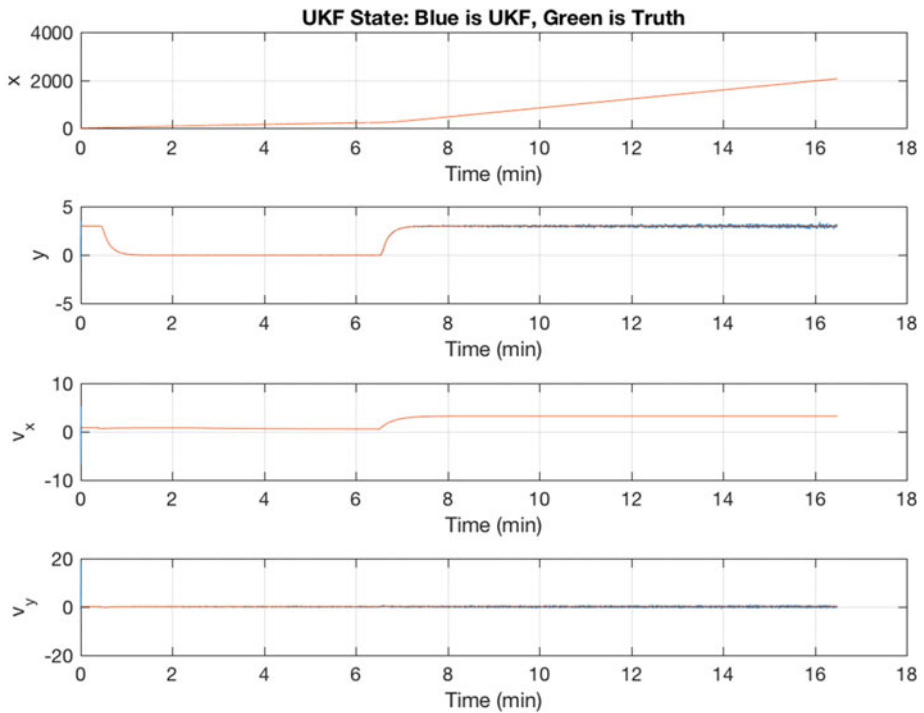

***Figure 13.8:*** The MHT distance between the automobiles during the simulation. Notice the spike in distance when the automobile maneuver starts.

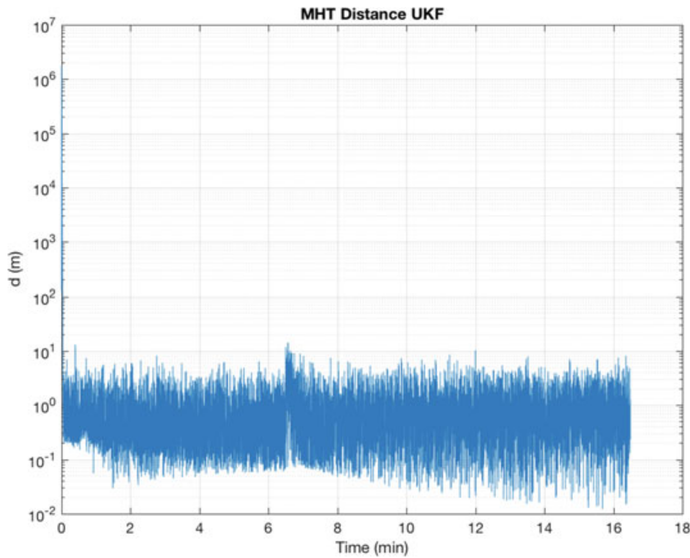


# 13.6 Automobile Target Tracking

## 13.6.1 Problem

We need to demonstrate target tracking for automobiles.

## 13.6.2 Solution

Build an automobile simulation with target tracking.

## 13.6.3 How It Works

The simulation is for a two-dimensional model of automobile dynamics. The primary car is driving along a highway at variable speeds. It carries a radar. Many cars pass the primary car, some of which change lanes from behind the car and cut in front. The MHT system tracks all cars. At the start of the simulation there are no cars in the radar field of view. One car passes and cuts in front of the radar car. The other two just pass in their lanes. This is a good test of track initiation.

The radar, covered in the first recipe of the chapter, measures range, range rate, and azimuth in the radar car frame. The model generates those values directly from the target and from the tracked cars' relative velocity and positions. The radar signal processing is not modeled, but the radar has field-of-view and range limitations. See `AutoRadar`.

The cars are driven by steering controllers that execute automobile maneuvers. Throttle (accelerator pedal) and steering angle can be controlled. Multiple maneuvers can be chained together. This provides a challenging test for the MHT system. You can try different maneuvers and add additional maneuver functions of your own.

The Unscented Kalman Filter described in Chapter 4 is used in this demo as the radar is a highly nonlinear measurement. The UKF dynamical model, `RHSAutomobileXY`, is a pair of double integrators in the inertial frame relative to the radar car. The model accommodates steering and throttle changes by making the plant covariance, both position and velocity, larger than would be expected by analyzing the relative accelerations. An alternative would be to use interactive multiple models (IMMs) with a "steering" model and "acceleration" model. This added complication does not appear to be necessary. A considerable amount of uncertainty would be retained even with IMM, since a steering model would be limited to one or two steering angles. The script implementing the simulation with MHT is `MHTAutomobileDemo`. There are four cars in the demo; car 4 will be passing. Figure 13.10 shows the radar measurement for car 3, which is the last car tracked. The MHT system handles vehicle acquisition well. The MHT GUI in Figure 13.11 shows a hypothesis with three tracks at the end of the simulation. This is the expected result.

***Figure 13.9:*** Automobile demo car trajectories.

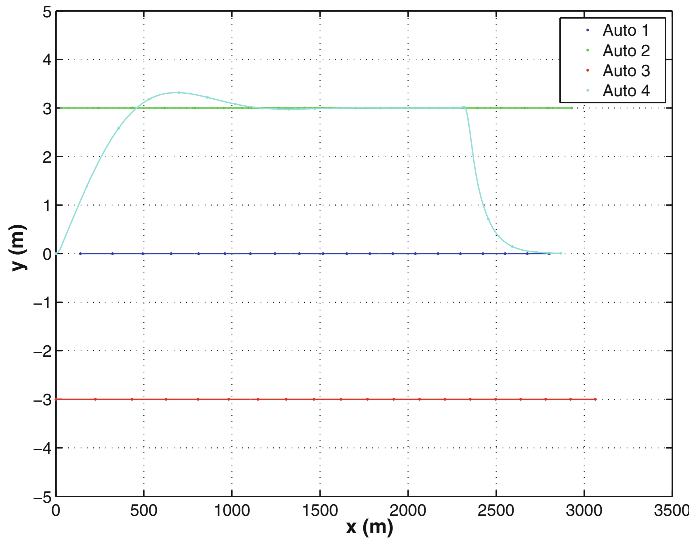


Figure 13.12 shows the final tree. There are several redundant tracks. These tracks can be removed, since they are clones of other tracks. This does not impact hypothesis generation.

