

# Computer Vision

## A Modern Approach

David A. Forsyth

*University of California at Berkeley*

Jean Ponce

*University of Illinois at Urbana-Champaign*

=====*An Alan R. Apt Book*=====



Prentice Hall  
Upper Saddle River, New Jersey 07458

## Library of Congress Cataloging-in-Publication Data

CIP data on file.

Vice President and Editorial Director, ECS: *Marcia J. Horton*  
Publisher: *Alan Apt*  
Associate Editor: *Toni D. Holm*  
Editorial Assistant: *Patrick Lindner*  
Vice President and Director of Production and Manufacturing, ESM: *David W. Riccardi*  
Executive Managing Editor: *Vince O'Brien*  
Assistant Managing Editor: *Camille Trentacoste*  
Production Editor: *Leslie Galen*  
Director of Creative Services: *Paul Belfanti*  
Creative Director: *Carole Anson*  
Art Director: *Wanda Espana*  
Art Editor: *Greg Dulles*  
Manufacturing Manager: *Trudy Piscioti*  
Manufacturing Buyer: *Lynda Castillo*  
Marketing Manager: *Pamela Shaffer*  
Marketing Assistant: *Barrie Reinhold*

*About the cover:* Image courtesy of Gamma/Superstock



© 2003 by Pearson Education, Inc.  
Pearson Education, Inc.  
Upper Saddle River, NJ 07458

All rights reserved. No part of this book may be reproduced, in any form or by any means, without permission in writing from the publisher.

The author and publisher of this book have used their best efforts in preparing this book. These efforts include the development, research, and testing of the theories and programs to determine their effectiveness. The author and publisher make no warranty of any kind, expressed or implied, with regard to these programs or the documentation contained in this book. The author and publisher shall not be liable in any event for incidental or consequential damages in connection with, or arising out of, the furnishing, performance, or use of these programs.

Printed in the United States of America

10 9 8 7 6 5 4 3 2 1

ISBN 0-13-085198-1

Pearson Education Ltd., *London*  
Pearson Education Australia Pty. Ltd., *Sydney*  
Pearson Education Singapore, Pte. Ltd.  
Pearson Education North Asia Ltd., *Hong Kong*  
Pearson Education Canada, Inc., *Toronto*  
Pearson Educación de México, S.A. de C.V.  
Pearson Education–Japan, *Tokyo*  
Pearson Education Malaysia, Pte. Ltd.  
Pearson Education, *Upper Saddle River, New Jersey*

*To my family—DAF*

*To Camille and Oscar—JP*



# *Contents*

<b>Part I Image Formation and Image Models</b>	<b>1</b>
<b>1 CAMERAS</b>	<b>3</b>
1.1 Pinhole Cameras 4	
1.1.1 Perspective Projection, 4	
1.1.2 Affine Projection, 6	
1.2 Cameras with Lenses 7	
1.2.1 Paraxial Geometric Optics, 8	
1.2.2 Thin Lenses, 9	
1.2.3 Real Lenses, 11	
1.3 The Human Eye 13	
1.4 Sensing 15	
1.4.1 CCD Cameras, 16	
1.4.2 Sensor Models, 17	
1.5 Notes 18	
Problems 19	
<b>2 GEOMETRIC CAMERA MODELS</b>	<b>20</b>
2.1 Elements of Analytical Euclidean Geometry 20	
2.1.1 Coordinate Systems and Homogeneous Coordinates, 20	
2.1.2 Coordinate System Changes and Rigid Transformations, 23	
2.2 Camera Parameters and the Perspective Projection 28	
2.2.1 Intrinsic Parameters, 29	

2.2.2	Extrinsic Parameters, 30	
2.2.3	A Characterization of Perspective Projection Matrices, 31	
2.3	Affine Cameras and Affine Projection Equations 32	
2.3.1	Affine Cameras, 32	
2.3.2	Affine Projection Equations, 33	
2.3.3	A Characterization of Affine Projection Matrices, 34	
2.4	Notes 35	
	Problems 35	
<b>3</b>	<b>GEOMETRIC CAMERA CALIBRATION</b>	<b>38</b>
3.1	Least-Squares Parameter Estimation 39	
3.1.1	Linear Least-Squares Methods, 39	
3.1.2	Nonlinear Least-Squares Methods, 42	
3.2	A Linear Approach to Camera Calibration 45	
3.2.1	Estimation of the Projection Matrix, 45	
3.2.2	Estimation of the Intrinsic and Extrinsic Parameters, 45	
3.2.3	Degenerate Point Configurations, 46	
3.3	Taking Radial Distortion into Account 47	
3.3.1	Estimation of the Projection Matrix, 48	
3.3.2	Estimation of the Intrinsic and Extrinsic Parameters, 48	
3.3.3	Degenerate Point Configurations, 49	
3.4	Analytical Photogrammetry 50	
3.5	An Application: Mobile Robot Localization 51	
3.6	Notes 53	
	Problems 53	
<b>4</b>	<b>RADIOMETRY—MEASURING LIGHT</b>	<b>55</b>
4.1	Light in Space 55	
4.1.1	Foreshortening, 55	
4.1.2	Solid Angle, 56	
4.1.3	Radiance, 58	
4.2	Light at Surfaces 60	
4.2.1	Simplifying Assumptions, 60	
4.2.2	The Bidirectional Reflectance Distribution Function, 60	
4.2.3	Example: The Radiometry of Thin Lenses, 62	
4.3	Important Special Cases 63	
4.3.1	Radiosity, 63	
4.3.2	Directional Hemispheric Reflectance, 64	
4.3.3	Lambertian Surfaces and Albedo, 65	
4.3.4	Specular Surfaces, 65	
4.3.5	The Lambertian + Specular Model, 66	
4.4	Notes 67	
	Problems 68	
<b>5</b>	<b>SOURCES, SHADOWS, AND SHADING</b>	<b>70</b>
5.1	Qualitative Radiometry 70	
5.2	Sources and Their Effects 71	
5.2.1	Radiometric Properties of Light Sources, 71	
5.2.2	Point Sources, 72	

- 5.2.3 Line Sources, 75
- 5.2.4 Area Sources, 75
- 5.3 Local Shading Models 77
  - 5.3.1 Local Shading Models for Point Sources, 77
  - 5.3.2 Area Sources and Their Shadows, 79
  - 5.3.3 Ambient Illumination, 79
- 5.4 Application: Photometric Stereo 80
  - 5.4.1 Normal and Albedo from Many Views, 82
  - 5.4.2 Shape from Normals, 84
- 5.5 Interreflections: Global Shading Models 86
  - 5.5.1 An Interreflection Model, 87
  - 5.5.2 Solving for Radiosity, 89
  - 5.5.3 The Qualitative Effects of Interreflections, 90
- 5.6 Notes 91
  - Problems 95

