

AN INTRODUCTION TO 3D COMPUTER VISION TECHNIQUES AND ALGORITHMS

Bogusław Cyganek

Department of Electronics, AGH University of Science and Technology, Poland

J. Paul Siebert

Department of Computing Science, University of Glasgow, Scotland, UK



A John Wiley and Sons, Ltd., Publication

**AN INTRODUCTION
TO 3D COMPUTER
VISION TECHNIQUES
AND ALGORITHMS**

AN INTRODUCTION TO 3D COMPUTER VISION TECHNIQUES AND ALGORITHMS

Bogusław Cyganek

Department of Electronics, AGH University of Science and Technology, Poland

J. Paul Siebert

Department of Computing Science, University of Glasgow, Scotland, UK

 **WILEY**

A John Wiley and Sons, Ltd., Publication

This edition first published 2009
© 2009 John Wiley & Sons, Ltd

Registered office

John Wiley & Sons Ltd, The Atrium, Southern Gate, Chichester, West Sussex, PO19 8SQ, United Kingdom

For details of our global editorial offices, for customer services and for information about how to apply for permission to reuse the copyright material in this book please see our website at www.wiley.com.

The right of the author to be identified as the author of this work has been asserted in accordance with the Copyright, Designs and Patents Act 1988.

All rights reserved. No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, recording or otherwise, except as permitted by the UK Copyright, Designs and Patents Act 1988, without the prior permission of the publisher.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic books.

Designations used by companies to distinguish their products are often claimed as trademarks. All brand names and product names used in this book are trade names, service marks, trademarks or registered trademarks of their respective owners. The publisher is not associated with any product or vendor mentioned in this book. This publication is designed to provide accurate and authoritative information in regard to the subject matter covered. It is sold on the understanding that the publisher is not engaged in rendering professional services. If professional advice or other expert assistance is required, the services of a competent professional should be sought.

Library of Congress Cataloging-in-Publication Data

Cyganek, Boguslaw.

An introduction to 3D computer vision techniques and algorithms / by Boguslaw

Cyganek and J. Paul Siebert.

p. cm.

Includes index.

ISBN 978-0-470-01704-3 (cloth)

1. Computer vision. 2. Three-dimensional imaging. 3. Computer algorithms. I. Siebert,

J. Paul. II. Title

TA1634.C94 2008

006.3'7—dc22

2008032205

A catalogue record for this book is available from the British Library.

ISBN 978-0-470-01704-3

Set in 10/12pt Times by Aptara Inc., New Delhi, India.

Printed in Great Britain by CPI Antony Rowe, Chippenham, Wiltshire

To Magda, Nadia and Kamil
From Bogusław

To Sabina, Konrad and Gustav
From Paul

Contents

Preface	xv
Acknowledgements	xvii
Notation and Abbreviations	xix
Part I	1
1 Introduction	3
1.1 Stereo-pair Images and Depth Perception	4
1.2 3D Vision Systems	4
1.3 3D Vision Applications	5
1.4 Contents Overview: The 3D Vision Task in Stages	6
2 Brief History of Research on Vision	9
2.1 Abstract	9
2.2 Retrospective of Vision Research	9
2.3 Closure	14
2.3.1 <i>Further Reading</i>	14
Part II	15
3 2D and 3D Vision Formation	17
3.1 Abstract	17
3.2 Human Visual System	18
3.3 Geometry and Acquisition of a Single Image	23
3.3.1 <i>Projective Transformation</i>	24
3.3.2 <i>Simple Camera System: the Pin-hole Model</i>	24
3.3.2.1 <i>Extrinsic Parameters</i>	26
3.3.2.2 <i>Intrinsic Parameters</i>	27
3.3.3 <i>Projective Transformation of the Pin-hole Camera</i>	28
3.3.4 <i>Special Camera Setups</i>	29
3.3.5 <i>Parameters of Real Camera Systems</i>	30