***Figure 13.10:*** Automobile demo radar measurement for car 3.

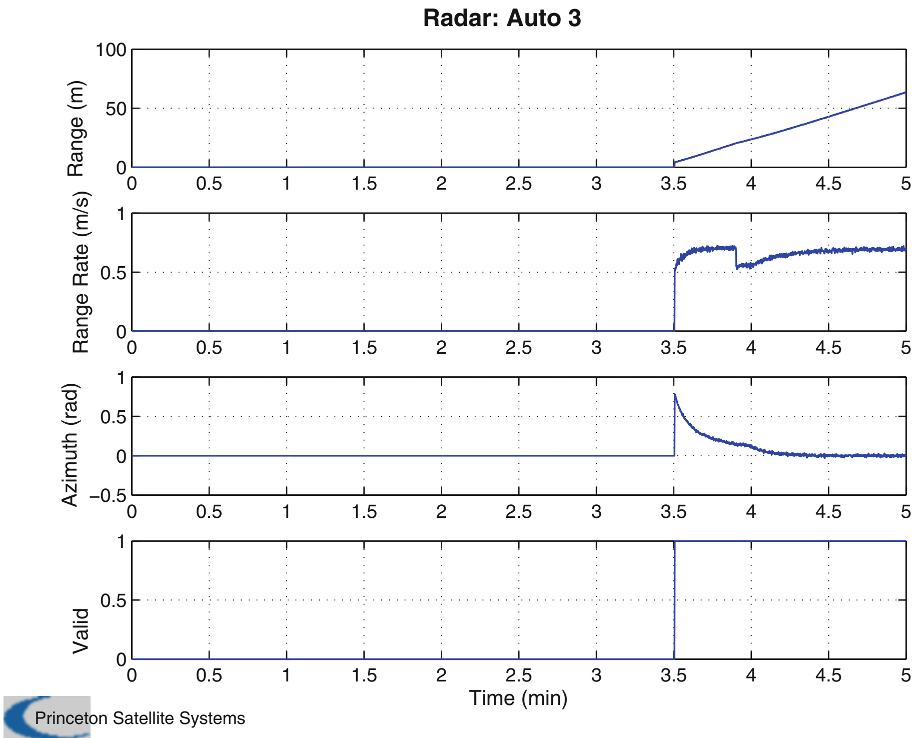


***Figure 13.11:*** The MHT GUI shows three tracks. Each track has consistent measurements.

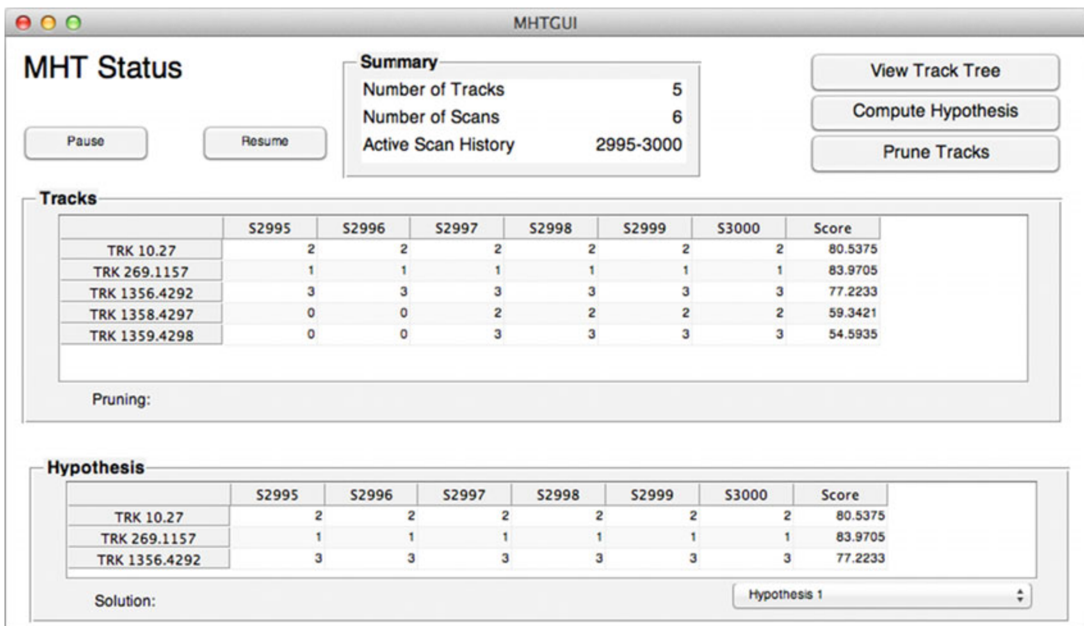

***Figure 13.12:*** The final tree for the automobile demo.

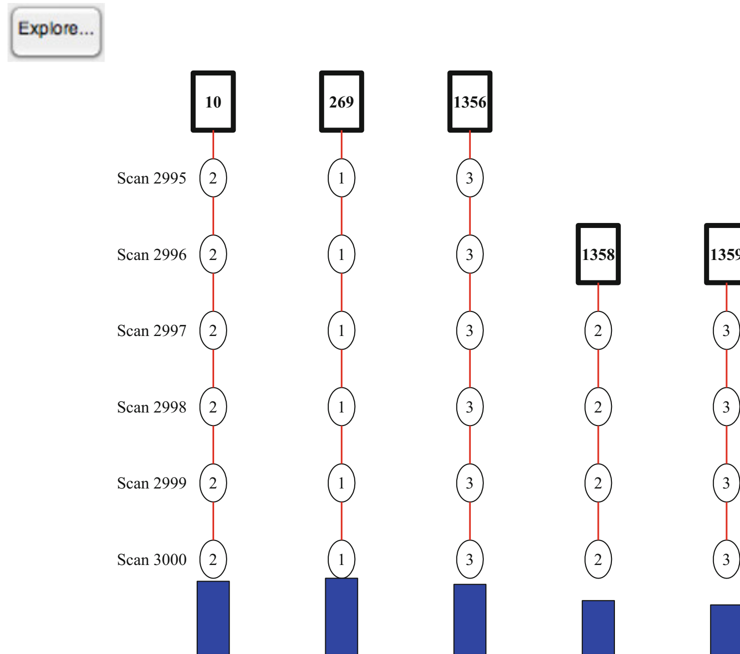


## 13.7 Summary

This chapter has demonstrated an automobile tracking problem. The automobile has a radar system that detects cars in its field of view. The system accurately assigns measurements to tracks and successfully learns the path of each neighboring car. You started by building an Unscented Kalman Filter to model the motion of an automobile and to incorporate measurements from a radar system. This is demonstrated in a simulated script. You then build a script that incorporates track-oriented multiple hypothesis testing to assign measurements taken by the radar of multiple automobiles. This allows our radar system to autonomously and reliably track multiple cars.

You also learned how to make simple automobile controllers. The two controllers steer the automobiles and allow them to pass other cars.

Table 13.1 lists the functions and scripts included in the companion code.

***Table 13.1:*** Chapter Code Listing

| File | Description |
|---|---|
| AutoRadar | Automobile radar model for simulation. |
| AutoRadarUKF | Automobile radar model for the UKF. |
| AutomobileDemo | Demonstrate automobile animation. |
| AutomobileInitialize | Initialize the automobile data structure. |
| AutomobileLaneChange | Automobile control algorithm for lane changes. |
| AutomobilePassing | Automobile control algorithm for passing. |
| DrawComponents | Draw a 3D model. |
| LoadOBJ | Load an .obj graphics file. |
| MHTAutomobileDemo | Demonstrate the use of multiple hypothesis testing for automobile radar systems. |
| RHSAutomobile | Automobile dynamical model for simulation. |
| RHSAutomobileXY | Automobile dynamical model for the UKF. |
| UKFAutomobileDemo | Demonstrate the UKF for an automobile. |

CHAPTER 14

# Case-Based Expert Systems

In this chapter, we will introduce case-based expert systems, an example of the Artificial Intelligence branch of our Autonomous Learning taxonomy. There are two broad classes of expert systems, rule-based and case-based. Rule-based systems have a set of rules that are applied to come to a decision; they are just a more organized way of writing decision statements in computer code. These systems provide a way of automating the process when decision-making involves hundreds or thousands of rules. Case-based systems decide by example, that is, a set of predefined cases.

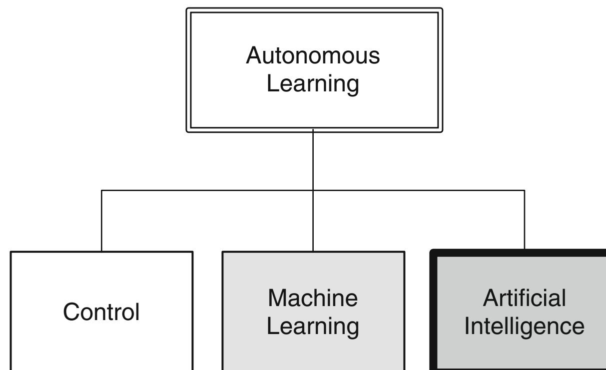


Learning in the context of an expert system depends strongly on the configuration of the expert system. There are three primary methods, which vary in the level of autonomy of learning and the average generalization of the new knowledge of the system.

The least autonomous method of learning is the introduction of new rule sets in simple rule-based expert systems. Learning of this sort can be highly tailored and focused, but is done entirely at the behest of external teachers. In general, quite specific rule-based systems with extremely general rules tend to have issues with edge cases that require exceptions to their rules. Thus, this type of learning, although easy to manage and implement, is neither autonomous nor generalizable.

The second method is fact-gathering. The expert system makes decisions based on the known cause and effect relationships, along with an evolving model of the world; learning,

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2_14

then is broken up into two sub-pieces. Learning new cause and effect system rules is very similar to the type of learning described above, requiring external instruction, but can be more generalizable (as it is combined with more general world knowledge than a simple rule-based system might have). Learning new facts, however, can be very autonomous and involves the refinement of the expert system's model of reality by increasing the amount of information that can be taken advantage of by the automated reasoning systems.

The third method is fully autonomous-based reasoning, where actions and their consequences are observed, leading to inferences about what prior and action combinations lead to what results. For instance, if two similar actions result in positive results, then those priors, which are the same in both cases, can begin to be inferred as necessary preconditions for a positive result from that action. As additional actions are seen, these inferences can be refined and confidence can increase in the predictions made.

The three methods are listed in increasing difficulty of implementation. Adding rules to a rule-based expert system is quite straightforward, although rule dependencies and priorities can become complicated. Fact-based knowledge expansion in automated reasoning systems is also fairly straightforward, once suitably generic sensing systems for handling incoming data are set up. The third method is by far the most difficult; however, rule-based systems can incorporate this type of learning. In addition, more general pattern recognition algorithms can be applied to training data (including on-line, unsupervised training data) to perform this function, learning to recognize, e.g., with a neural network, patterns of conditions that would lead to positive or negative results from a given candidate action. The system can then check possible actions against these learned classification systems to gauge the potential outcome of the candidate actions.

In this chapter, we will explore case-based reasoning systems. This is a collection of cases with their states and values given by strings. We do not address the problem of having databases with thousands of cases. The code we present would be too slow. We will not deal with a system that autonomously learns. However, the code in this chapter can be made to learn by feeding back the results of new cases into the case-based system.

## 14.1 Building Expert Systems

### 14.1.1 Problem

We want a tool to build a case-based expert system. Our tool needs to work for small sets of cases.

### 14.1.2 Solution

Build a function, `BuildExpertSystem`, that accepts parameter pairs to create the case-based expert system.

### 14.1.3 How It Works

The knowledge base consists of states, values, and production rules. There are four parts of a new case: the case name, the states and values, and the outcome. A state can have multiple values.

The state catalog is a list of all of the information that will be available to the reasoning system. It is formatted as states and state values. Only string values are permitted. Cell arrays store all the data.

The default catalog is shown below for reaction wheel control system. The cell array of acceptable or possible values for each state follows the state definition:

```
{
      {'wheel-turning'},      {'yes','no'};
      {'power'},              {'on','off'};
      {'torque-command'},     {'yes','no'}
}
```

Our database of cases is designed to detect failures. We have three things to check to see if the wheel is working. If the wheel is turning and power is on and there is a torque command, then it is working. The wheel can be turning without a torque command or with the power off because it would just be spinning down from prior commands. If the wheel is not turning, the possibilities are that there is no torque command or that the power is off.

## 14.2 Running an Expert System

### 14.2.1 Problem

We want to create a case-based expert system and run it.

### 14.2.2 Solution

Build an expert system engine that implements a case-based reasoning system. It should be designed to handle small numbers of cases and be capable of updating the case database.

### 14.2.3 How It Works

Once you have defined a few cases from your state catalog, you can test the system. The function `CBREngine` implements the case-based reasoning engine. The idea is to pass it a case, `newCase`, and see if it matches any existing cases stored in the system data structure. For our problem we think that we have all the cases necessary to detect any failure. We do string matching with a built-in function using `strcmpi`. We then find the first value that matches.