## 6 COLOR

97

- 6.1 The Physics of Color 97
  - 6.1.1 Radiometry for Colored Lights: Spectral Quantities, 97
  - 6.1.2 The Color of Sources, 98
  - 6.1.3 The Color of Surfaces, 100
- 6.2 Human Color Perception 102
  - 6.2.1 Color Matching, 102
  - 6.2.2 Color Receptors, 104
- 6.3 Representing Color 105
  - 6.3.1 Linear Color Spaces, 105
  - 6.3.2 Non-linear Color Spaces, 110
  - 6.3.3 Spatial and Temporal Effects, 113
- 6.4 A Model for Image Color 114
  - 6.4.1 Cameras, 114
  - 6.4.2 A Model for Image Color, 115
  - 6.4.3 Application: Finding Specularities, 118
- 6.5 Surface Color from Image Color 119
  - 6.5.1 Surface Color Perception in People, 120
  - 6.5.2 Inferring Lightness, 122
  - 6.5.3 Surface Color from Finite-Dimensional Linear Models, 126
- 6.6 Notes 128
  - Problems 130

## Part II Early Vision: Just One Image

133

### 7 LINEAR FILTERS

135

- 7.1 Linear Filters and Convolution 135
  - 7.1.1 Convolution, 136
- 7.2 Shift Invariant Linear Systems 140
  - 7.2.1 Discrete Convolution, 141
  - 7.2.2 Continuous Convolution, 142
  - 7.2.3 Edge Effects in Discrete Convolutions, 145

- 7.3 Spatial Frequency and Fourier Transforms 145
  - 7.3.1 Fourier Transforms, 146
- 7.4 Sampling and Aliasing 148
  - 7.4.1 Sampling, 150
  - 7.4.2 Aliasing, 152
  - 7.4.3 Smoothing and Resampling, 153
- 7.5 Filters as Templates 157
  - 7.5.1 Convolution as a Dot Product, 157
  - 7.5.2 Changing Basis, 158
- 7.6 Technique: Normalized Correlation and Finding Patterns 158
  - 7.6.1 Controlling the Television by Finding Hands by Normalized Correlation, 158
- 7.7 Technique: Scale and Image Pyramids 159
  - 7.7.1 The Gaussian Pyramid, 159
  - 7.7.2 Applications of Scaled Representations, 161
- 7.8 Notes 162
  - Problems 164

## **8 EDGE DETECTION**

165

- 8.1 Noise 165
  - 8.1.1 Additive Stationary Gaussian Noise, 166
  - 8.1.2 Why Finite Differences Respond to Noise, 168
- 8.2 Estimating Derivatives 169
  - 8.2.1 Derivative of Gaussian Filters, 169
  - 8.2.2 Why Smoothing Helps, 170
  - 8.2.3 Choosing a Smoothing Filter, 171
  - 8.2.4 Why Smooth with a Gaussian? 172
- 8.3 Detecting Edges 175
  - 8.3.1 Using the Laplacian to Detect Edges, 175
  - 8.3.2 Gradient-Based Edge Detectors, 176
  - 8.3.3 Technique: Orientation Representations and Corners, 181
- 8.4 Notes 185
  - Problems 187

## **9 TEXTURE**

189

- 9.1 Representing Texture 190
  - 9.1.1 Extracting Image Structure with Filter Banks, 191
  - 9.1.2 Representing Texture Using the Statistics of Filter Outputs, 193
- 9.2 Analysis (and Synthesis) Using Oriented Pyramids 196
  - 9.2.1 The Laplacian Pyramid, 197
  - 9.2.2 Filters in the Spatial Frequency Domain, 199
  - 9.2.3 Oriented Pyramids, 202
- 9.3 Application: Synthesizing Textures for Rendering 202
  - 9.3.1 Homogeneity, 205
  - 9.3.2 Synthesis by Sampling Local Models, 205
- 9.4 Shape from Texture 207
  - 9.4.1 Shape from Texture for Planes, 208
- 9.5 Notes 210
  - Problems 211

<b>Part III Early Vision: Multiple Images</b>	<b>213</b>
<b>10 THE GEOMETRY OF MULTIPLE VIEWS</b>	<b>215</b>
10.1 Two Views 216	
10.1.1 Epipolar Geometry, 216	
10.1.2 The Calibrated Case, 217	
10.1.3 Small Motions, 218	
10.1.4 The Uncalibrated Case, 218	
10.1.5 Weak Calibration, 219	
10.2 Three Views 222	
10.2.1 Trifocal Geometry, 223	
10.2.2 The Calibrated Case, 223	
10.2.3 The Uncalibrated Case, 226	
10.2.4 Estimation of the Trifocal Tensor, 226	
10.3 More Views 227	
10.4 Notes 230	
Problems 232	
<b>11 STEREOPSIS</b>	<b>234</b>
11.1 Reconstruction 235	
11.1.1 Image Rectification, 236	
11.2 Human Stereopsis 237	
11.3 Binocular Fusion 240	
11.3.1 Correlation, 240	
11.3.2 Multi-Scale Edge Matching, 242	
11.3.3 Dynamic Programming, 243	
11.4 Using More Cameras 246	
11.4.1 Three Cameras, 246	
11.4.2 Multiple Cameras, 247	
11.5 Notes 247	
Problems 250	
<b>12 AFFINE STRUCTURE FROM MOTION</b>	<b>251</b>
12.1 Elements of Affine Geometry 253	
12.1.1 Affine Spaces and Barycentric Combinations, 253	
12.1.2 Affine Subspaces and Affine Coordinates, 254	
12.1.3 Affine Transformations and Affine Projection Models, 257	
12.1.4 Affine Shape, 258	
12.2 Affine Structure and Motion from Two Images 258	
12.2.1 Geometric Scene Reconstruction, 259	
12.2.2 Algebraic Motion Estimation, 260	
12.3 Affine Structure and Motion from Multiple Images 263	
12.3.1 The Affine Structure of Affine Image Sequences, 263	
12.3.2 A Factorization Approach to Affine Structure from Motion, 263	
12.4 From Affine to Euclidean Images 266	
12.4.1 Euclidean Constraints and Calibrated Affine Cameras, 266	
12.4.2 Computing Euclidean Upgrades from Multiple Views, 267	
12.5 Affine Motion Segmentation 269	
12.5.1 The Reduced Row-Echelon Form of the Data Matrix, 269	