3.4	Stereoscopic Acquisition Systems	31
3.4.1	<i>Epipolar Geometry</i>	31
3.4.1.1	<i>Fundamental Matrix</i>	34
3.4.1.2	<i>Epipolar Lines and Epipoles</i>	35
3.4.2	<i>Canonical Stereoscopic System</i>	36
3.4.3	<i>Disparity in the General Case</i>	38
3.4.4	<i>Bifocal, Trifocal and Multifocal Tensors</i>	39
3.4.5	<i>Finding the Essential and Fundamental Matrices</i>	41
3.4.5.1	<i>Point Normalization for the Linear Method</i>	44
3.4.5.2	<i>Computing F in Practice</i>	46
3.4.6	<i>Dealing with Outliers</i>	49
3.4.7	<i>Catadioptric Stereo Systems</i>	54
3.4.8	<i>Image Rectification</i>	55
3.4.9	<i>Depth Resolution in Stereo Setups</i>	59
3.4.10	<i>Stereo Images and Reference Data</i>	61
3.5	Stereo Matching Constraints	66
3.6	Calibration of Cameras	70
3.6.1	<i>Standard Calibration Methods</i>	71
3.6.2	<i>Photometric Calibration</i>	73
3.6.3	<i>Self-calibration</i>	73
3.6.4	<i>Calibration of the Stereo Setup</i>	74
3.7	Practical Examples	75
3.7.1	<i>Image Representation and Basic Structures</i>	75
3.7.1.1	<i>Computer Representation of Pixels</i>	76
3.7.1.2	<i>Representation of Images</i>	78
3.7.1.3	<i>Image Operations</i>	87
3.8	Appendix: Derivation of the Pin-hole Camera Transformation	91
3.9	Closure	93
3.9.1	<i>Further Reading</i>	93
3.9.2	<i>Problems and Exercises</i>	94
4	Low-level Image Processing for Image Matching	95
4.1	Abstract	95
4.2	Basic Concepts	95
4.2.1	<i>Convolution and Filtering</i>	95
4.2.2	<i>Filter Separability</i>	97
4.3	Discrete Averaging	99
4.3.1	<i>Gaussian Filter</i>	100
4.3.2	<i>Binomial Filter</i>	101
4.3.2.1	<i>Specification of the Binomial Filter</i>	101
4.3.2.2	<i>Spectral Properties of the Binomial Filter</i>	102
4.4	Discrete Differentiation	105
4.4.1	<i>Optimized Differentiating Filters</i>	105
4.4.2	<i>Savitzky–Golay Filters</i>	108
4.4.2.1	<i>Generation of Savitzky–Golay Filter Coefficients</i>	114

4.5	Edge Detection	115
4.5.1	<i>Edges from Signal Gradient</i>	117
4.5.2	<i>Edges from the Savitzky–Golay Filter</i>	119
4.5.3	<i>Laplacian of Gaussian</i>	120
4.5.4	<i>Difference of Gaussians</i>	126
4.5.5	<i>Morphological Edge Detector</i>	127
4.6	Structural Tensor	127
4.6.1	<i>Locally Oriented Neighbourhoods in Images</i>	128
4.6.1.1	<i>Local Neighbourhood with Orientation</i>	130
4.6.1.2	<i>Definition of a Local Neighbourhood of Pixels</i>	130
4.6.2	<i>Tensor Representation of Local Neighbourhoods</i>	133
4.6.2.1	<i>2D Structural Tensor</i>	136
4.6.2.2	<i>Computation of the Structural Tensor</i>	140
4.6.3	<i>Multichannel Image Processing with Structural Tensor</i>	143
4.7	Corner Detection	144
4.7.1	<i>The Most Common Corner Detectors</i>	144
4.7.2	<i>Corner Detection with the Structural Tensor</i>	149
4.8	Practical Examples	151
4.8.1	<i>C++ Implementations</i>	151
4.8.1.1	<i>Convolution</i>	151
4.8.1.2	<i>Implementing the Structural Tensor</i>	155
4.8.2	<i>Implementation of the Morphological Operators</i>	157
4.8.3	<i>Examples in Matlab: Computation of the SVD</i>	161
4.9	Closure	162
4.9.1	<i>Further Reading</i>	163
4.9.2	<i>Problems and Exercises</i>	163
5	Scale-space Vision	165
5.1	Abstract	165
5.2	Basic Concepts	165
5.2.1	<i>Context</i>	165
5.2.2	<i>Image Scale</i>	166
5.2.3	<i>Image Matching Over Scale</i>	166
5.3	Constructing a Scale-space	168
5.3.1	<i>Gaussian Scale-space</i>	168
5.3.2	<i>Differential Scale-space</i>	170
5.4	Multi-resolution Pyramids	172
5.4.1	<i>Introducing Multi-resolution Pyramids</i>	172
5.4.2	<i>How to Build Pyramids</i>	175
5.4.3	<i>Constructing Regular Gaussian Pyramids</i>	175
5.4.4	<i>Laplacian of Gaussian Pyramids</i>	177
5.4.5	<i>Expanding Pyramid Levels</i>	178
5.4.6	<i>Semi-pyramids</i>	179
5.5	Practical Examples	181
5.5.1	<i>C++ Examples</i>	181
5.5.1.1	<i>Building the Laplacian and Gaussian Pyramids in C++</i>	181