The algorithm finds the total fraction of the cases that match to determine if the example matches the stored cases. The engine is matching values for states in the new case against values for states in the case database. It weights the results by the number of states. If the new case has more states than an existing case, it biases the result by the number of states in the database case divided by the number of states in the new case. If more than one case matches the new case and the outcomes for the matching cases are different, the outcome is declared

"ambiguous". If they are the same, it gives the new case that outcome. The case names make it easier to understand the results. We use `strcmpi` to make string matches case insensitive.

```
function [outcome, pMatch] = CBREngine( newCase, system )

% Find the cases that most closely match the given state values
pMatch  = zeros(1,length(system.case));
pMatchF = length(newCase.state); % Number of states in the new case
for k = 1:length(system.case)
  f = min([1 length(system.case(k).activeStates)/pMatchF]);
  for j = 1:length(newCase.state)
    % Does state j match any active states?
    q = StringMatch( newCase.state(j), system.case(k).activeStates );
    if( ~isempty(q) )
      % See if our values match
      i = strcmpi(newCase.values{j},system.case(k).values{q});
      if( i )
        pMatch(k) = pMatch(k) + f/pMatchF;
      end
    end
  end
end

i = find(pMatch == 1);
if( isempty(i) )
  i = max(pMatch,1);
end

outcome = system.case(i(1)).outcome;

for k = 2:length(i)
  if( ~strcmp(system.case(i(k)).outcome,outcome))
    outcome = 'ambiguous';
  end
end
```

The demo script, `ExpertSystemDemo`, is quite simple. The first part builds the system. The remaining code runs some cases. `'id'` denotes the index of the following data in its cell array. For example, the first three entries are for the catalog and they are items 1 through 3. The next three are for cases and they are items 1 through 4. As `BuildExpertSystem` goes through the list of parameter pairs, it uses the last `id` as the index for subsequent parameter pairs.

```
system = BuildExpertSystem( [], 'id',1,...
                            'catalog␣state␣name','wheel-turning',...
                            'catalog␣value',{'yes','no'},...
                            'id',2,...
                            'catalog␣state␣name','power',...
```

```
                              'catalog␣value',{'on' 'off'},...
                              'id',3,...
                              'catalog␣state␣name','torque-command',...
                              'catalog␣value',{'yes','no'},...
                              'id',1,...
                              'case␣name', 'Wheel␣operating',...
                              'case␣states',{'wheel-turning', 'power',
                                 'torque-command'},...
                              'case␣values',{'yes' 'on' 'yes'},...
                              'case␣outcome','working',...
                              'id',2,...
                              'case␣name', 'Wheel␣power␣ambiguous',...
                              'case␣states',{'wheel-turning', 'power',
                                 'torque-command'},...
                              'case␣values',{'yes' {'on' 'off'} 'no'
                                 },...
                              'case␣outcome','working',...
                              'id',3,...
                              'case␣name', 'Wheel␣broken',...
                              'case␣states',{'wheel-turning', 'power',
                                 'torque-command'},...
                              'case␣values',{'no' 'on' 'yes'},...
                              'case␣outcome','broken',...
                              'id',4,...
                              'case␣name', 'Wheel␣turning',...
                              'case␣states',{'wheel-turning', 'power'
                                 },...
                              'case␣values',{'yes' 'on'},...
                              'case␣outcome','working',...
                              'match␣percent',80);

newCase.state  = {'wheel-turning', 'power', 'torque-command'};
newCase.values = {'yes','on','no'};
newCase.outcome = '';

[newCase.outcome, pMatch] = CBREngine( newCase, system );

fprintf(1,'New␣case␣outcome:␣%s\n\n',newCase.outcome);

fprintf(1,'Case␣ID␣Name␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣␣Percentage␣Match\n');
for k = 1:length(pMatch)
  fprintf(1,'Case␣%d:␣%-30s␣%4.0f\n',k,system.case(k).name,pMatch(k)
     *100);
end
```

As you can see, we match two cases, but because their outcome is the same the wheel is declared working. The wheel power ambiguous is called that because the power could be on or off, hence ambiguous. We could add this new case to the database using `BuildExpertSystem`.

We used `fprintf` in the script to print the following results into the command window.

```
>> ExpertSystemDemo
New case outcome: working

Case ID Name                                  Percentage Match
Case 1: Wheel working                           67
Case 2: Wheel power ambiguous                   67
Case 3: Wheel broken                            33
Case 4: Wheel turning                           44
```

This example is for a very small case-based expert system with a binary outcome. Multiple outcomes can be handled without any changes to the code. However, the matching process is slow, as it cycles through all the cases. A more robust system, handling thousands of cases, would need some kind of decision tree to cull the cases tested. For example, suppose we had several different components that we were testing. For example, with a landing gear we need to know that the tire is not flat, the brakes are working, the gear is deployed, and the gear is locked. If the gear is not deployed, we no longer have to test the brakes or the tires or that the gear is locked.

## 14.3 Summary

This chapter has demonstrated a simple case-based reasoning expert system. The system can be configured to add new cases based on the results of previous cases. An alternative would be a rule-based system. Table 14.1 lists the functions and scripts included in the companion code.

***Table 14.1:*** Chapter Code Listing

| File | Description |
|---|---|
| BuildExpertSystem | Function to build a case-based expert system database. |
| CBREngine | Case-based reasoning engine. |
| ExpertSystemDemo | Expert system demonstration. |

## APPENDIX A

# A Brief History of Autonomous Learning

## A.1 Introduction

In the first chapter of this book, you were introduced to autonomous learning. You saw that autonomous learning could be divided into the areas of machine learning, controls, and artificial intelligence (AI). In this appendix we will provide some background on how each area evolved. Automatic control predates artificial intelligence. However, we are interested in adaptive or learning control, which is a relatively new development and really began evolving around the time that artificial intelligence had its foundations. Machine learning is considered an offshoot of artificial intelligence. However, many of the methods used in machine learning came from different fields of study, such as statistics and optimization.

## A.2 Artificial Intelligence

Artificial Intelligence research began shortly after World War II [22]. Early work was based on knowledge of the structure of the brain, propositional logic, and Turing's theory of computation. Warren McCulloch and Walter Pitts created a mathematical formulation for neural networks based on threshold logic. This allowed neural network research to be split into two approaches. One centered on biological processes in the brain and the other on the application of neural networks to artificial intelligence. It was demonstrated that any function could be implemented through a set of such neurons and that a neural net could learn. In 1948, Weiner's book, "Cybernetics," was published, which described concepts in control, communications, and statistical signal processing. The next major step in neural networks was Hebb's book in 1949, "The Organization of Behavior," connecting connectivity with learning in the brain. His book became a source of learning and adaptive systems. Marvin Minsky and Dean Edmonds built the first neural computer in 1950.

In 1956, Allen Newell and Herbert Simon designed a reasoning program, the Logic Theorist (LT), which worked non-numerically. The first version was hand simulated using index cards. It could prove mathematical theorems and even improve on human derivations. It solved 38 of

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2

the 52 theorems in Principia Mathematica. LT employed a search tree with heuristics to limit the search. LT was implemented on a computer using IPL, a programming language that led to Lisp, a programming language that is discussed below.

Blocks World was one of the first attempts to demonstrate general computer reasoning. The blocks world was a micro world. A set of blocks would sit on a table, some sitting on other blocks. The AI systems could rearrange blocks in certain ways. Blocks under other blocks could not be moved until the block on top was moved. This is not unlike the Towers of Hanoi problem. Blocks World was a spectacular advancement as it showed that a machine could reason at least in a limited environment. Blocks World was an early example of the use of machine vision. The computer had to process an image of Blocks World and determine what was a block and where they were located.

Blocks World and Newell's and Simons' LT was followed up by the General Problem Solver (GPS). It was designed to imitate human problem solving methods. Within its limited class of puzzles it could solve them much like a human. Although GPS solved simple problems such as the Towers of Hanoi, Figure A-1, it could not solve real-world problems because the search was lost in a combinatorial explosion that represented the enumeration of all choices in a vast decision space.

***Figure A-1:*** Towers of Hanoi. The disks must be moved from the first peg to the last without ever putting a bigger diameter disk on top of a smaller diameter disk.

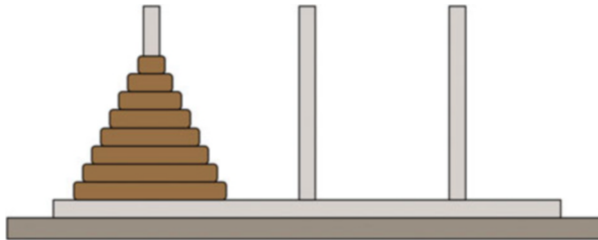

In 1959, Herman Gelernter wrote the Geometry Theorem prover, which could prove theorems that were quite tricky. The first game-playing programs were written at this time. In 1958, John McCarthy invented the language Lisp (LISt Processing), which was to become the main AI language. It is now available as Scheme and Common Lisp. Lisp was introduced only one year after FORTRAN. A typical Lisp expression is:

```
(defun sqrt-iter (guess x)
  (if (good-enough-p guess x)
      guess
      (sqrt-iter (improve guess x) x)))
```

This computes a square root through recursion. Eventually, dedicated Lisp machines were built, but they went out of favor when general purpose processors became faster.

Time sharing was invented at MIT to facilitate AI research. Professor McCarthy created a hypothetical computer program, Advice Taker, a complete AI system that could embody general world information. It would have used a formal language such as predicate calculus.

For example, it could come up with a route to the airport from simple rules. Marvin Minsky arrived at MIT in 1958 and began working on micro-worlds. Within these limited domains AI could solve problems, such as closed-form integrals in calculus. Minsky wrote the book "Perceptrons" (with Seymour Papert), which was fundamental in the analysis of artificial neural networks. The book contributed to the movement toward symbolic processing in AI. The book noted that single neurons could not implement some logical functions such as exclusive-or and erroneously implied that multi-layer networks would have the same issue. It was later found that three-layer networks could implement such functions.

■ **TIP** Three-layer networks are the minimum to solve most learning problems.