- 12.5.2 The Shape Interaction Matrix, 270
- 12.6 Notes 272
- Problems 272

### **13 PROJECTIVE STRUCTURE FROM MOTION 274**

- 13.1 Elements of Projective Geometry 275
  - 13.1.1 Projective Spaces, 275
  - 13.1.2 Projective Subspaces and Projective Coordinates, 276
  - 13.1.3 Affine and Projective Spaces, 278
  - 13.1.4 Hyperplanes and Duality, 279
  - 13.1.5 Cross-Ratios and Projective Coordinates, 280
  - 13.1.6 Projective Transformations, 282
  - 13.1.7 Projective Shape, 284
- 13.2 Projective Structure and Motion from Binocular Correspondences 284
  - 13.2.1 Geometric Scene Reconstruction, 284
  - 13.2.2 Algebraic Motion Estimation, 285
- 13.3 Projective Motion Estimation from Multilinear Constraints 287
  - 13.3.1 Motion Estimation from Fundamental Matrices, 288
  - 13.3.2 Motion Estimation from Trifocal Tensors, 288
- 13.4 Projective Structure and Motion from Multiple Images 289
  - 13.4.1 A Factorization Approach to Projective Structure from Motion, 289
  - 13.4.2 Bundle Adjustment, 292
- 13.5 From Projective to Euclidean Images 292
- 13.6 Notes 294
- Problems 296

## **Part IV Mid-Level Vision 299**

### **14 SEGMENTATION BY CLUSTERING 301**

- 14.1 What Is Segmentation? 301
  - 14.1.1 Model Problems, 303
  - 14.1.2 Segmentation as Clustering, 304
- 14.2 Human Vision: Grouping and Gestalt 304
- 14.3 Applications: Shot Boundary Detection and Background Subtraction 309
  - 14.3.1 Background Subtraction, 309
  - 14.3.2 Shot Boundary Detection, 310
- 14.4 Image Segmentation by Clustering Pixels 313
  - 14.4.1 Segmentation Using Simple Clustering Methods, 313
  - 14.4.2 Clustering and Segmentation by K-means, 315
- 14.5 Segmentation by Graph-Theoretic Clustering 317
  - 14.5.1 Terminology for Graphs, 317
  - 14.5.2 The Overall Approach, 318
  - 14.5.3 Affinity Measures, 319
  - 14.5.4 Eigenvectors and Segmentation, 321
  - 14.5.5 Normalized Cuts, 323
- 14.6 Notes 326
- Problems 327

<b>15</b>	<b>SEGMENTATION BY FITTING A MODEL</b>	<b>329</b>
15.1	The Hough Transform 329	
15.1.1	Fitting Lines with the Hough Transform, 330	
15.1.2	Practical Problems with the Hough Transform, 330	
15.2	Fitting Lines 333	
15.2.1	Line Fitting with Least Squares, 333	
15.2.2	Which Point Is on Which Line? 335	
15.3	Fitting Curves 337	
15.3.1	Implicit Curves, 337	
15.3.2	Parametric Curves, 340	
15.4	Fitting as a Probabilistic Inference Problem 341	
15.5	Robustness 342	
15.5.1	M-estimators, 343	
15.5.2	RANSAC, 346	
15.6	Example: Using RANSAC to Fit Fundamental Matrices 348	
15.6.1	An Expression for Fitting Error, 348	
15.6.2	Correspondence as Noise, 348	
15.6.3	Applying RANSAC, 348	
15.6.4	Finding the Distance, 349	
15.6.5	Fitting a Fundamental Matrix to Known Correspondences, 351	
15.7	Notes 351	
	Problems 352	
<b>16</b>	<b>SEGMENTATION AND FITTING USING PROBABILISTIC METHODS</b>	<b>354</b>
16.1	Missing Data Problems, Fitting, and Segmentation 354	
16.1.1	Missing Data Problems, 355	
16.1.2	The EM Algorithm, 357	
16.1.3	The EM Algorithm in the General Case, 359	
16.2	The EM Algorithm in Practice 359	
16.2.1	Example: Image Segmentation, Revisited, 359	
16.2.2	Example: Line Fitting with EM, 362	
16.2.3	Example: Motion Segmentation and EM, 363	
16.2.4	Example: Using EM to Identify Outliers, 365	
16.2.5	Example: Background Subtraction Using EM, 367	
16.2.6	Example: EM and the Fundamental Matrix, 368	
16.2.7	Difficulties with the EM Algorithm, 368	
16.3	Model Selection: Which Model Is the Best Fit? 369	
16.3.1	Basic Ideas, 369	
16.3.2	AIC—An Information Criterion, 369	
16.3.3	Bayesian Methods and Schwartz' BIC, 370	
16.3.4	Description Length, 370	
16.3.5	Other Methods for Estimating Deviance, 370	
16.4	Notes 371	
	Problems 372	
<b>17</b>	<b>TRACKING WITH LINEAR DYNAMIC MODELS</b>	<b>373</b>
17.1	Tracking as an Abstract Inference Problem 374	
17.1.1	Independence Assumptions, 374	

- 17.1.2 Tracking as Inference, 375
- 17.1.3 Overview, 376
- 17.2 Linear Dynamic Models 376
  - 17.2.1 Drifting Points, 377
  - 17.2.2 Constant Velocity, 377
  - 17.2.3 Constant Acceleration, 377
  - 17.2.4 Periodic Motion, 379
  - 17.2.5 Higher Order Models, 379
- 17.3 Kalman Filtering 380
  - 17.3.1 The Kalman Filter for a 1D State Vector, 380
  - 17.3.2 The Kalman Update Equations for a General State Vector, 383
  - 17.3.3 Forward–Backward Smoothing, 383
- 17.4 Data Association 388
  - 17.4.1 Choosing the Nearest—Global Nearest Neighbours, 389
  - 17.4.2 Gating and Probabilistic Data Association, 390
- 17.5 Applications and Examples 393
  - 17.5.1 Vehicle Tracking, 393
- 17.6 Notes 397
  - Problems 397

## Part V High-Level Vision: Geometric Methods

399

### 18 MODEL-BASED VISION

401

- 18.1 Initial Assumptions 401
  - 18.1.1 Obtaining Hypotheses, 402
- 18.2 Obtaining Hypotheses by Pose Consistency 403
  - 18.2.1 Pose Consistency for Perspective Cameras, 403
  - 18.2.2 Affine and Projective Camera Models, 404
  - 18.2.3 Linear Combinations of Models, 406
- 18.3 Obtaining Hypotheses by Pose Clustering 407
- 18.4 Obtaining Hypotheses Using Invariants 410
  - 18.4.1 Invariants for Plane Figures, 410
  - 18.4.2 Geometric Hashing, 413
  - 18.4.3 Invariants and Indexing, 413
- 18.5 Verification 416
  - 18.5.1 Edge Proximity, 416
  - 18.5.2 Similarity in Texture, Pattern, and Intensity, 418
- 18.6 Application: Registration in Medical Imaging Systems 418
  - 18.6.1 Imaging Modes, 419
  - 18.6.2 Applications of Registration, 420
  - 18.6.3 Geometric Hashing Techniques in Medical Imaging, 421
- 18.7 Curved Surfaces and Alignment 422
- 18.8 Notes 423
  - Problems 425