5.5.2	<i>Matlab Examples</i>	186
5.5.2.1	<i>Building the Gaussian Pyramid in Matlab</i>	190
5.5.2.2	<i>Building the Laplacian of Gaussians Pyramid in Matlab</i>	190
5.6	Closure	191
5.6.1	<i>Chapter Summary</i>	191
5.6.2	<i>Further Reading</i>	191
5.6.3	<i>Problems and Exercises</i>	192
6	Image Matching Algorithms	193
6.1	Abstract	193
6.2	Basic Concepts	193
6.3	Match Measures	194
6.3.1	<i>Distances of Image Regions</i>	194
6.3.2	<i>Matching Distances for Bit Strings</i>	198
6.3.3	<i>Matching Distances for Multichannel Images</i>	199
6.3.3.1	<i>Statistical Distances</i>	201
6.3.4	<i>Measures Based on Theory of Information</i>	202
6.3.5	<i>Histogram Matching</i>	205
6.3.6	<i>Efficient Computations of Distances</i>	206
6.3.7	<i>Nonparametric Image Transformations</i>	209
6.3.7.1	<i>Reduced Census Coding</i>	212
6.3.7.2	<i>Sparse Census Relations</i>	214
6.3.7.3	<i>Fuzzy Relationships Among Pixels</i>	215
6.3.7.4	<i>Implementation of Nonparametric Image Transformations</i>	216
6.3.8	<i>Log-polar Transformation for Image Matching</i>	218
6.4	Computational Aspects of Matching	222
6.4.1	<i>Occlusions</i>	222
6.4.2	<i>Disparity Estimation with Subpixel Accuracy</i>	224
6.4.3	<i>Evaluation Methods for Stereo Algorithms</i>	226
6.5	Diversity of Stereo Matching Methods	229
6.5.1	<i>Structure of Stereo Matching Algorithms</i>	233
6.5.1.1	<i>Aggregation of the Cost Values</i>	234
6.5.1.2	<i>Computation of the Disparity Map</i>	235
6.5.1.3	<i>Disparity Map Postprocessing</i>	237
6.6	Area-based Matching	238
6.6.1	<i>Basic Search Approach</i>	239
6.6.2	<i>Interpreting Match Cost</i>	241
6.6.3	<i>Point-oriented Implementation</i>	245
6.6.4	<i>Disparity-oriented Implementation</i>	250
6.6.5	<i>Complexity of Area-based Matching</i>	256
6.6.6	<i>Disparity Map Cross-checking</i>	257
6.6.7	<i>Area-based Matching in Practice</i>	259
6.6.7.1	<i>Intensity Matching</i>	260
6.6.7.2	<i>Area-based Matching in Nonparametric Image Space</i>	260
6.6.7.3	<i>Area-based Matching with the Structural Tensor</i>	262

6.7	Area-based Elastic Matching	273
6.7.1	<i>Elastic Matching at a Single Scale</i>	273
6.7.1.1	<i>Disparity Match Range</i>	274
6.7.1.2	<i>Search and Subpixel Disparity Estimation</i>	275
6.7.2	<i>Elastic Matching Concept</i>	278
6.7.3	<i>Scale-based Search</i>	280
6.7.4	<i>Coarse-to-fine Matching Over Scale</i>	283
6.7.5	<i>Scale Subdivision</i>	284
6.7.6	<i>Confidence Over Scale</i>	285
6.7.7	<i>Final Multi-resolution Matcher</i>	286
6.8	Feature-based Image Matching	288
6.8.1	<i>Zero-crossing Matching</i>	289
6.8.2	<i>Corner-based Matching</i>	292
6.8.3	<i>Edge-based Matching: The Shirai Method</i>	295
6.9	Gradient-based Matching	296
6.10	Method of Dynamic Programming	298
6.10.1	<i>Dynamic Programming Formulation of the Stereo Problem</i>	301
6.11	Graph Cut Approach	306
6.11.1	<i>Graph Cut Algorithm</i>	306
6.11.1.1	<i>Graphs in Computer Vision</i>	309
6.11.1.2	<i>Optimization on Graphs</i>	310
6.11.2	<i>Stereo as a Voxel Labelling Problem</i>	311
6.11.3	<i>Stereo as a Pixel Labelling Problem</i>	312
6.12	Optical Flow	314
6.13	Practical Examples	318
6.13.1	<i>Stereo Matching Hierarchy in C++</i>	318
6.13.2	<i>Log-polar Transformation</i>	319
6.14	Closure	321
6.14.1	<i>Further Reading</i>	321
6.14.2	<i>Problems and Exercises</i>	322
7	Space Reconstruction and Multiview Integration	323
7.1	Abstract	323
7.2	General 3D Reconstruction	323
7.2.1	<i>Triangulation</i>	324
7.2.2	<i>Reconstruction up to a Scale</i>	325
7.2.3	<i>Reconstruction up to a Projective Transformation</i>	327
7.3	Multiview Integration	329
7.3.1	<i>Implicit Surfaces and Marching Cubes</i>	330
7.3.1.1	<i>Range Map Pre-segmentation</i>	331
7.3.1.2	<i>Volumetric Integration Algorithm Overview</i>	332
7.3.1.3	<i>Hole Filling</i>	332
7.3.1.4	<i>Marching Cubes</i>	333
7.3.1.5	<i>Implementation Considerations</i>	338
7.3.2	<i>Direct Mesh Integration</i>	338