More challenging problems were tried in the 1960s. Limitations in the AI techniques became evident. The first language translation programs had mixed results. Trying to solve problems by working through massive numbers of possibilities (such as in chess) ran into computation problems. Human chess play has many forms. Some involve memorization of patterns, including openings where the board positions are well-defined and end games when the number of pieces is relatively low. Positional play involves seeing patterns on the board through the human brain's ability to process patterns. Someone who is a good positional player will arrange her pieces on the board so that the other player's options are restricted. Localized pattern recognition is seen in mate-in-n problems. Human approaches are not really used in computer chess. Computer chess programs have become very capable primarily because of faster processors and the ability to store openings and end games. Multi-layer neural networks were discovered in the 1960s, but not really studied until the 1980s.

In the 1970s, self-organizing maps using competitive learning were introduced [12]. A resurgence in neural networks happened in the 1980s. Knowledge-based systems were also introduced in the 1980s. From Jackson [14],

> An expert system is a computer program that represents and reasons with knowledge of some specialized subject with a view to solving problems or giving advice.

This included expert systems that could store massive amounts of domain knowledge. These could also incorporate uncertainty in their processing. Expert systems are applied to medical diagnoses and other problems. Unlike AI techniques up to this time, expert systems could deal with problems of realistic complexity and attain high performance. They also explain their reasoning. This last feature is critical in their operational use. Sometimes these are called knowledge-based systems. A well-known open source expert system is CLIPS (C Language Integrated Production System).

Back propagation for neural networks was re-invented in the 1980s leading to renewed progress in this field. Studies began both of human neural networks (i.e., the human brain) and the creation of algorithms for effective computational neural networks. This eventually led to deep learning networks in machine learning applications.

Advances were made in the 1980s as AI researchers began to apply rigorous mathematical and statistical analysis to develop algorithms. Hidden Markov Models were applied to speech.

A Hidden Markov Model is a model with unobserved (i.e., hidden) states. Combined with massive databases, they have resulted in vastly more robust speech recognition. Machine translation has also improved. Data mining, the first form of machine learning as it is known today, was developed. Chess programs improved initially through the use of specialized computers, such as IBM's Deep Blue. With the increase in processing power, powerful chess programs that are better than most human players are now available on personal computers.

The Bayesian network formalism was invented to allow for the rigorous application of uncertainty in reasoning problems. In the late 1990s, intelligent agents were introduced. Search engines, bots, and web site aggregators are examples of intelligent agents used on the Internet. Figure A-2 gives a timeline of selected events in the history of autonomous systems.

***Figure A-2:*** Artificial intelligence timeline.

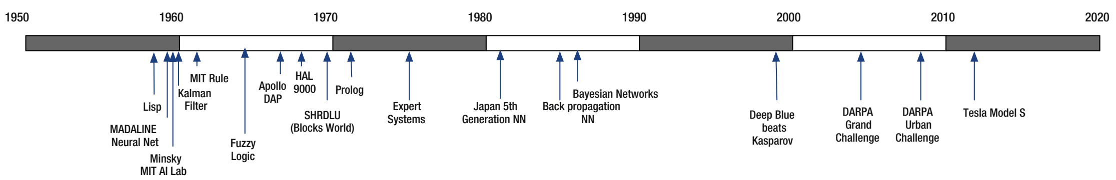


Today, the state of the art in AI includes autonomous cars, speech recognition, planning and scheduling, game playing, robotics, and machine translation. All of these are based on AI technology. They are in constant use today. You can take a PDF document and translate it into any language using Google translate. The translations are not perfect. One certainly would not use them to translate literature.

Recent advances in AI include IBM's Watson. Watson is a question-answering computing system with advanced natural language processing and information retrieval from massive databases. It defeated champion Jeopardy players in 2011. It is currently being applied to medical problems and many other complex problems.

## A.3 Learning Control

Adaptive or intelligent control was motivated in the 1950s [2] by the problems of aircraft control. Control systems of that time worked very well for linear systems. Aircraft dynamics could be linearized about a particular speed. For example, a simple equation for total velocity in level flight is:

$$m\frac{dv}{dt} = T - \frac{1}{2}\rho C_D S v^2 \tag{A-1}$$

This says the mass $m$ times the change in velocity per time, $\frac{dv}{dt}$, equals the thrust $T$, the force from the aircraft engine, minus the drag. $C_D$ is the aerodynamic drag coefficient and $S$ is the wetted area (i.e., the area that causes drag such as the wings and fuselage). The thrust is used

for control. This is a nonlinear equation in velocity $v$ because of the $v^2$ term. We can linearize it around a particular velocity $v_s$ so that $v = v_\delta + v_s$ and get:

$$m\frac{dv_\delta}{dt} = T - \rho C_D S v_s v_\delta \tag{A-2}$$

This equation is linear in $v_\delta$. We can control velocity with a simple thrust control law:

$$T = T_s - cv_\delta \tag{A-3}$$

where $T_s = \frac{1}{2}\rho C_D S v_s^2$. $c$ is the damping coefficient. $\rho$ is the atmospheric density and is a nonlinear function of altitude. For the linear control to work the control must be adaptive. If we want to guarantee a certain damping value, which is the quantity in parentheses:

$$m\frac{dv_\delta}{dt} = -\left(c + \rho C_D S v_s\right) v_\delta \tag{A-4}$$

we need to know $\rho$, $C_D$, $S$, and $v_s$. This approach leads to a gain scheduling control system where we measure the flight condition – altitude, velocity – and schedule the linear gains based on the where the aircraft is in the gain schedule.

In the 1960s, progress was made on adaptive control. State Space theory was developed, which made it easier to design multi-loop control systems, that is, control systems that controlled more than one state at a time with different control loops. The general space controller is:

$$\dot{x} = Ax + Bu \tag{A-5}$$

$$y = Cx + Du \tag{A-6}$$

$$u = -Ky \tag{A-7}$$

where $A$, $B$, $C$, and $D$ are matrices, $x$ is the state, $y$ is the measurement, and $u$ is the control input. A state is a quantity that changes with time that is needed to define what the system is doing. For a point mass that can only move in one direction, the position and velocity make up the two states. If $A$ completely models the system and $y$ contains all of the information about the state vector $x$, then this system is stable. Full state feedback would be $x = -Kx$, where $K$ can be computed to have guaranteed phase and gain margins (that is, tolerance to delays and tolerance to amplification errors). This was a major advance in control theory. Before this, multi-loop systems had to be designed separately and combined very carefully.

Learning control and adaptive control were found to be realizable from a common framework. The Kalman Filter, also known as linear quadratic estimation, was introduced.

Spacecraft required autonomous control, since they were often out of contact with the ground or the time delays were too long for effective ground supervision. The first digital autopilots were on the Apollo spacecraft, which first flew in 1968 on Apollo 7. Don Eyles book [9] gives the history of the Lunar Module Digital Autopilot. Geosynchronous communications satellites were automated to the point where one operator could fly a dozen satellites.

Advances in system identification, the process of just determining parameters of a system (such as the drag coefficient above) were made. Adaptive control was applied to real problems. Autopilots have progressed from fairly simple mechanical pilot augmentation systems to sophisticated control systems than can takeoff, cruise, and land under computer control.

In the 1970s, proof of adaptive control stability was made. Stability of linear control systems was well established, but adaptive systems are inherently nonlinear. Universally stabilizing controllers were studied. Progress was made in the robustness of adaptive control. Robustness is the ability of a system to deal with changes in parameters that were assumed to be known, sometimes due to failures in the systems. It was in the 1970s that digital control became widespread, replacing traditional analog circuits composed of transistors and operational amplifiers.

Adaptive controllers started to appear commercially in the 1980s. Most modern single-loop controllers have some form of adaptation. Adaptive techniques were also found to be useful for tuning controllers.

More recently, there has been a melding of artificial intelligence and control. Expert systems have been proposed that determine what algorithms (not just parameters) to use depending on the environment. For example, during a winged reentry of a glider, the control system would use one system in orbit, a second at high altitudes, a third during high Mach (Mach is the ratio of the velocity to the speed of sound) flight, and a fourth at low Mach numbers and during landing. An F3D Skynight used the Automatic Carrier Landing System on 12 August 1957. This was the first shipboard test of the landing system designed to land aircraft on board autonomously. Naira Hovakimyan (U of IL U-C), and also Nahn Nguyen (NASA) were pioneers in this area. Adaptive control was demonstrated on sub-scale F-18s, which controlled and landed the aircraft after most of one wing was lost!

## A.4 Machine Learning

Machine learning started as a branch of artificial intelligence. However, many techniques are much older. Thomas Bayes created Bayes theorem in 1763. Bayes theorem is:

$$P(A_i|B) = \frac{P(B|A_i)P(A_i)}{\sum P(B|A_i)} \tag{A-8}$$

$$P(A_i|B) = \frac{P(B|A_i)P(A_i)}{P(B)}$$

which is just the probability of $A_i$ given $B$. This assumes that $P(B) \neq 0$. In the Bayesian interpretation, the theorem introduces the effect of evidence on belief. One technique, regression, was discovered by Legendre in 1805 and Gauss in 1809.

As noted in the section on artificial intelligence, modern machine learning began with data mining, which is the process of getting new insights from data. In the early days of AI, there was considerable work on machines learning from data. However, this lost favor and in the 1990s it was reinvented as the field of machine learning. The goal was to solve practical problems of pattern recognition using statistics. This was greatly aided by the massive amounts

of data available online along with the tremendous increase in processing power available to developers. Machine learning is closely related to statistics.

In the early 1990s, Vapnik and co-workers invented a computationally powerful class of supervised learning networks known as support vector machines (SVMs). These networks could solve problems of pattern recognition, regression, and other machine learning problems.

A growing application of machine learning is autonomous driving. Autonomous driving makes use of all aspects of autonomous learning, including controls, artificial intelligence and machine learning. Machine vision is used in most systems as cameras are inexpensive and provide more information than lidar, radar or sonar (which are also useful). It isn't possible to build really safe autonomous driving systems without learning through experience. Thus, designers of such systems put their cars on the roads and collect experiences that are used to fine tune the system.