### 19 SMOOTH SURFACES AND THEIR OUTLINES

427

- 19.1 Elements of Differential Geometry 429
  - 19.1.1 Curves, 429
  - 19.1.2 Surfaces, 434

- 19.2 Contour Geometry 438
  - 19.2.1 The Occluding Contour and the Image Contour, 439
  - 19.2.2 The Cusps and Inflections of the Image Contour, 439
  - 19.2.3 Koenderink's Theorem, 441
- 19.3 Notes 442
  - Problems 443

## **20 ASPECT GRAPHS 444**

- 20.1 Visual Events: More Differential Geometry 447
  - 20.1.1 The Geometry of the Gauss Map, 448
  - 20.1.2 Asymptotic Curves, 449
  - 20.1.3 The Asymptotic Spherical Map, 450
  - 20.1.4 Local Visual Events, 451
  - 20.1.5 The Bitangent Ray Manifold, 453
  - 20.1.6 Multilocal Visual Events, 455
- 20.2 Computing the Aspect Graph 456
  - 20.2.1 Step 1: Tracing Visual Events, 457
  - 20.2.2 Step 2: Constructing the Regions, 459
  - 20.2.3 Remaining Steps of the Algorithm, 459
  - 20.2.4 An Example, 460
- 20.3 Aspect Graphs and Object Localization 460
- 20.4 Notes 464
  - Problems 465

## **21 RANGE DATA 467**

- 21.1 Active Range Sensors 467
- 21.2 Range Data Segmentation 469
  - 21.2.1 Elements of Analytical Differential Geometry, 469
  - 21.2.2 Finding Step and Roof Edges in Range Images, 471
  - 21.2.3 Segmenting Range Images into Planar Regions, 476
- 21.3 Range Image Registration and Model Acquisition 477
  - 21.3.1 Quaternions, 478
  - 21.3.2 Registering Range Images Using the Iterative Closest-Point Method, 479
  - 21.3.3 Fusing Multiple Range Images, 481
- 21.4 Object Recognition 483
  - 21.4.1 Matching Piecewise-Planar Surfaces Using Interpretation Trees, 483
  - 21.4.2 Matching Free-Form Surfaces Using Spin Images, 486
- 21.5 Notes 490
  - Problems 491

## **Part VI High-Level Vision: Probabilistic and Inferential Methods 493**

### **22 FINDING TEMPLATES USING CLASSIFIERS 495**

- 22.1 Classifiers 496
  - 22.1.1 Using Loss to Determine Decisions, 496
  - 22.1.2 Overview: Methods for Building Classifiers, 497

- 22.1.3 Example: A Plug-in Classifier for Normal Class-conditional Densities, 499
- 22.1.4 Example: A Nonparametric Classifier Using Nearest Neighbors, 500
- 22.1.5 Estimating and Improving Performance, 501
- 22.2 Building Classifiers from Class Histograms 502
  - 22.2.1 Finding Skin Pixels Using a Classifier, 503
  - 22.2.2 Face Finding Assuming Independent Template Responses, 504
- 22.3 Feature Selection 505
  - 22.3.1 Principal Component Analysis, 507
  - 22.3.2 Identifying Individuals with Principal Components Analysis, 508
  - 22.3.3 Canonical Variates, 512
- 22.4 Neural Networks 516
  - 22.4.1 Key Ideas, 516
  - 22.4.2 Minimizing the Error, 518
  - 22.4.3 When to Stop Training, 520
  - 22.4.4 Finding Faces Using Neural Networks, 520
  - 22.4.5 Convolutional Neural Nets, 522
- 22.5 The Support Vector Machine 523
  - 22.5.1 Support Vector Machines for Linearly Separable Datasets, 524
  - 22.5.2 Finding Pedestrians Using Support Vector Machines, 527
- 22.6 Notes 528
  - Problems 531
- 22.7 Appendix I: Backpropagation 531
- 22.8 Appendix II: Support Vector Machines for Datasets That Are Not Linearly Separable 534
- 22.9 Appendix III: Using Support Vector Machines with Non-Linear Kernels 535

## **23 RECOGNITION BY RELATIONS BETWEEN TEMPLATES**

537

- 23.1 Finding Objects by Voting on Relations between Templates 538
  - 23.1.1 Describing Image Patches, 538
  - 23.1.2 Voting and a Simple Generative Model, 539
  - 23.1.3 Probabilistic Models for Voting, 539
  - 23.1.4 Voting on Relations, 541
  - 23.1.5 Voting and 3D Objects, 542
- 23.2 Relational Reasoning Using Probabilistic Models and Search 542
  - 23.2.1 Correspondence and Search, 543
  - 23.2.2 Example: Finding Faces, 545
- 23.3 Using Classifiers to Prune Search 546
  - 23.3.1 Identifying Acceptable Assemblies Using Projected Classifiers, 547
  - 23.3.2 Example: Finding People and Horses Using Spatial Relations, 548
- 23.4 Technique: Hidden Markov Models 550
  - 23.4.1 Formal Matters, 550
  - 23.4.2 Computing with Hidden Markov Models, 552
  - 23.4.3 Varieties of HMMs, 557
- 23.5 Application: Hidden Markov Models and Sign Language Understanding 559
  - 23.5.1 Language Models: Sentences from Words, 560
- 23.6 Application: Finding People with Hidden Markov Models 563
- 23.7 Notes 566