7.4 Closure	342
7.4.1 <i>Further Reading</i>	342
8 Case Examples	343
8.1 Abstract	343
8.2 3D System for Vision-Impaired Persons	343
8.3 Face and Body Modelling	345
8.3.1 <i>Development of Face and Body Capture Systems</i>	345
8.3.2 <i>Imaging Resolution, 3D Resolution and Implications for Applications</i>	346
8.3.3 <i>3D Capture and Analysis Pipeline for Constructing Virtual Humans</i>	350
8.4 Clinical and Veterinary Applications	352
8.4.1 <i>Development of 3D Clinical Photography</i>	352
8.4.2 <i>Clinical Requirements for 3D Imaging</i>	353
8.4.3 <i>Clinical Assessment Based on 3D Surface Anatomy</i>	353
8.4.4 <i>Extraction of Basic 3D Anatomic Measurements</i>	354
8.4.5 <i>Vector Field Surface Analysis by Means of Dense Correspondences</i>	357
8.4.6 <i>Eigenspace Methods</i>	359
8.4.7 <i>Clinical and Veterinary Examples</i>	362
8.4.8 <i>Multimodal 3D Imaging</i>	367
8.5 Movie Restoration	370
8.6 Closure	374
8.6.1 <i>Further Reading</i>	374
Part III	375
9 Basics of the Projective Geometry	377
9.1 Abstract	377
9.2 Homogeneous Coordinates	377
9.3 Point, Line and the Rule of Duality	379
9.4 Point and Line at Infinity	380
9.5 Basics on Conics	382
9.5.1 <i>Conics in \wp^2</i>	382
9.5.1.1 <i>The Dual Conic</i>	383
9.5.1.2 <i>Circular Points</i>	383
9.5.2 <i>Conics in \wp^2</i>	384
9.5.2.1 <i>The Absolute Conic</i>	384
9.5.2.2 <i>The Dual Absolute Conic</i>	385
9.6 Group of Projective Transformations	385
9.6.1 <i>Projective Base</i>	385
9.6.2 <i>Hyperplanes</i>	386
9.6.3 <i>Projective Homographies</i>	386
9.7 Projective Invariants	387
9.8 Closure	388
9.8.1 <i>Further Reading</i>	389

10	Basics of Tensor Calculus for Image Processing	391
10.1	Abstract	391
10.2	Basic Concepts	391
	10.2.1 <i>Linear Operators</i>	392
	10.2.2 <i>Change of Coordinate Systems: Jacobians</i>	393
10.3	Change of a Base	394
10.4	Laws of Tensor Transformations	396
10.5	The Metric Tensor	397
	10.5.1 <i>Covariant and Contravariant Components in a Curvilinear Coordinate System</i>	397
	10.5.2 <i>The First Fundamental Form</i>	399
10.6	Simple Tensor Algebra	399
	10.6.1 <i>Tensor Summation</i>	399
	10.6.2 <i>Tensor Product</i>	400
	10.6.3 <i>Contraction and Tensor Inner Product</i>	400
	10.6.4 <i>Reduction to Principal Axes</i>	400
	10.6.5 <i>Tensor Invariants</i>	401
10.7	Closure	401
	10.7.1 <i>Further Reading</i>	401
11	Distortions and Noise in Images	403
11.1	Abstract	403
11.2	Types and Models of Noise	403
11.3	Generating Noisy Test Images	405
11.4	Generating Random Numbers with Normal Distributions	407
11.5	Closure	408
	11.5.1 <i>Further Reading</i>	408
12	Image Warping Procedures	409
12.1	Abstract	409
12.2	Architecture of the Warping System	409
12.3	Coordinate Transformation Module	410
	12.3.1 <i>Projective and Affine Transformations of a Plane</i>	410
	12.3.2 <i>Polynomial Transformations</i>	411
	12.3.3 <i>Generic Coordinates Mapping</i>	412
12.4	Interpolation of Pixel Values	412
	12.4.1 <i>Bilinear Interpolation</i>	412
	12.4.2 <i>Interpolation of Non-scalar-Valued Pixels</i>	414
12.5	The Warp Engine	414
12.6	Software Model of the Warping Schemes	415
	12.6.1 <i>Coordinate Transformation Hierarchy</i>	415
	12.6.2 <i>Interpolation Hierarchy</i>	416
	12.6.3 <i>Image Warp Hierarchy</i>	416
12.7	Warp Examples	419
12.8	Finding the Linear Transformation from Point Correspondences	420
	12.8.1 <i>Linear Algebra on Images</i>	424

12.9	Closure	427
	12.9.1 <i>Further Reading</i>	428
13	Programming Techniques for Image Processing and Computer Vision	429
13.1	Abstract	429
13.2	Useful Techniques and Methodology	430
	13.2.1 <i>Design and Implementation</i>	430
	13.2.1.1 <i>Comments and Descriptions of 'Ideas'</i>	430
	13.2.1.2 <i>Naming Conventions</i>	431
	13.2.1.3 <i>Unified Modelling Language (UML)</i>	431
	13.2.2 <i>Template Classes</i>	436
	13.2.2.1 <i>Expression Templates</i>	437
	13.2.3 <i>Asserting Code Correctness</i>	438
	13.2.3.1 <i>Programming by Contract</i>	438
	13.2.4 <i>Debugging Issues</i>	440
13.3	Design Patterns	441
	13.3.1 <i>Template Function Objects</i>	441
	13.3.2 <i>Handle-body or Bridge</i>	442
	13.3.3 <i>Composite</i>	445
	13.3.4 <i>Strategy</i>	447
	13.3.5 <i>Class Policies and Traits</i>	448
	13.3.6 <i>Singleton</i>	450
	13.3.7 <i>Proxy</i>	450
	13.3.8 <i>Factory Method</i>	451
	13.3.9 <i>Prototype</i>	452
13.4	Object Lifetime and Memory Management	453
13.5	Image Processing Platforms	455
	13.5.1 <i>Image Processing Libraries</i>	455
	13.5.2 <i>Writing Software for Different Platforms</i>	455
13.6	Closure	456
	13.6.1 <i>Further Reading</i>	456
14	Image Processing Library	457
	References	459
	Index	475

Preface

Recent decades have seen rapidly growing research in many areas of computer science, including computer vision. This comes from the natural interest of researchers as well as demands from industry and society for qualitatively new features to be afforded by computers. One especially desirable capability would be automatic reconstruction and analysis of the surrounding 3D environment and recognition of objects in that space. Effective 3D computer vision methods and implementations would open new possibilities such as automatic navigation of robots and vehicles, scene surveillance and monitoring (which allows automatic recognition of unexpected behaviour of people or other objects, such as cars in everyday traffic), medical reasoning, remote surgery and many, many more.