Other applications include high-speed stock trading and algorithms to guide investments. These are under rapid development and are now available to the consumer. Data mining and machine learning are in use to predict events, both human and natural. Searches on the internet have been used to track disease outbreaks. If there are a lot of data, and the internet makes gathering massive data easy, then you can be sure that machine learning techniques are being applied to mine the data.

## A.5 The Future

Autonomous learning in all its branches is undergoing rapid development today. Many of the technologies are used operationally even in low-cost consumer technology. Virtually every automobile company in the world, and many non-automotive companies, are working to perfect autonomous driving. Military organizations are extremely interested in artificial intelligence and machine learning. Combat aircraft today have systems to take over from the pilot to prevent planes from crashing to the ground.

Although completely autonomous systems are the goal in many areas, the meshing of human and machine intelligence is also an area of active research. Much AI research has been to study how the human mind works. In addition to improving performance by tapping into the extraordinary power of the brain, this work will also enable machine learning systems to mesh more seamlessly with human beings. This is critical for autonomous control involving people, but may also allow people to augment their own abilities.

This is an exciting time for machine learning! We hope that this book helps you to bring your own advances to machine learning!

# APPENDIX B

# Software for Machine Learning

## B.1 Autonomous Learning Software

There are many sources for machine learning software. Machine learning encompasses both software to help the user learn from data and software that helps machines learn and adapt to their environment. This book gives you a sampling of software that you can use immediately. However, the software is not designed for industrial applications. This chapter describes software that is available for the MATLAB environment. Both professional and open source MATLAB software is discussed. The book may not cover every available package as new packages are continually becoming available and older packages may become obsolete.

The packages you select for your project depend on your goal and your level of software expertise. Many of the packages in this chapter are research tools that can be used to design, analyze, and improve various types of machine learning systems, but to turn them into deployable, production quality systems would require compiling, integrating and testing software that is custom developed for each application. Other packages, such as commercial expert system shells, can be deployed as your application. You'll look for packages that are most compatible with your deployment environment. For example, if your goal is an embedded system, you will need development tools for the embedded processors and packages that are most compatible with that development environment.

This chapter includes software for what is conventionally called "machine learning." These are the statistics functions that help to give us an insight into data. These are often used in the context of "big data." It also includes descriptions of packages for other branches of autonomous learning systems, such as system identification. System identification is a branch of automatic control that learns about the systems under control, allowing for better and more precise control.

The chapter, for completeness, also covers popular software that is MATLAB-compatible, but requires extra steps to use it from within MATLAB. Examples include R, Python, and SNOPT. In all cases, it is straightforward to write MATLAB interfaces to these packages. Using MATLAB as a front end can be very helpful and allow you to create integrated packages that include MATLAB, Simulink, and the machine learning package of your choice.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2

You will note that we include optimization software. Optimization is a tool used as part of machine learning to find the best or "optimal" parameters. We use it in this book in our decision tree chapter.

Don't be upset if we didn't include your favorite package, or your package! We apologize in advance.

# B.2 Commercial MATLAB Software

## B.2.1 MathWorks Products

The MathWorks sells several packages for machine learning. These are in the Machine Learning branch of our taxonomy shown in Figure 1.3. The MathWorks products provide high-quality algorithms for data analysis along with graphics tools to visualize the data. Visualization tools are a critical part of any machine learning system. They can be used for data acquisition, for example, for image recognition or as part of systems for autonomous control of vehicles, or for diagnosis and debugging during development. All of these packages can be integrated with each other and with other MATLAB functions to produce powerful systems for machine learning. The most applicable toolboxes that we will discuss are:

- Statistics and Machine Learning Toolbox
- Neural Network Toolbox
- Computer Vision System Toolbox
- System Identification Toolbox
- MATLAB for Deep Learning
- Fuzzy Logic Toolbox

### B.2.1.1 Statistics and Machine Learning Toolbox

The Statistics and Machine Learning Toolbox provides data analytics methods for gathering trends and patterns from massive amounts of data. These methods do not require a model for analyzing the data. The toolbox functions can be broadly divided into Classification Tools, Regression Tools, and Clustering Tools.

Classification methods are used to place data into different categories. For example, data, in the form of an image, might be used to classify an image of an organ as having a tumor. Classification is used for handwriting recognition, credit scoring, and face identification. Classification methods include support vector machines (SVMs), decision trees and neural networks.

Regression methods let you build models from current data to predict future data. The models can then be updated as new data become available. If the data are only used once to create the model, then it is a batch method. A regression method that incorporates data as it becomes available is a recursive method.

Clustering finds natural groupings in data. Object recognition is an application of clustering methods. For example, if you want to find a car in an image you look for data that are associated with the part of an image that is a car. Although cars are of different shapes and sizes, they have many features in common.

The toolbox has many functions to support these areas and many that do not fit neatly into these categories. The Statistics and Machine Learning Toolbox is an excellent place to start for professional tools that are seamlessly integrated into the MATLAB environment.

#### B.2.1.2 Neural Network Toolbox

The MATLAB Neural Network Toolbox is a comprehensive neural net toolbox that seamlessly integrates with MATLAB. The toolbox provides functions to create, train, and simulate neural networks. The toolbox includes convolutional neural networks and deep learning networks. Neural networks can be computationally intensive owing to the large numbers of nodes and associated weights, especially during training. The Neural Network Toolbox allows you to distribute computation across multicore processors and graphical processing units (GPUs) if you have the Parallel Computing Toolbox, another MATLAB add-on. You can extend this even further to a network cluster of computers using the MATLAB Distributed Computing Server™. As with all MATLAB products, the Neural Network Toolbox provides extensive graphics and visualization capabilities that make it easier to understand your results.

The Neural Network Toolbox is capable of handling large data sets. This could be gigabytes or terabytes of data. This makes it suitable for industrial strength problems and complex research. MATLAB also provides videos, webinars, and tutorials, including a full suite of resources for applying deep learning.

#### B.2.1.3 Computer Vision System Toolbox

The MATLAB Computer Vision System Toolbox provides functions for developing computer vision systems. The toolbox provides extensive support for video processing, but also includes functions for feature detection and extraction. It also supports 3D vision and can process information from stereo cameras. 3D motion detection is supported. This is particularly helpful if your inputs are from cameras.

#### B.2.1.4 System Identification Toolbox

The System Identification Toolbox provides MATLAB functions and Simulink blocks for constructing mathematical models of systems. You can identify transfer functions from input/output data and perform parameter identification for models. Both linear and nonlinear system identification is supported. This could be used as part of a control system. It is often helpful to divide your system into parts that can have equations as models and parts that can't be modeled with equations. Your system can learn the parameters of the former part with the System Identification Toolbox.

#### B.2.1.5 MATLAB for Deep Learning

MATLAB for Deep Learning allows you to design, build, and visualize convolutional neural networks. You can easily implement models such as GoogLENet, VGG-16, VGG_19, AlexNet, and ResNet-59. It has extensive capabilities for visualization and debugging of neural networks. This is important to ensure that your system is behaving properly. It includes a number of pre-trained models. Deep Learning is a type of neural net. You can use this when the Neural Net Toolbox functions aren't sufficient for your system. You can use both toolboxes together, along with all the other MATLAB toolboxes.

### B.2.2 Princeton Satellite Systems Products

Several of our own commercial packages provide tools within the purview of autonomous learning.

#### B.2.2.1 Core Control Toolbox

The Core Control Toolbox provides the control and estimation functions of our Spacecraft Control Toolbox with general industrial dynamics examples, including robotics and chemical processing. The suite of Kalman Filter routines includes conventional filters, Extended Kalman Filters, and Unscented Kalman Filters. The Unscented Filters have a fast sigma point calculation algorithm. All of the Kalman Filters use a common code format with separate prediction and update functions. This allows the two steps to be used independently. The filters can handle multiple measurement sources that can be changed dynamically, with measurements arriving at different times.

Add-ons for the Core Control Toolbox include Imaging and Target Tracking modules. Imaging includes lens models, image processing, ray tracing, and image analysis tools.

#### B.2.2.2 Target Tracking

The Target Tracking module employs track-oriented multiple hypothesis testing. Track-oriented multiple hypothesis testing is a powerful technique for assigning measurements to tracks of objects when the number of objects is unknown or changing. It is absolutely essential for accurate tracking of multiple objects.

In many situations a sensor system must track multiple targets, like in rush hour traffic. This leads to the problem of associating measurements with objects, or tracks. This is a crucial element of any practical tracking system.

The track-oriented approach recomputes the hypotheses using the newly updated tracks after each scan of data is received. Rather than maintaining, and expanding, hypotheses from scan to scan, the track-oriented approach discards the hypotheses formed on scan $k-1$. The tracks that survive pruning are propagated to the next scan $k$, where new tracks are formed, using the new observations, and reformed into hypotheses. The hypothesis formation step is formulated as a mixed integer linear program (MILP) and solved using a GLPK (GNU Linear Programming Kit). Except for the necessity to delete some tracks based upon low probability, no information is lost because the track scores that are maintained contain all the relevant statistical data.

The MHT Module uses a powerful track-pruning algorithm that does the pruning in one step. Because of its speed, ad-hoc pruning methods are not required, leading to more robust and reliable results. The track management software is, as a consequence, quite simple.

All three Kalman Filters in the Core Toolbox, including Extended and Unscented, can be used independently or as part of the MHT system. The Unscented Kalman Filter automatically uses sigma points and does not require derivatives to be taken of the measurement functions or linearized versions of the measurement models.

Interactive multiple model systems (IMMs) can also be used as part of the MHT system. IMMs employ multiple dynamic models to facilitate tracking maneuvering objects. One model might involve maneuvering while another models constant motion. Measurements are assigned to all of the models. The IMMs are based on Jump Markovian Systems.

## B.3 MATLAB Open Source Resources

MATLAB open source tools are a great resource for implementing state-of-the-art machine learning. Machine learning and convex optimization packages are available.

### B.3.1 DeepLearnToolbox

The Deep Learn Toolbox by Rasmus Berg Palm is a MATLAB toolbox for Deep Learning. It includes Deep Belief Nets, Stacked Autoencoders, Convolutional Neural Nets, and other neural net functions. It is available through the MathWorks File Exchange.