<b>24</b>	<b><i>GEOMETRIC TEMPLATES FROM SPATIAL RELATIONS</i></b>	<b>568</b>
24.1	Simple Relations between Object and Image 568	
24.1.1	Relations for Curved Surfaces, 569	
24.1.2	Class-Based Grouping, 575	
24.2	Primitives, Templates, and Geometric Inference 576	
24.2.1	Generalized Cylinders as Volumetric Primitives, 576	
24.2.2	Ribbons, 579	
24.2.3	What Can One Represent with Ribbons? 584	
24.2.4	Linking 3D and 2D for Cylinders of Known Length, 586	
24.2.5	Linking 3D and Image Data Using Explicit Geometric Reasoning, 588	
24.3	Afterword: Object Recognition 591	
24.3.1	The Facts on the Ground, 592	
24.3.2	Current Approaches to Object Recognition, 593	
24.3.3	Limitations, 594	
24.4	Notes 594	
	Problems 595	
<b>Part VII</b>	<b>Applications</b>	<b>597</b>
<b>25</b>	<b><i>APPLICATION: FINDING IN DIGITAL LIBRARIES</i></b>	<b>599</b>
25.1	Background: Organizing Collections of Information 601	
25.1.1	How Well Does the System Work? 601	
25.1.2	What Do Users Want? 602	
25.1.3	Searching for Pictures, 602	
25.1.4	Structuring and Browsing, 604	
25.2	Summary Representations of the Whole Picture 604	
25.2.1	Histograms and Correlograms, 605	
25.2.2	Textures and Textures of Textures, 606	
25.3	Representations of Parts of the Picture 609	
25.3.1	Segmentation, 610	
25.3.2	Template Matching, 612	
25.3.3	Shape and Correspondence, 613	
25.3.4	Clustering and Organizing Collections, 614	
25.4	Video 617	
25.5	Notes 618	
<b>26</b>	<b><i>APPLICATION: IMAGE-BASED RENDERING</i></b>	<b>620</b>
26.1	Constructing 3D Models from Image Sequences 620	
26.1.1	Scene Modeling from Registered Images, 621	
26.1.2	Scene Modeling from Unregistered Images, 627	
26.2	Transfer-Based Approaches to Image-Based Rendering 629	
26.2.1	Affine View Synthesis, 630	
26.2.2	Euclidean View Synthesis, 632	
26.3	The Light Field 635	
26.4	Notes 639	
	Problems 639	
	<b><i>BIBLIOGRAPHY</i></b>	<b>643</b>
	<b><i>INDEX</i></b>	<b>673</b>



# *Preface*

Computer vision as a field is an intellectual frontier. Like any frontier, it is exciting and disorganised; there is often no reliable authority to appeal to—many useful ideas have no theoretical grounding, and some theories are useless in practice; developed areas are widely scattered, and often one looks completely inaccessible from the other. Nevertheless, we have attempted in this book to present a fairly orderly picture of the field.

We see computer vision—or just “vision”; apologies to those who study human or animal vision—as an enterprise that uses statistical methods to disentangle data using models constructed with the aid of geometry, physics and learning theory. Thus, in our view, vision relies on a solid understanding of cameras and of the physical process of image formation (part I of this book) to obtain simple inferences from individual pixel values (part II), combine the information available in multiple images into a coherent whole (part III), impose some order on groups of pixels to separate them from each other or infer shape information (part IV), and recognize objects using geometric information (part V) or probabilistic techniques (part VI). Computer vision has a wide variety of applications, old (e.g., mobile robot navigation, industrial inspection, and military intelligence) and new (e.g., human computer interaction, image retrieval in digital libraries, medical image analysis, and the realistic rendering of synthetic scenes in computer graphics). We discuss some of these applications in part VII.

## **WHY STUDY VISION?**

Computer vision’s great trick is extracting descriptions of the world from pictures or sequences of pictures. This is unequivocally useful. Taking pictures is usually non-destructive and some-

times discreet. It is also easy and (now) cheap. The descriptions that users seek can differ widely between applications. For example, a technique known as structure from motion makes it possible to extract a representation of what is depicted and how the camera moved from a series of pictures. People in the entertainment industry use these techniques to build three-dimensional (3D) computer models of buildings, typically keeping the structure and throwing away the motion. These models are used where real buildings cannot be; they are set fire to, blown up, etc. Good, simple, accurate and convincing models can be built from quite small sets of photographs. People who wish to control mobile robots usually keep the motion and throw away the structure. This is because they generally know something about the area where the robot is working, but don't usually know the precise robot location in that area. They can determine it from information about how a camera bolted to the robot is moving.

There are a number of other, important applications of computer vision. One is in medical imaging: One builds software systems that can enhance imagery, or identify important phenomena or events, or visualize information obtained by imaging. Another is in inspection: One takes pictures of objects to determine whether they are within specification. A third is in interpreting satellite images, both for military purposes—a program might be required to determine what militarily interesting phenomena have occurred in a given region recently; or what damage was caused by a bombing—and for civilian purposes—what will this year's maize crop be? How much rainforest is left? A fourth is in organizing and structuring collections of pictures. We know how to search and browse text libraries (though this is a subject that still has difficult open questions) but don't really know what to do with image or video libraries.

Computer vision is at an extraordinary point in its development. The subject itself has been around since the 1960s, but it is only recently that it has been possible to build useful computer systems using ideas from computer vision. This flourishing has been driven by several trends: Computers and imaging systems have become very cheap. Not all that long ago, it took tens of thousands of dollars to get good digital color images; now it takes a few hundred, at most. Not all that long ago, a color printer was something one found in few, if any, research labs; now they are in many homes. This means it is easier to do research. It also means that there are many people with problems to which the methods of computer vision apply. For example, people would like to organize their collection of photographs, make 3D models of the world around them, and manage and edit collections of videos. Our understanding of the basic geometry and physics underlying vision and, what is more important, what to do about it, has improved significantly. We are beginning to be able to solve problems that lots of people care about, but none of the hard problems have been solved and there are plenty of easy ones that have not been solved either (to keep one intellectually fit while trying to solve hard problems). It is a great time to be studying this subject.

### **What Is in This Book?**

This book covers what we feel a computer vision professional ought to know. However, it is addressed to a wider audience. We hope that those engaged in computational geometry, computer graphics, image processing, imaging in general, and robotics will find it an informative reference. We have tried to make the book accessible to senior undergraduates or graduate students with a passing interest in vision. Each chapter covers a different part of the subject, and, as a glance at Table 1 will confirm, chapters are relatively independent. This means that one can dip into the book as well as read it from cover to cover. Generally, we have tried to make chapters run from easy material at the start to more arcane matters at the end. Each chapter has brief notes at the end, containing historical material and assorted opinions. We have tried to produce a book that describes ideas that are useful, or likely to be so in the future. We have put emphasis on understanding the basic geometry and physics of imaging, but have tried to link this with actual

applications. In general, the book reflects the enormous recent influence of geometry and various forms of applied statistics on computer vision.

A reader who goes from cover to cover will hopefully be well informed, if exhausted; there is too much in this book to cover in a one-semester class. Of course, prospective (or active) computer vision professionals should read every word, do all the exercises, and report any bugs found for the second edition (of which it is probably a good idea to plan buying a copy!). While the study of computer vision does not require deep mathematics, it does require facility with a lot of different mathematical ideas. We have tried to make the book self contained, in the sense that readers with the level of mathematical sophistication of an engineering senior should be comfortable with the material of the book, and should not need to refer to other texts. We have also tried to keep the mathematics to the necessary minimum—after all, this book is about computer vision, not applied mathematics—and have chosen to insert what mathematics we have kept in the main chapter bodies instead of a separate appendix.