This book is a result of our long fascination with computers and vision algorithms. It started many years ago as a set of short notes with the only purpose ‘to remember this or that’ or to have a kind of ‘short reference’ just for ourselves. However, as this diary grew with the years we decided to make it available to other people. We hope that it was a good decision! It is our hope that this book facilitates access to this enthralling area, especially for students and young researchers. Our intention is to provide a very concise, though as far as possible complete, overview of the basic concepts of 2D and 3D computer vision. However, the best way to get into the field is to try it oneself! Therefore, in parallel with explaining basic concepts, we provide also a basic programming framework with the hope of making this process easier. We greatly encourage the reader to take the next step and try the techniques in practice.

Bogusław Cyganek, Kraków, Poland
J. Paul Siebert, Glasgow, UK

Acknowledgements

We would like to express our gratitude to all the people who helped in the preparation of this book!

In particular, we are indebted to the whole Wiley team who helped in the preparation of the manuscript. In this group special thanks go to Simone Taylor who believed in this project and made it happen. We would also like to express our gratitude to Sian Andrews, Laura Bell, Liz Benson, Emily Bone, Lucy Bryan, Kate Griffiths, Wendy Hunter, Alex King, Erica Peters, Kathryn Sharples, and Nicky Skinner.

We are also very grateful to the individuals and organizations who agreed to the use of their figures in the book. These are Professor Yuichi Ohta from Tsukuba University, as well as Professor Ryszard Szeliski from Microsoft Research. Likewise we would like to thank Dimensional Imaging Ltd. and Precision 3D Ltd. for use of their images. In this respect we would also like to express our gratitude to Springer Science and Business Media, IEEE Computer Society Press, the IET, Emerald Publishing, the ACM, Maney Publishing and Elsevier Science.

We would also like to thank numerous colleagues from the AGH University of Science and Technology in Kraków. We owe a special debt of gratitude to Professor Ryszard Tadeusiewicz and Professor Kazimierz Wiatr, as well as to Lidia Krawentek for their encouragement and continuous support.

We would also like to thank members of the former Turing Institute in Glasgow (Dr Tim Niblett, Joseph Jin, Dr Peter Mowforth, Dr Colin Urquhart and also Arthur van Hoff) as well as members of the Computer Vision and Graphics Group in the Department of Computing Science, University of Glasgow, for access to and use of their research material (Dr John Patterson, Dr Paul Cockshott, Dr Xiangyang Ju, Dr Yijun Xiao, Dr Zhili Mao, Dr Zhifang Mao (posthumously), Dr J.C Nebel, Dr Tim Boyling, Janet Bowman, Susanne Oehler, Stephen Marshall, Don Whiteford and Colin McLaren). Similarly we would like to thank our collaborators in the Glasgow Dental Hospital and School (Professor Khursheed Moos, Professor Ashraf Ayoub and Dr Balvinder Khambay), Canniesburn Plastic Surgery Unit (Mr Arup Ray), Glasgow, the Department of Statistics (Professor Adrian Bowman and Dr Mitchum Bock), Glasgow University, Professor Donald Hadley, Institute of Neurological Sciences, Southern General Hospital, Glasgow, and also those colleagues formerly at the Silsoe Research Institute (Dr Robin Tillett, Dr Nigel McFarlane and Dr Jerry Wu), Silsoe, UK.

Special thanks are due to Dr Sumitha Balasuriya for use of his Matlab codes and graphs. Particular thanks are due to Professor “Keith” van Rijsbergen and Professor Ray Welland without whose support much of the applied research we report would not have been possible.

We wish to express our special thanks and gratitude to Steve Brett from Pandora Inc. for granting rights to access their software platform.

Some parts of the research for which results are provided in this book were possible due to financial support of the European Commission under RACINE-S (IST-2001-37117) and IP-RACINE (IST-2-511316-IP) as well as Polish funds for scientific research in 2007–2008. Research described in these pages has also been funded by the UK DTI and the EPSRC & BBSRC funding councils, the Chief Scientist Office (Scotland), Wellcome Trust, Smith's Charity, the Cleft Lip and Palate Association, the National Lottery (UK) and the Scottish Office. Their support is greatly appreciated.

Finally, we would like to thank Magda and Sabina for their encouragement, patience and understanding over the three-year period it took to write this book.