### B.3.2 Deep Neural Network

The Deep Neural Network by Masayuki Tanaka provides deep learning tools of deep belief networks of stacked restricted Boltzmann machines. It has functionality for both unsupervised and supervised learning. It is available through the MathWorks File Exchange.

### B.3.3 MatConvNet

MatConvNet implements Convolutional Neural Networks for image processing. It includes a range of pre-trained networks for image processing functions. You can find it by searching on the name, or at the time of printing at http://www.vlfeat.org/matconvnet/. This package is open source and is open to contributors.

## B.4 Non- MATLAB Products for Machine Learning

There are many products, both open source and commercial, for machine learning. We cover some of the more popular open-source products. Both machine learning and convex optimization packages are discussed.

### B.4.1 R

R is open-source software for statistical computing. It compiles on MacOS, UNIX, and Windows. It is similar to the Bell Labs S language developed by John Chambers and colleagues. It includes many statistical functions and graphics techniques.

You can use R in batch mode from MATLAB using the `system` command. Write

```
system('R CMD BATCH inputfile outputfile');
```

This runs the code in `inputfile` and puts it into `outputfile`. You can then read the `outputfile` into MATLAB.

### B.4.2 scikit-learn

scikit-learn is a machine learning library for use in Python. It includes a wide variety of tools including:

1. Classification
2. Regression
3. Clustering
4. Dimensionality reduction
5. Model selection
6. Preprocessing

scikit-learn is well suited to a wide variety of data mining and data analysis problems.

MATLAB supports the reference implementation of Python, CPython. Mac users and Linux users already have Python installed. Windows users need to install a distribution.

### B.4.3 LIBSVM

LIBSVM [6] is a library for SVMs. It has an extensive collection of tools for SVMs, including extensions by many users of LIBSVM. LIBSVM Tools include distributed processing and multi-core extensions. The authors are Chih-Chung Chang and Chih-Jen Lin. You can find it at https://www.csie.ntu.edu.tw/~cjlin/libsvm/.

## B.5 Products for Optimization

Optimization tools often are used as part of machine learning systems. Optimizers minimize a cost given a set of constraints on the variables that are optimized. The maximum or minimum value for a variable is one type of constraint. Constraints and costs may be linear or nonlinear.

### B.5.1 LOQO

LOQO [27] is a system for solving smooth constrained optimization problems, available from Princeton University. The problems can be linear or nonlinear, convex or nonconvex, constrained or unconstrained. The only real restriction is that the functions defining the problem be smooth (at the points evaluated by the algorithm). If the problem is convex, LOQO finds a globally optimal solution. Otherwise, it finds a locally optimal solution near to a given starting point.

Once you compile the mex-file interface to LOQO, you must pass it an initial guess and sparse matrices for the problem definition variables. You may also pass in a function handle to provide animation of the algorithm at each iteration of the solution.

### B.5.2 SNOPT

SNOPT [10] is a software package for solving large-scale optimization problems (linear and nonlinear programs) hosted at the University of California, San Diego. It is especially effective for nonlinear problems whose functions and gradients are expensive to evaluate. The functions should be smooth but need not be convex. SNOPT is designed to take advantage of the sparsity of the Jacobian matrix, effectively reducing the size of the problem being solved. For optimal control problems, the Jacobian is very sparse because you have a matrix with rows and columns that span a large number of time points, but only adjacent time points can have nonzero entries.

SNOPT makes use of nonlinear function and gradient values. The solution obtained will be a local optimum (which may or may not be a global optimum). If some of the gradients are unknown, they will be estimated by finite differences. Infeasible problems are treated methodically via elastic bounds. SNOPT allows the nonlinear constraints to be violated and minimizes the sum of such violations. Efficiency is improved in large problems if only some of the variables are nonlinear, or if the number of active constraints is nearly equal to the number of variables.

### B.5.3 GLPK

GLPK (GNU Linear Programming Kit) solves a variety of linear programming problems. It is part of the GNU project (https://www.gnu.org/software/glpk/). The most well-known one is solving the linear program:

$$Ax = b \tag{B-1}$$

$$y = cx \tag{B-2}$$

where it is desired to find $x$, which, when it is multiplied by $A$, equals $b$. $c$ is the cost vector, which, when multiplied by $x$, gives the scalar cost of applying $x$. If $x$ is the same length as $b$ the solution is:

$$x = A^{-1}b \tag{B-3}$$

Otherwise, we can use GLPK to solve for $x$, which minimizes $y$. GLPK can solve this problem and others where one or more elements of $x$ has to be an integer or even just 0 or 1.

### B.5.4 CVX

CVX [4] is a MATLAB-based modeling system for convex optimization. CVX turns MATLAB into a modeling language, allowing constraints and objectives to be specified using standard MATLAB expression syntax.

In its default mode, CVX supports a particular approach to convex optimization that we call disciplined convex programming. Under this approach, convex functions and sets are built up from a small set of rules from convex analysis, starting from a base library of convex functions and sets. Constraints and objectives that are expressed using these rules are automatically transformed to a canonical form and solved. CVX can be used for free with solvers such as SeDuMi or with commercial solvers if a license is obtained from CVX Research.

### B.5.5 SeDuMi

SeDuMi [25] is MATLAB software for optimization over second-order cones, currently hosted at Lehigh University. It can handle quadratic constraints. SeDuMi was used in Acikmese [1]. SeDuMi stands for Self-Dual-Minimization. It implements the *self-dual* embedding technique over *self-dual* homogeneous cones. This makes it possible to solve certain optimization problems in one phase. SeDuMi is available as part of YALMIP and as a standalone package.

### B.5.6 YALMIP

YALMIP is free MATLAB software by Johan Lofberg that provides an easy-to-use interface to other solvers. It interprets constraints and can select the solver based on the constraints. SeDuMi and MATLAB's *fmincon* from the Optimization Toolbox are available solvers.

## B.6 Products for Expert Systems

There are dozens, if not hundreds, of expert system shells. For MATLAB users, the most useful shells are ones for which the C or C++ code is available. It is straightforward to write an interface in C using a .mex file. CLIPS is an expert system shell https:http://www.clipsrules.net/?q=AboutCLIPS. It stands for 'C' Language Integrated Production System. It is a rule-based language for creating expert systems. It allows users to implement heuristic solutions easily. It has been used for many applications including:

- An intelligent training system for space shuttle flight controllers
- Applications of artificial intelligence to space shuttle mission control
- PI-in-a-Box: A knowledge-based system for space science experimentation
- The DRAIR Advisor: A knowledge-based system for materiel deficiency analysis
- The Multi-mission VICAR Planner: Image processing for scientific data
- IMPACT: Development and deployment experience of network event correlation applications

- The NASA personnel security processing expert system
- Expert system technology for nondestructive waste assay
- Hybrid knowledge-based system for automatic classification of B-scan images from ultrasonic rail inspection
- An expert system for recognition of facial actions and their intensity
- Development of a hybrid knowledge-based system for multi-objective optimization of power distribution system operations

CLIPS is currently maintained by Gary Riley who has written the book, "Expert Systems: Principles and Programming," now in its fourth edition. In the next section we will learn about .mex files to interface CLIPS, and other software, with MATLAB.

## B.7 MATLAB MEX files

### B.7.1 Problem

CLIPS needs to be connected to MATLAB.

### B.7.2 Solution

The solution is to create a .mex file to interface MATLAB and CLIPS.

### B.7.3 How It Works

Look at the file MEXTest.c. This is a C file with the accompanying header, .h, file.

```
//
//  MEXTest.c
//

#include "MEXTest.h"
#include "mex.h"

void mexFunction( int nlhs, mxArray *plhs[], int nrhs, const mxArray
    *prhs[] )
{
    // Check the arguments
    if( nrhs != 2 )
    {
        mexErrMsgTxt("Two inputs required.");
    }
```

```
    if( nlhs != 1 )
    {
        mexErrMsgTxt("One output required.");
    }
}
```

and MEXTest.h

```
//
//  MEXTest.h
//

#ifndef MEXTest_h
#define MEXTest_h

#include <stdio.h>

#endif /* MEXTest_h */
```

You can edit these files in MATLAB or in any text editor.

Now type in the command line

```
>> mex MEXTest.c
Building with 'Xcode␣Clang++'.
MEX completed successfully.
```

mex just calls your development system's compiler and linker. "XCode" is the MacOS development environment. "Clang" is a C/C++ compiler. You end up with the file `MEXTest.mexmaci64` if you are using MacOS. Typing `help MEXTest` in the command window does not return anything. If you want help, add a file `MEXTest.m` such as:

```
function CLIPS
%% Help for the CLIPS.cpp file
```

If you try and run the function you will get:

```
>> MEXTest
Error using MEXTest
Two inputs required.
```

The CLIPS .mex file isn't much more complicated. The CLIPS .mex file reads in a rules file and then takes inputs to which the rule is applied. The rule file is `Rules.CLP` and is:

```
(defrule troubleshoot-car
  (wheel-turning no) (power yes) (torque-command yes)
  =>
  (cbkFunction))
```

`defrule` defines a rule. `cbkFunction` is the call back function. It is called (also known as fired) if the rule `troubleshoot-car` is true. If all conditions are true it fires `cbkFunction`. You can write the rules file using any text editor. To run the .mex file you need to type:

```
>> mex CLIPS.c -lCLIPS
Building with 'Xcode␣Clang++'.
MEX completed successfully.
```

“-lCLIPS” loads the dynamics library `libCLIPS.dylib`. This dynamics library was built using XCode on the MacOS from CLIPS source code, which is a set of C files. Now you are ready to run the function. Type the facts into the command window:

```
>> facts = '(wheel-turning␣no)␣(power␣yes)␣(torque-command␣yes)'

facts =
    '(wheel-turning␣no)␣(power␣yes)␣(torque-command␣yes)'
```

To run the function, call CLIPS with this facts variable,

```
>> CLIPS(facts)

ans =
    'Wheel␣failed'
```

Now see if it knows when the wheel is working. Change the first fact to “yes.”

```
>> facts =  '(wheel-turning␣yes)␣(power␣yes)␣(torque-command␣yes)';
>> CLIPS(facts)

ans =
    'Wheel␣working'
```

You can use this function to create any expert system by making a more elaborate rule base.