Generally, we have tried to reduce the interdependence between chapters, so that readers interested in particular topics can avoid wading through the whole book. It is not possible to make each chapter entirely self contained, and Table 1 indicates the dependencies between chapters.

**TABLE 1** Dependencies between chapters: It will be difficult to read a chapter if you don't have a good grasp of the material in the chapters it "requires." If you have not read the chapters labeled "helpful," you may need to look one or two things up.

Part	Chapter	Requires	Helpful
I	1: Cameras		
	2: Geometric camera models	1	
	3: Geometric camera calibration	2	
	4: Radiometry—measuring light		
	5: Sources, shadows and shading		4, 1
	6: Color		5
II	7: Linear filters		
	8: Edge detection	7	
	9: Texture	7	8
III	10: The geometry of multiple views	3	
	11: Stereopsis	10	
	12: Affine structure from motion	10	
	13: Projective structure from motion	12	
IV	14: Segmentation by clustering		9, 6, 5
	15: Segmentation by fitting a model		14
	16: Segmentation and fitting using probabilistic methods		15, 10
	17: Tracking with linear dynamic models		
V	18: Model-based vision	3	
	19: Smooth surfaces and their outlines	2	
	20: Aspect graphs	19	
	21: Range data		20, 19, 3
VI	22: Finding templates using classifiers		9, 8, 7, 6, 5
	23: Recognition by relations between templates		9, 8, 7, 6, 5
	24: Geometric templates from spatial relations	2, 1	16, 15, 14
VII	25: Application: Finding in digital libraries		16, 15, 14, 6
	26: Application: Image-based rendering	10	13, 12, 11, 6, 5, 3, 2, 1

## What Is Not in This Book

The computer vision literature is vast, and it was not easy to produce a book about computer vision that can be lifted by ordinary mortals. To do so, we had to cut material, ignore topics, and so on. We cut two entire chapters close to the last moment: One is an introduction to probability and inference, the other an account of methods for tracking objects with non-linear dynamics. These chapters appear on the book's web page <http://www.cs.berkeley.edu/~daf/book.html>.

We left out some topics because of personal taste, or because we became exhausted and stopped writing about a particular area, or because we learned about them too late to put them in, or because we had to shorten some chapter, or any of hundreds of other reasons. We have tended to omit detailed discussions of material that is mainly of historical interest, and offer instead some historical remarks at the end of each chapter. Neither of us claims to be a fluent intellectual archaeologist, meaning that ideas may have deeper histories than we have indicated. We just didn't get around to writing up deformable templates and mosaics, two topics of considerable practical importance; we will try to put them into the second edition.

## ACKNOWLEDGMENTS

In preparing this book, we have accumulated a significant set of debts. A number of anonymous reviewers have read several drafts of the book and have made extremely helpful contributions. We are grateful to them for their time and efforts. Our editor, Alan Apt, organized these reviews with the help of Jake Warde. We thank them both. Leslie Galen, Joe Albrecht, and Dianne Parish, of Integre Technical Publishing, helped us over numerous issues with proofreading and illustrations. Some images used herein were obtained from IMSI's Master Photos Collection, 1895 Francisco Blvd. East, San Rafael, CA 94901-5506, USA. In preparing the bibliography, we have made extensive use of Keith Price's excellent computer vision bibliography, which can be found at <http://iris.usc.edu/Vision-Notes/bibliography/contents.html>.

Both the overall coverage of topics and several chapters were reviewed by various colleagues, who made valuable and detailed suggestions for their revision. We thank Kobus Barnard, Margaret Fleck, David Kriegman, Jitendra Malik and Andrew Zisserman. A number of our students contributed suggestions, ideas for figures, proofreading comments, and other valuable material. We thank Okan Arıkan, Sébastien Blind, Martha Cepeda, Stephen Chenney, Frank Cho, Yakup Genc, John Haddon, Sergey Ioffe, Svetlana Lazebnik, Cathy Lee, Sung-il Pae, David Parks, Fred Rothganger, Attawith Sudsang, and the students in several offerings of our vision classes at U.C. Berkeley and UIUC. We have been very lucky to have colleagues at various universities use (often rough) drafts of our book in their vision classes. Institutions whose students suffered through these drafts include, besides ours, Carnegie-Mellon University, Stanford University, the University of Wisconsin at Madison, the University of California at Santa Barbara and the University of Southern California; there may be others we are not aware of. We are grateful for all the helpful comments from adopters, in particular Chris Bregler, Chuck Dyer, Martial Hebert, David Kriegman, B.S. Manjunath, and Ram Nevatia, who sent us many detailed and helpful comments and corrections. The book has also benefitted from comments and corrections from Aydin Alaylioglu, Srinivas Akella, Marie Banich, Serge Belongie, Ajit M. Chaudhari, Navneet Dalal, Richard Hartley, Glen Healey, Mike Heath, Hayley Iben, Stéphanie Jonquière, Tony Lewis, Benson Limketkai, Simon Maskell, Brian Milch, Tamara Miller, Cordelia Schmid, Brigitte and Gerry Serlin, Ilan Shimshoni, Eric de Sturler, Camillo J. Taylor, Jeff Thompson, Claire Vallat, Daniel S. Wilkerson, Jinghan Yu, Hao Zhang, and Zhengyou Zhang. If you find an apparent typographic error, please email DAF ([daf@cs.berkeley.edu](mailto:daf@cs.berkeley.edu)) with the details, using the phrase "book typo" in your email; we will try to credit the first finder of each typo in the second edition.

We also thank P. Besl, B. Boufama, J. Costeira, P. Debevec, O. Faugeras, Y. Genc, M. Hebert, D. Huber, K. Ikeuchi, A.E. Johnson, T. Kanade, K. Kutulakos, M. Levoy, S. Mahamud, R. Mohr, H. Moravec, H. Murase, Y. Ohta, M. Okutami, M. Pollefeys, H. Saito, C. Schmid, S. Sullivan, C. Tomasi, and M. Turk for providing the originals of some of the figures shown in this book.

DAF acknowledges a wide range of intellectual debts, starting at kindergarten. Important figures in the very long list of his creditors include Gerald Alanthwaite, Mike Brady, Tom Fair, Margaret Fleck, Jitendra Malik, Joe Mundy, Mike Rodd, Charlie Rothwell and Andrew Zisserman. JP cannot even remember kindergarten, but acknowledges his debts to Olivier Faugeras, Mike Brady, and Tom Binford. He also wishes to thank Sharon Collins for her help. Without her, this book, like most of his work, probably would have never been finished. Both authors would also like to acknowledge the profound influence of Jan Koenderink's writings on their work at large and on this book in particular.