Notation and Abbreviations

$I_k(x, y)$	Intensity value of a k -th image at a point with local image coordinates (x, y)
$\overline{I_k(x, y)}$	Average intensity value of a k -th image at a point with local image coordinates (x, y)
I	Identity matrix; image treated as a matrix
P	A vector (a point), matrix, tensor, etc.
$T[\mathbf{I}, \mathbf{P}]$	The Census transformation T for a pixel \mathbf{P} in the image \mathbf{I}
i, j	Free coordinates
d_x, d_y	Displacements (offset) in the x and y directions
$D(\mathbf{p}_l, \mathbf{p}_r)$	Disparity between points \mathbf{p}_l and \mathbf{p}_r
D	Disparity map (a matrix)
$U(x, y)$	Local neighbourhood of pixels around a point (x, y)
O_c	Optical centre point
$\mathbf{P}_c = [X_c, Y_c, Z_c]^T$	Coordinates of a 3D point in the camera coordinate system
Π	Camera plane; a projective plane
$\mathbf{o} = (o_x, o_y)$	Central point of a camera plane
f	Focus length of a camera
b	Base line in a stereo system (a distance between cameras)
h_x, h_y	Physical horizontal and vertical dimensions of a pixel
$\mathbf{P} = [X, Y, Z]^T$	3D point and its coordinates
\wp^n	N -dimensional projective space
$\mathbf{P} = [X_h, Y_h, Z_h, 1]^T$	Homogenous coordinates of a point
M	Camera matrix
M_i	Intrinsic parameters of a camera
M_e	Extrinsic parameters of a camera
E	Essential matrix.
F	Fundamental matrix.
e_i	Epipole in an i -th image
SAD	Sum of absolute differences
SSD	Sum of squared differences
ZSAD	Zero-mean sum of absolute differences
ZSSD	Zero-mean sum of squared differences
ZSSD-N	Zero-mean sum of squared differences, normalized

SCP	Sum of cross products
SCP-N	Sum of cross products, normalized
RMS	Root mean square
RMSE	Root mean square error
$\langle L_{xx}, L_{yy} \rangle$	Code lines from a line L_{xx} to L_{yy}
HVS	Human Visual System
SDK	Software Development Kit
\wedge	logical 'and'
\vee	logical 'or'
LRC	Left-right checking (cross-checking)
OCC	Occlusion constraint
ORD	Point ordering constraint
BMD	Bimodality rule
MGJ	Match goodness jumps
NM	Null method
GT RMS	Ground-truth RMS
WTA	Winner-takes-all
*	Convolution operator

Part I

1

Introduction

The purpose of this text on stereo-based imaging is twofold: it is to give students of computer vision a thorough grounding in the image analysis and projective geometry techniques relevant to the task of recovering three-dimensional (3D) surfaces from stereo-pair images; and to provide a complete reference text for professional researchers in the field of computer vision that encompasses the fundamental mathematics and algorithms that have been applied and developed to allow 3D vision systems to be constructed.

Prior to reviewing the contents of this text, we shall set the context of this book in terms of the underlying objectives and the explanation and design of 3D vision systems. We shall also consider briefly the historical context of optics and vision research that has led to our contemporary understanding of 3D vision.

Here we are specifically considering 3D vision systems that base their operation on acquiring stereo-pair images of a scene and then decoding the depth information *implicitly* captured within the stereo-pair as parallaxes, i.e. relative displacements of the contents of one of the images of the stereo-pair with respect to the other image. This process is termed *stereo-photogrammetry*, i.e. measurement from stereo-pair images. For readers with normal functional binocular vision, the everyday experience of observing the world with both of our eyes results in the perception of the relative distance (depth) to points on the surfaces of objects that enter our field of view. For over a hundred years it has been possible to configure a stereo-pair of cameras to capture stereo-pair images, in a manner analogous to mammalian binocular vision, and thereafter view the developed photographs to observe a miniature 3D scene by means of a stereoscope device (used to present the left and right images of the captured stereo-pair of photographs to the appropriate eye). However, in this scenario it is the brain of the observer that must decode the depth information locked within the stereo-pair and thereby experience the perception of depth. In contrast, in this book we shall present underlying mechanisms by which a computer program can be devised to analyse digitally formatted images captured by a stereo-pair of cameras and thereby recover an *explicit* measurement of distances to points *sampling* surfaces in the imaged field of view. Only by explicitly recovering depth estimates does it become possible to undertake useful tasks such as 3D measurement or reverse engineering of object surfaces as elaborated below. While the science of stereo-photogrammetry is a well-established field and it has indeed been possible to undertake 3D

measurement by means of stereo-pair images using a manually operated measurement device (the stereo-comparator) since the beginning of the twentieth century, we present fully automatic approaches for 3D imaging and measurement in this text.