# Bibliography

[1] Behcet Acikmese and Scott R. Ploen. Convex Programming Approach to Powered Descent Guidance for Mars Landing. *Journal of Guidance, Control, and Dynamics*, 30(5):1353–1366, 2007.

[2] K. J. Åström and B. Wittenmark. *Adaptive Control, Second Edition*. Addison-Wesley, 1995.

[3] S.S. Blackman and R.F. Popoli. *Design and Analysis of Modern Tracking Systems*. Artech House, 1999.

[4] S. Boyd. CVX: Matlab Software for Disciplined Convex Programming. http://cvxr.com/cvx/, 2015.

[5] A. E. Bryson Jr. *Control of Spacecraft and Aircraft*. Princeton, 1994.

[6] Chih-Chung Chang and Chih-Jen Lin. LIBSVM – A Library for Support Vector Machines. https://www.csie.ntu.edu.tw/~cjlin/libsvm/, 2015.

[7] Ka Cheok et al. Fuzzy Logic-Based Smart Automatic Windshield Wiper. *IEEE Control Systems*, December 1996.

[8] Corinna Cortes and Vladimir Vapnik. Support-Vector Networks. *Machine Learning*, 20:273–297, 1995.

[9] Don Eyles. *Sunburst and Luminary: An Apollo Memoir*. Fort Point Press, 2018.

[10] Philip Gill, Walter Murray, and Michael Saunders. SNOPT 6.0 Description. http://www.sbsi-sol-optimize.com/asp/sol_products_snopt_desc.htm, 2013.

[11] J. Grus. *Data Science from Scratch*. O'Reilly, 2015.

[12] S. Haykin. *Neural Networks*. Prentice-Hall, 1999.

[13] Matthijs Hollemans. Convolutional Neural Networks on the iPhone with VGGNet. http://matthijshollemans.com/2016/08/30/vggnet-convolutional-neural-network-iphone/, 2016.

[14] P. Jackson. *Introduction to Expert Systems, Third Edition*. Addison-Wesley, 1999.

© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2

[15] Byoung S. Kim and Anthony J. Calise. Nonlinear Flight Control Using Neural Networks. *Journal of Guidance, Control, and Dynamics*, 20(1):26–33, 1997.

[16] J. B. Mueller. *Design and Analysis of Optimal Ascent Trajectories for Stratospheric Airships*. PhD thesis, University of Minnesota, 2013.

[17] J. B. Mueller and G. J. Balas. Implementation and Testing of LPV Controllers for the F/A-18 System Research Aircraft. In *Proceedings*, number AIAA-2000- 4446. AIAA, August 2000.

[18] Andrew Ng, Jiquan Ngiam, Chuan Yu Foo, Yifan Mai, Caroline Suen, Adam Coates, Andrew Maas, Awni Hannun, Brody Huval, Tao Wang, and Sameep Tandon. UFDL Tutorial. http://ufldl.stanford.edu/tutorial/, 2016.

[19] Nils J. Nilsson. *Artificial Intelligence: A New Synthesis*. Morgran Kaufmann Publishers, 1998.

[20] Sebastian Raschka. *Python Machine Learning*. [PACKT], 2015.

[21] D. B. Reid. An Algorithm for Tracking Multiple Targets. *IEEE Transactions on Automatic Control*, AC=24(6):843–854, December 1979.

[22] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach, Third Edition*. Prentice-Hall, 2010.

[23] S. Sarkka. Lecture 3: Bayesian Optimal Filtering Equations and the Kalman Filter. Technical Report, Department of Biomedical Engineering and Computational Science, Aalto University School of Science, February 2011.

[24] L. D. Stone, C. A. Barlow, and T. L. Corwin. *Bayesian Multiple Target Tracking*. Artech House, 1999.

[25] Jos F. Sturm. Using SeDuMi 1.02, a MATLAB Toolbox for Optimization Over Symmetric Cones. http://sedumi.ie.lehigh.edu/wp-content/sedumi-downloads/usrguide.ps, 1998.

[26] K Terano, Asai T., and M. Sugeno. *Fuzzy Systems Theory and Its Applications*. Academic Press, 1992.

[27] R. J. Vanderbei. LOQO USERS MANUAL VERSION 4.05. http://www.princeton.edu/~rvdb/tex/loqo/loqo405.pdf, September 2013.

[28] M. C. VanDyke, J. L. Schwartz, and C. D. Hall. Unscented Kalman Filtering for Spacecraft Attitude State and Parameter Estimation. *Advances in Astronautical Sciences*, 2005.

[29] Matthew G. Villella. *Nonlinear Modeling and Control of Automobiles with Dynamic Wheel-Road Friction and Wheel Torque Inputs*. PhD thesis, Georgia Institute of Technology, April 2004.

[30] Peggy S. Williams-Hayes. Flight Test Implementation of a Second Generation Intelligent Flight Control System. Technical Report NASA/TM-2005-213669, NASA Dryden Flight Research Center, November 2005.

# Index


## A

Adaptive control, 9
- MRAC (*see* Model Reference Adaptive Control (MRAC))
- self tuning
  - modeling an oscillator, 110–111
  - tuning an oscillator, 112–116
- ship steering, 126–130
- spacecraft pointing, 130–133
- square wave, 121–122

Artificial intelligence (AI)
- back propagation, 319
- Bayesian network, 320
- Blocks World, 318
- chess programs, 319, 320
- Cybernetics, 317
- definition of, 16
- expert systems, 17
- Google translate, 320
- GPS, 318
- Hidden Markov Models, 319–320
- intelligent cars, 16–17
- knowledge-based systems, 319
- limitations, 319
- Lisp, 318
- LT, 317–318
- military organizations, 323
- neural networks, 317
- timeline of, 320
- time sharing, 318–319
- Towers of Hanoi, 318

Automobile animation, 299–302

Automobile demo
- car trajectories, 307
- final tree, 309
- radar measurement, 308

Automobile dynamics
- planar model, 293–294
- RungeKutta, 292
- vehicle states, 292
- wheel force and torque, 294–295

AutomobilePassing, 297–299

Automobile radar, 295–297

Automobile simulation
- Kalman Filter, 303–306
- snap shots, 302

Automobile target tracking, 306–309

Automobile 3D model, 300

Autonomous driving, 323
- automobile animation, 299–302
- automobile dynamics, 292–295
- AutomobilePassing, 297–299
- automobile radar, 295–297
- automobile simulation and Kalman Filter, 303–306
- technology, 16

Autonomous learning
- AI, 317–320
- categories of, 7–8
- learning control, 320–322
- machine learning, 322–323
- software, 325–326

AutoRadar function, 296

## B

Bayesian network, 17, 320

Bayes theorem, 322

Billiard ball Kalman filter, 274–279


© Michael Paluszek and Stephanie Thomas 2019

M. Paluszek and S. Thomas, *MATLAB Machine Learning Recipes*,
https://doi.org/10.1007/978-1-4842-3916-2

Binary decision trees
- autonomous learning taxonomy, 147
- box data structure fields, 162
- child boxes, 161
- classification error, 156
- ClassifierSet, 148–151
- distinct values, 158
- entropy, 156
- FindOptimalAction, 159
- fminbnd, 159
- Gini impurity, 155, 156
- homogeneity measure, 156–158
- IG, 155
- MATLAB function patch, 150
- parent/child nodes, 159
- PointInPolygon, 149
- testing data, 148, 165–168
- training, 160–162, 167–168

Blocks World, 318


## C


Case-based expert systems
- autonomous learning taxonomy, 311
- building, 312–313
- functions and scripts, 316
- running
  - BuildExpertSystem, 314–316
  - CBREngine, 313
  - ExpertSystemDemo, 314
  - fprintf, 316
  - strcmpi, 313–314

catColorReducer, 34

Cat images
- grayscale photographs, 211
- ImageArray, 212–213
- ScaleImage, 214–215
- 64x64 pixel, 213

Cell arrays, 20–21

Chapman–Kolmogorov equation, 83

Cholesky factorization, 102

C Language Integrated Production System (CLIPS), 319, 332–333

Classification tree, 13

Comma-Separated Lists, 20–21

Commercial software
- MathWorks products, 326–328
- PSS products, 328–329

Computer Vision System Toolbox, 327

ConnectNode, 51

Convolution process, 216
- deep learning, 210
- layers, 210
- stages, 225

Core Control Toolbox, 328

CVX, 332

Cybernetics, 317


## D


Damped oscillator, 114

Data mining, 323

Datastores
- functions, 26
- properties, 25

Data structures, 21–22
- parameters, 30–33

Daylight detector, 171–173

Decision trees, 13–14

Deep learning, 14, 328
- convolutional neural net, 210
- neural net, 209

Deep Learn Toolbox, 329

Deep Neural Network, 329

Digits
- CreateDigitImage function, 188–190
- DigitTrainingData, 190–192
- feed-forward neural network, 195
- grayscale, conversion, 189
- GUI, 192
- MLFF neural network function, 193
- multiple outputs
  - MAT-file, 206
  - multiple-digit neural net, 205–207
  - training data, 204
- NeuralNetDeveloper tool, 192

NeuralNetMLFF, 196–197
NeuralNetTraining, 196
Neuron activation functions, 192–195
Poisson or shot noise, 188
SaveTS function, 190
single output node
default parula map, 200
Digit0FontsTS, 197–198
NeuralNetTrainer, 200
node weights, 202
RMSE, 198, 200
sigmoid function, 197
single digit training error, 200
testing, 202–203
DrawBinaryTree
cell array, 152
data structure, 151
DefaultDataStructure, 154
demo, 155
DrawBox, 152
lines, 153
patch function, 152
resize rows, 152–153
RGB numbers, 152
sprintf, 154
text function, 152
DrawNode, 51
dynamicExpression, 22