## SAMPLE SYLLABI

The whole book can be covered in two (rather intense) semesters, by starting at the first page and plunging on. Ideally, one would cover one application chapter—probably the chapter on image-based rendering—in the first semester, and the other one in the second. Few departments will experience heavy demand for so detailed a sequence of courses. We have tried to structure this book so that instructors can choose areas according to taste. Sample syllabi for busy 15-week semesters appear in Tables 2 to 6, structured according to needs that can reasonably be expected. We would encourage (and expect!) instructors to rearrange these according to taste.

Table 2 contains a suggested syllabus for a one-semester introductory class in computer vision for seniors or first-year graduate students in computer science, electrical engineering, or other engineering or science disciplines. The students receive a broad presentation of the field, including application areas such as digital libraries and image-based rendering. Although the

**TABLE 2** A one-semester introductory class in computer vision for seniors or first-year graduate students in computer science, electrical engineering, or other engineering or science disciplines.

Week	Chapter	Sections	Key topics
1	1, 4	1.1, 4 (summary only)	pinhole cameras, radiometric terminology
2	5	5.1–5.5	local shading models; point, line and area sources; photometric stereo
3	6	all	color
4	7, 8	7.1–7.5, 8.1–8.3	linear filters; smoothing to suppress noise; edge detection
5	9	all	texture: as statistics of filter outputs; synthesis; shape from
6	10, 11	10.1, 11	basic multi-view geometry; stereo
7	14	all	segmentation as clustering
8	15	15.1–15.4	fitting lines, curves; fitting as maximum likelihood; robustness
9	16	16.1, 16.2	hidden variables and EM
10	17	all	tracking with a Kalman filter; data association
11	2, 3	2.1, 2.2, all of 3	camera calibration
12	18	all	model-based vision using correspondence and camera calibration
13	22	all	template matching using classifiers
14	23	all	matching on relations
15	25, 26	all	finding images in digital libraries; image based rendering

**TABLE 3** A syllabus for students of computer graphics who want to know the elements of vision that are relevant to their topic.

Week	Chapter	Sections	Key topics
1	1, 4	1.1, 4 (summary only)	pinhole cameras, radiometric terminology
2	5	5.1–5.5	local shading models; point, line and area sources; photometric stereo
3	6.1–6.4	all	color
4	7, 8	7.1–7.5, 8.1–8.3	linear filters; smoothing to suppress noise; edge detection
5	9	9.1–9.3	texture: as statistics of filter outputs; synthesis
6	2, 3	2.1, 2.2, all of 3	camera calibration
7	10, 11	10.1, 11	basic multi-view geometry; stereo
8	12	all	affine structure from motion
9	13	all	projective structure from motion
10	26	all	image-based rendering
11	15	all	fitting; robustness; RANSAC
12	16	all	hidden variables and EM
13	19	all	surfaces and outlines
14	21	all	range data
15	17	all	tracking, the Kalman filter and data association

**TABLE 4** A syllabus for students who are primarily interested in the applications of computer vision.

Week	Chapter	Sections	Key topics
1	1, 4	1.1, 4 (summary only)	pinhole cameras, radiometric terminology
2	5, 6	5.1,5.3, 5.4, 5.5, 6.1–6.4	local shading models; point, line and area sources; photometric stereo; color—physics, human perception, color spaces
3	2, 3	all	camera models and their calibration
4	7, 9	all of 7; 9.1–9.3	linear filters; texture as statistics of filter outputs; texture synthesis
5	10, 11	all	multiview geometry, stereo as an example
6	12,13	all	affine structure from motion; projective structure from motion
7	13, 26	all	projective structure from motion; image-based rendering
8	14	all	segmentation as clustering, particular emphasis on shot boundary detection and background subtraction
9	15	all	fitting lines, curves; robustness; RANSAC
10	16	all	hidden variables and EM
11	25	all	finding images in digital libraries
12	17	all	tracking, the Kalman filter and data association
13	18	all	model-based vision
14	22	all	finding templates using classifiers
15	20	all	range data

hardest theoretical material is omitted, there is a thorough treatment of the basic geometry and physics of image formation. We assume that students will have a wide range of backgrounds, and can be assigned background readings in probability (we suggest the chapter on the book's web page) around week 2 or 3. We have put off the application chapters to the end, but many may prefer to do chapter 20 around week 10 and chapter 21 around week 6.

Table 3 contains a syllabus for students of computer graphics who want to know the elements of vision that are relevant to their topic. We have emphasized methods that make it possible to recover object models from image information; understanding these topics needs a working knowledge of cameras and filters. Tracking is becoming useful in the graphics world, where it is particularly important for motion capture. We assume that students will have a wide range of backgrounds, and have some exposure to probability.

Table 4 shows a syllabus for students who are primarily interested in the applications of computer vision. We cover material of most immediate practical interest. We assume that students will have a wide range of backgrounds, and can be assigned background reading on probability around week 2 or 3.

Table 5 is a suggested syllabus for students of cognitive science or artificial intelligence who want a basic outline of the important notions of computer vision. This syllabus is less aggressively paced, and assumes less mathematical experience. Students will need to read some material on probability (e.g., the chapter on the book's web page) around week 2 or 3.

Table 6 shows a sample syllabus for students who have a strong interest in applied mathematics, electrical engineering or physics. This syllabus makes for a very busy semester; we move fast, assuming that students can cope with a lot of mathematical material. We assume that students will have a wide range of backgrounds, and can be assigned some reading on probability around week 2 or 3. As a break in a pretty abstract and demanding syllabus, we have inserted a brief review of digital libraries; the chapter on image-based rendering or that on range data could be used instead.

**TABLE 5** For students of cognitive science or artificial intelligence who want a basic outline of the important notions of computer vision.

Week	Chapter	Sections	Key topics
1	1, 4	1, 4 (summary only)	pinhole cameras; lenses; cameras and the eye; radiometric terminology
2	5	all	local shading models; point, line and area sources; photometric stereo; interreflections; lightness computations
3	6	all	color: physics, human perception, spaces; image models; color constancy
4	7	7.1–7.5, 7.7	linear filters; sampling; scale
5	8	all	edge detection
6	9	all	texture; representation, synthesis, shape from
7	10.1, 10.2	all	basic multiple view geometry
8	11	all	stereopsis
9	14	all	segmentation by clustering
10	15	all	fitting lines, curves; robustness; RANSAC
11	16	all	hidden variables and EM
12	18	all	model-based vision
13	22	all	finding templates using classifiers
14	23	all	recognition by relations between templates
15	24	all	geometric templates from spatial relations

**TABLE 6** A syllabus for students who have a strong interest in applied mathematics, electrical engineering or physics.