1.1 Stereo-pair Images and Depth Perception

To appreciate the structure of 3D vision systems based on processing stereo-pair images, it is first necessary to grasp, at least in outline, the most basic principles involved in the formation of stereo-pair images and their subsequent analysis. As outlined above, when we observe a scene with both eyes, an image of the scene is formed on the retina of each eye. However, since our eyes are horizontally displaced with respect to each other, the images thus formed are not identical. In fact this stereo-pair of retinal images contains slight displacements between the relative locations of local parts of the image of the scene with respect to each image of the pair, depending upon how close these local scene components are to the point of *fixation* of the observer's eyes. Accordingly, it is possible to reverse this process and deduce how far away scene components were from the observer according to the magnitude and direction of the parallaxes within the stereo-pairs when they were captured. In order to accomplish this task two things must be determined: firstly, those local parts of one image of the stereo-pair that match the corresponding parts in the other image of the stereo-pair, in order to find the local parallaxes; secondly, the precise geometric properties and configuration of the eyes, or cameras. Accordingly, a process of *calibration* is required to discover the requisite geometric information to allow the imaging process to be inverted and relative distances to surfaces observed in the stereo-pair to be recovered.

1.2 3D Vision Systems

By definition, a stereo-photogrammetry-based 3D vision system will require stereo-pair image acquisition hardware, usually connected to a computer hosting software that automates acquisition control. Multiple stereo-pairs of cameras might be employed to allow all-round coverage of an object or person, e.g. in the context of whole-body scanners. Alternatively, the object to be imaged could be mounted on a computer-controlled turntable and overlapping stereo-pairs captured from a fixed viewpoint for different turntable positions. Accordingly, sequencing capture and image download from multiple cameras can be a complex process, and hence the need for a computer to automate this process.

The stereo-pair acquisition process falls into two categories, active illumination and passive illumination. Active illumination implies that some form of pattern is projected on to the scene to facilitate finding and disambiguating parallaxes (also termed *correspondences* or *disparities*) between the stereo-pair images. Projected patterns often comprise grids or stripes and sometimes these are even colour coded. In an alternative approach, a random speckle texture pattern is projected on to the scene in order to augment the texture already present on imaged surfaces. Speckle projection can also guarantee that that imaged surfaces appear to be randomly textured and are therefore locally uniquely distinguishable and hence able to be matched successfully using certain classes of image matching algorithm. With the advent of 'high-resolution' digital cameras the need for pattern projection has been reduced, since the surface texture naturally present on materials, having even a matte finish, can serve to facilitate

matching stereo-pairs. For example, stereo-pair images of the human face and body can be matched successfully using ordinary studio flash illumination when the pixel sampling density is sufficient to resolve the natural texture of the skin, e.g. skin-pores. A camera resolution of approximately 8–13M pixels is adequate for stereo-pair capture of an area corresponding to the adult face or half-torso.

The acquisition computer may also host the principal 3D vision software components:

- An image matching algorithm to find correspondences between the stereo-pairs.
- Photogrammetry software that will perform system calibration to recover the geometric configuration of the acquisition cameras and perform 3D point reconstruction in world coordinates.
- 3D surface reconstruction software that builds complete manifolds from 3D *point-clouds* captured by each imaging stereo-pair.

3D visualisation facilities are usually also provided to allow the reconstructed surfaces to be displayed, often *draped* with an image to provide a *photorealistic* surface model. At this stage the 3D shape and surface appearance of the imaged object or scene has been captured in explicit digital metric form, ready to feed some subsequent application as described below.

1.3 3D Vision Applications

This book has been motivated in part by the need to provide a manual of techniques to serve the needs of the computer vision practitioner who wishes to construct 3D imaging systems configured to meet the needs of practical applications. A wide variety of applications are now emerging which rely on the fast, efficient and low-cost capture of 3D surface information. The traditional role for image-based 3D surface measurement has been the reserve of *close-range* photogrammetry systems, capable of recovering surface measurements from objects in the range of a few tens of millimetres to a few metres in size. A typical example of a classical close-range photogrammetry task might comprise surface measurement for manufacturing quality control, applied to high-precision engineered products such as aircraft wings.

Close-range video-based photogrammetry, having a lower spatial resolution than traditional plate-camera film-based systems, initially found a niche in imaging the human face and body for clinical and creative media applications. 3D clinical photographs have the potential to provide quantitative measurements that reduce subjectivity in assessing the surface anatomy of a patient (or animal) before and after surgical intervention by providing numeric, possibly automated, scores for the shape, symmetry and longitudinal change of anatomic structures. Creative media applications include whole-body 3D imaging to support creation of human avatars of specific individuals, for 3D gaming and cine special effects requiring virtual actors. Clothing applications include body or foot scanning for the production of custom clothing and shoes or as a means of sizing customers accurately. An innovative commercial application comprises a ‘virtual catwalk’ to allow customers to visualize themselves in clothing prior to purchasing such goods on-line via the Internet.

There are very many more emerging uses for 3D imaging beyond the above and commercial ‘reverse engineering’ of premanufactured goods. 3D vision systems have the potential to revolutionize autonomous vehicles and the capabilities of robot vision systems. Stereo-pair cameras could be mounted on a vehicle to facilitate autonomous navigation or configured

within a robot workcell to endow a ‘blind’ pick-and-place robot, both object recognition capabilities based on 3D cues and simultaneously 3D spatial quantification of object locations in the workspace.