## E

Euler integration, 94
Extended Kalman Filter (EKF), 92–97

## F

Fact-gathering, 311–312
Fast Fourier Transform (FFT), 9, 35, 110
FFTEnergy, 112
Filter covariances, 278
Filter errors, 279, 287
Flexible Image Transport System (FITS), 23
F-16 model, aircraft, 237–238
FORTRAN, XVII, 318
Frequency spectrum, 114
without noise, 115
Function PlotSet, XVIII–XIX
Fuzzy logic
AND and OR, 141
BuildFuzzySystem, 137–138
Defuzzify, 142
description, 135
Fire, 141
Fuzzify, 140–141
MATLAB data structure, 136
membership functions, 138–139
set structure, 137
smart wipers
rain wetness and intensity, 144, 145
wiper speed and interval, 144
wiper speed and interval *vs.* droplet frequency and wetness, 145

## G

Gaussian membership function, 138
General bell function, 138
General Problem Solver (GPS), 318
GNU Linear Programming Kit (GLPK), XIX, 328, 331
Google translate, 320
Graphical user interface (GUI), 45, 280
blank, 60
inspector, 61
snapshot
editing window, 62
simulation, 63, 64
GraphicConverter application, 214
Graphics
animation, bar chart, 63–67
building GUI, 58–63
custom two-dimensional diagrams, 50–51
general 2D, 48–49
three-dimensional box, 51–54
3D graphics, 56–58
3D object with texture, 54–56
2D line plots, 45–47

## H
Hidden Markov Models (HMM), 82, 319–320

## I, J
Images
  display options, 24
  formats, 23
  functions, 25
  information, 23–24
Inclined plane, 2
Information gain (IG), 155
Interactive multiple model systems (IMMs), 329

## K
Kalman Filters, 8
  automobile simulation, 303–306
  Chapman–Kolmogorov equation, 83
  Cholesky factorization, 102
  derivation, 80
  Euler integration, 94
  extended, angle measurement, 97
  family tree, 81
  HMM, 82
  implementation, 87
  linear, 74–92
  Monte Carlo methods, 80–81
  noise matrix, 91, 92
  normal/Gaussian random variable, 82
  OscillatorDamping RatioSim, 76
  OscillatorSim, 78
  parameter estimation, UKF, 104–107
  RHSOscillator, 78
  Spring-mass-damper system, 75, 77, 79
  state estimation
    EKF, 92–97
    linear, 74–92
    UKF, 97–103
  undamped natural frequency, 76
Kernel function, 15
Knowledge-based systems, 17, 319

## L
Large MAT-files, 29
Learning control, aircraft, 320–322
  dynamic pressure, 245
  Kalman Filter, 247
  least squares solution, 245
  longitudinal dynamics, 261–264
  neural net, 243
  PID controller, 244
  pinv function, 245
  recursive learning algorithm, 246, 247
  sigma-pi neural net, 243, 244
LIBSVM, 330
Linear Kalman Filter, 74–92
Linear regression, 12
Lisp, 318
Logic Theorist (LT), 317–318
Log-likelihood ratio, 269–270
Longitudinal control, aircraft, 231
  differential equations, 235
  drag polar, 233
  dynamics symbols, 233–234
  F-16 model, 237–238
  learning approach, 232
  longitudinal dynamics, 232, 233
  Oswald efficiency factor, 234
  RHSAircraft, 235–236
  sigma-pi type network, 232–233
  training algorithm, 233
LOQO, 331

## M
Machine learning
  AI, 322
  autonomous driving, 323
  Bayes theorem, 322
  concept of learning, 4–6
  data mining, 323
  definition of, 2
  elements
    data, 2
    models, 3

training, 3
examples, XVII
feedback control, 8–9
FORTRAN, XVII
SVMs, 323
taxonomy, 6–8
Mapreduce
datastore, 33–35
framework, 26
progress, 35
valueIterator class, 34
MatConvNet, 329
MAT-file function, 29
MathWorks products
Computer Vision System Toolbox, 327
Deep Learning, 328
Neural Network Toolbox, 327
Statistics and Machine Learning Toolbox, 326–327
System Identification Toolbox, 327
MATLAB toolbox
functions, XVIII
html help, XVIII
scripts, XVIII
Matrices, 19–20
Membership functions, fuzzy logic
Gaussian, 138
general bell, 138
sigmoidal, 138
trapezoid, 138
triangular, 138
MEX files, 333–335
Mixed integer linear program (MILP), 272
Model Reference Adaptive Control (MRAC)
implementation, 117–121
rotor, 123–125
Monte Carlo methods, 80–81
Multi-layer feed-forward (MLFF), 14, 193
Multiple hypothesis testing (MHT), 269
estimated states, 289
GUI, 283, 308
information window, 284
measurement and gates, 271
object states, 287, 289
testing parameters, 282
tree, 284


## N


Nelder–Meade simplex, 229
Neural aircraft control
activation function, 242–243
Combination function, 248–249
equilibrium state, 238–240
learning control (*see* Learning control, aircraft)
longitudinal dynamics simulation, 232
nonlinear simulation, 261–264
numerical simulation, 240–242
pitch angle, PID controller, 256–258
sigma-pi net neural function, 249–251
Neural networks/nets, 14–15
convolution layer, 217–218
daylight detector, 171–173
description, 171
fully connected layer, 220–222
image processing, 224
image recognition, 228–230
matrix convolution, 215–217
number recognition, 225–228
pendulum (*see* Pendulum)
pitch dynamics, 258–261
pooling to outputs, 218–220
probability determination, 222–223
single neuron angle estimator, 177–181
testing, 223–225
training image generation, 211–215
Neural Network Toolbox, 327
New track measurements, 268
Nonlinear simulation, aircraft control, 261–264
Non-MATLAB products
LIBSVM, 330
R, 330
scikit-learn, 330

Normal/Gaussian random variable, 82
Numerics, 23

## O

One-dimensional motion, MHT, 285–287
    track association, 287–289
Online learning, 4
Open source resources
    Deep Learn Toolbox, 329
    Deep Neural Network, 329
    MatConvNet, 329
Optimization tools
    CVX, 332
    GLPK, 331
    LOQO, 331
    SeDuMi, 332
    SNOPT, 331
    YALMIP, 332
OscillatorDamping RatioSim, 76–77
OscillatorSim, 78

## P, Q

Parallel Computing Toolbox, 26, 33, 327
patch function, 50, 52, 54
Pattern recognition, 187
Pendulum
    activation function, 183
    dynamics, 173
    linear equations, 175, 177
    magnitude oscillation, 185–186
    NeuralNetMLFF, 182–184
    NeuralNetTraining, 182
    NNPendulumDemo, 182
    nonlinear equations, 177
    PendulumSim, 176
    RungeKutta integration, 174, 175
    Taylor's series expansion, 175
    torque, 174
    xDot, 176
Perceptrons, 319
Pitch angle, PID controller, 256–258
Pitch dynamics, 231
    neural net, 258–261
Planar automobile dynamical model, 293
PlotSet function, 46, 47
plotXY function, 47
Pluto, 3D globe, 55
Princeton Satellite Systems (PSS) products
    Core Control Toolbox, 328
    Target Tracking, 328–329
Processing table data, 37–41
Proportional integral derivative (PID) controller
    closed loop transfer function, 253
    coding, 254–255
    derivative operator, 253
    design, 254
    double integrator equations, 255
    feedback controller, 252
    nonlinear inversion controller, 258
    pitch angle, 251, 256–258
    recursive training, 252

## R

R, 330
Recursive learning algorithm, 246, 247
Regression, 10–13
RHSAircraft, 240
RHSOscillator, 78
Riccati equation, 128
Root mean square error (RMSE), 198, 200, 205
Rotor, MRAC
    gain convergence, 125
    RungeKutta, 123–124
    speed control, 117
    SquareWave, 123
Rule-based expert systems, 311, 312
RungeKutta, 240

## S

SCARA robot, 69, 70
scikit-learn, 330
Second-order system, 58, 59
SeDuMi, 332
Semi-supervised learning, 4
Ship steering

gains and rudder angle, 129
Gaussian white noise, 130
heading control, 126
parameters, 127
Riccati equation, 128–129
ShipSim, 127
Sigma-pi neural net function, 243, 244, 248–251
Sigmoidal membership function, 138
Sigmoid function, 242
Simple binary tree, 147–148
Simple machines, 2
Single neuron angle estimator
activation functions, 178–179
linear estimator, 177
OneNeuron, 180
tanh neuron output, 181
SNOPT, 331
Softmax function, 222
Software
autonomous learning, 325–326
commercial MATLAB, 326–329
expert systems, 332–333
MATLAB MEX files, 333–335
MATLAB open source resources, 329
non-MATLAB products, 329–330
optimization tools, 330–332
Solar flux, 172
Spacecraft model, 131
Spacecraft simulation, 133
Sparse matrices, 27, 28
sphere function, 55
Spring-mass-damper system, 75, 77, 79, 111
Spurious measurement, tracking, 267
Square wave, 122
Statistics and Machine Learning Toolbox, 326–327
Strings
arrays of, 41
concatenation, 41
substrings, 42
Supervised learning, 3
Support vector machines (SVMs), 15, 323
Synonym set, 211
System Identification Toolbox, 327


## T


Table creation, FFTs, 35–37
Tables and categoricals, 27–28
TabularTextDatastore, 38–41
Tall arrays, 26–27
Target Tracking, 328–329
Towers of Hanoi, 318
Tracking
algorithm, 269–270
definition of, 265
hypothesis formation, 271–272
measurements, 269
assignment, 270–271
new track, 268
spurious, 267
valid, 268
problem, 268
track pruning, 272–273
Track-oriented multiple hypothesis testing (MHT), 17, 265, 266, 328
Trapezoid membership function, 138
Tree diagrams, graphics functions, 50
Triangular membership function, 138
Two by three bar chart, 65, 67
2D plot types, 48–49


## U


Undamped natural frequency, 76
Unscented Kalman Filter (UKF), 8, 303
non-augmented Kalman Filter, 97–103
parameter estimation, 104–107
true and estimated states, 305
Unsupervised learning, 4


## V, W, X


Valid measurements, tracking, 268
varargin, 30–32, 46


## Y, Z


YALMIP, 332