Week	Chapter	Sections	Key topics
1	1, 4	all	cameras, radiometry
2	5	all	shading models; point, line and area sources; photometric stereo; interreflections and shading primitives
3	6	all	color:—physics, human perception, spaces, color constancy
4	2, 3	all	camera parameters and calibration
5	7, 8	all	linear filters and edge detection
6	8, 9	all	finish edge detection; texture: representation, synthesis, shape from
7	10, 11	all	multiple view geometry, stereopsis as an example
8	12, 13	all	structure from motion
9	14, 15	all	segmentation as clustering; fitting lines, curves; robustness; RANSAC
10	15, 16	all	finish fitting; hidden variables and EM
11	17, 25	all	tracking: Kalman filters, data association; finding images in digital libraries
12	18	all	model-based vision
13	19	all	surfaces and their outlines
14	20	all	aspect graphs
15	22	all	template matching

## NOTATION

We use the following notation throughout the book: points, lines, and planes are denoted by Roman or Greek letters in italic font (e.g.,  $P$ ,  $\Delta$ , or  $\Pi$ ). Vectors are usually denoted by Roman or Greek bold-italic letters (e.g.,  $\mathbf{v}$ ,  $\mathbf{P}$ , or  $\boldsymbol{\xi}$ ), but the vector joining two points  $P$  and  $Q$  is often denoted by  $\overrightarrow{PQ}$ . Lower-case letters are normally used to denote geometric figures in the image plane (e.g.,  $p$ ,  $\mathbf{p}$ ,  $\delta$ ), and upper-case letters are used for scene objects (e.g.,  $P$ ,  $\Pi$ ). Matrices are denoted by Roman letters in calligraphic font (e.g.,  $\mathcal{U}$ ).

The familiar three-dimensional Euclidean space is denoted by  $\mathbb{E}^3$ , and the vector space formed by  $n$ -tuples of real numbers with the usual laws of addition and multiplication by a scalar is denoted by  $\mathbb{R}^n$ , with  $\mathbf{0}$  being used to denote the zero vector. Likewise, the vector space formed by  $m \times n$  matrices with real entries is denoted by  $\mathbb{R}^{m \times n}$ . When  $m = n$ ,  $\text{Id}$  is used to denote the identity matrix—that is, the  $n \times n$  matrix whose diagonal entries are equal to 1 and nondiagonal entries are equal to 0. The transpose of the  $m \times n$  matrix  $\mathcal{U}$  with coefficients  $u_{ij}$  is the  $n \times m$  matrix denoted by  $\mathcal{U}^T$  with coefficients  $u_{ji}$ . Elements of  $\mathbb{R}^n$  are often identified with column vectors or  $n \times 1$  matrices, e.g.,  $\mathbf{a} = (a_1, a_2, a_3)^T$  is the transpose of a  $1 \times 3$  matrix (or *row vector*), i.e., an  $3 \times 1$  matrix (or *column vector*), or equivalently an element of  $\mathbb{R}^3$ .

The *dot product* (or *inner product*) of two vectors  $\mathbf{a} = (a_1, \dots, a_n)^T$  and  $\mathbf{b} = (b_1, \dots, b_n)^T$  in  $\mathbb{R}^n$  is defined by

$$\mathbf{a} \cdot \mathbf{b} = a_1 b_1 + \dots + a_n b_n,$$

and it can also be written as a matrix product, i.e.,  $\mathbf{a} \cdot \mathbf{b} = \mathbf{a}^T \mathbf{b} = \mathbf{b}^T \mathbf{a}$ . We denote by  $|\mathbf{a}|^2 = \mathbf{a} \cdot \mathbf{a}$  the square of the Euclidean norm of the vector  $\mathbf{a}$  and denote by  $d$  the distance function induced

by the Euclidean norm in  $\mathbb{E}^n$ , i.e.,  $d(P, Q) = |\overrightarrow{PQ}|$ . Given a matrix  $U$  in  $\mathbb{R}^{m \times n}$ , we generally use  $|U|$  to denote its *Frobenius norm*, i.e., the square root of the sum of its squared entries.

When the vector  $\mathbf{a}$  has unit norm, the dot product  $\mathbf{a} \cdot \mathbf{b}$  is equal to the (signed) length of the projection of  $\mathbf{b}$  onto  $\mathbf{a}$ . More generally,

$$\mathbf{a} \cdot \mathbf{b} = |\mathbf{a}| |\mathbf{b}| \cos \theta,$$

where  $\theta$  is the angle between the two vectors, which shows that a necessary and sufficient condition for two vectors to be orthogonal is that their dot product be zero.

The *cross product* (or *outer product*) of two vectors  $\mathbf{a} = (a_1, a_2, a_3)^T$  and  $\mathbf{b} = (b_1, b_2, b_3)^T$  in  $\mathbb{R}^3$  is the vector

$$\mathbf{a} \times \mathbf{b} \stackrel{\text{def}}{=} \begin{pmatrix} a_2 b_3 - a_3 b_2 \\ a_3 b_1 - a_1 b_3 \\ a_1 b_2 - a_2 b_1 \end{pmatrix}.$$

Note that  $\mathbf{a} \times \mathbf{b} = [\mathbf{a}_\times] \mathbf{b}$ , where

$$[\mathbf{a}_\times] \stackrel{\text{def}}{=} \begin{pmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{pmatrix}.$$

The cross product of two vectors  $\mathbf{a}$  and  $\mathbf{b}$  in  $\mathbb{R}^3$  is orthogonal to these two vectors, and a necessary and sufficient condition for  $\mathbf{a}$  and  $\mathbf{b}$  to have the same direction is that  $\mathbf{a} \times \mathbf{b} = \mathbf{0}$ . If  $\theta$  denotes as before the angle between the vectors  $\mathbf{a}$  and  $\mathbf{b}$ , it can be shown that

$$|\mathbf{a} \times \mathbf{b}| = |\mathbf{a}| |\mathbf{b}| |\sin \theta|.$$

## PROGRAMMING ASSIGNMENTS AND RESOURCES

The programming assignments given throughout the book sometimes require routines for numerical linear algebra, singular value decomposition, and linear and nonlinear least squares. An extensive set of such routines is available in MATLAB as well as in public-domain libraries such as LINPACK, LAPACK, and MINPACK, which can be downloaded from the Netlib repository (<http://www.netlib.org/>). We offer some pointers to other software on the book's web page <http://www.cs.berkeley.edu/~daf/book.html>. Datasets—or pointers to datasets—for the programming assignment are also available there.