1.4 Contents Overview: The 3D Vision Task in Stages

The organization of this book reflects the underlying principles, structural components and uses of 3D vision systems as outlined above, starting with a brief historical view of vision research in Chapter 2. We deal with the basic existence proof that binocular 3D vision is possible, in an overview of the human visual system in Chapter 3. The basic projective geometry techniques that underpin 3D vision systems are also covered here, including the geometry of monocular and binocular image formation which relates how binocular parallaxes are produced in stereo-pair images as a result of imaging scenes containing variation in depth. Camera calibration techniques are also presented in Chapter 3, completing the introduction of the role of image formation and geometry in the context of 3D vision systems.

We deal with fundamental 2D image analysis techniques required to undertake image filtering and feature detection and localization in Chapter 4. These topics serve as a precursor to perform image matching, the process of detecting and quantifying parallaxes between stereo-pair images, a prerequisite to recovering depth information. In Chapter 5 the issue of spatial scale in images is explored, namely how to structure algorithms capable of efficiently processing images containing structures of varying scales which are unknown in advance. Here the concept of an image *scale-space* and the *multi-resolution image pyramid* data structure is presented, analysed and explored as a precursor to developing matching algorithms capable of operating over a wide range of visual scales. The core algorithmic issues associated with stereo-pair image matching are contained in Chapter 6 dealing with distance measures for comparing image patches, the associated parametric issues for matching and an in-depth analysis of area-based matching over scale-space within a practical matching algorithm. Feature-based approaches to matching are also considered and their combination with area-based approaches. Then two solutions to the stereo problem are discussed: the first, based on the *dynamic programming*, and the second one based on the *graph cuts* method. The chapter ends with discussion of the *optical flow* methods which allow estimation of local displacements in a sequence of images.

Having dealt with the recovery of disparities between stereo-pairs, we progress logically to the recovery of 3D surface information in Chapter 7. We consider the process of *triangulation* whereby 3D points in world coordinates are computed from the disparities recovered in the previous chapter. These 3D points can then be organized into surfaces represented by *polygonal meshes* and the *3D point-clouds* recovered from *multi-view* systems acquiring more than one stereo-pair of the scene can be fused into a coherent surface model either directly or via volumetric techniques such as *marching cubes*. In Chapter 8 we conclude the progression from theory to practice, with a number of case examples of 3D vision applications covering areas such as face and body imaging for clinical, veterinary and creative media applications and also 3D vision as a visual prosthetic. An application based only on image matching is also presented that utilizes motion-induced inter-frame disparities within a cine sequence to synthesize missing or damaged frames, or sets of frames, in digitized historic archive footage.

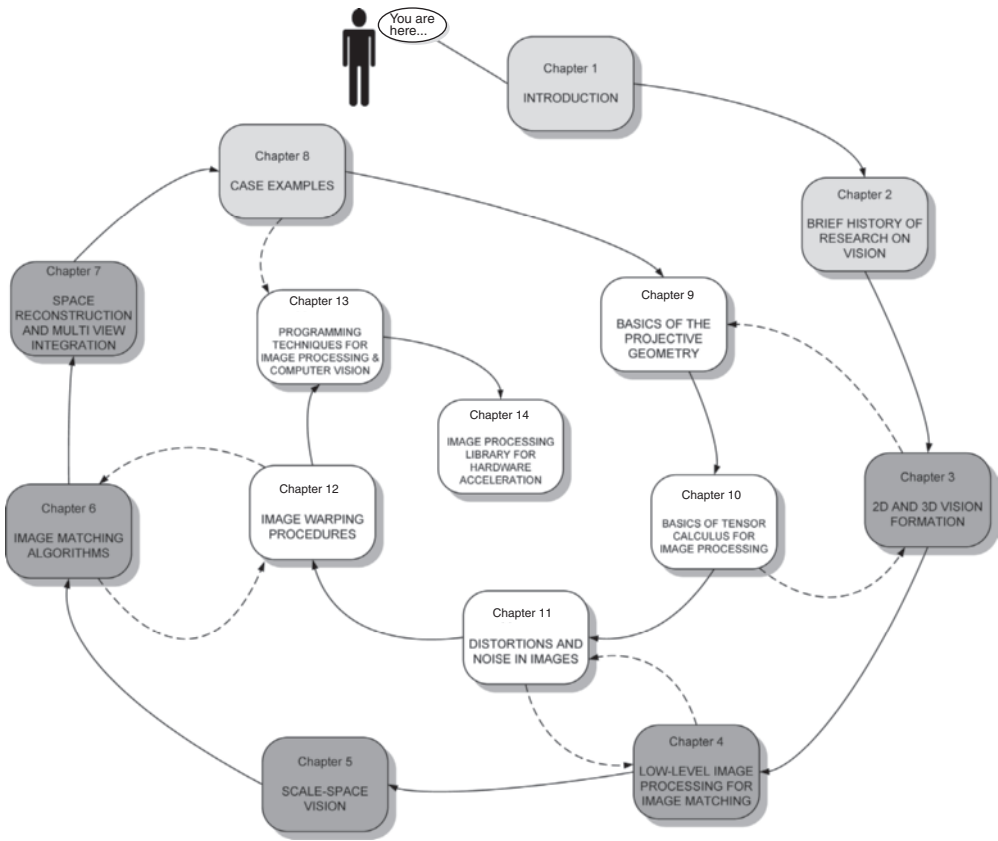


Figure 1.1 Organization of the book

The remaining chapters provide a series of detailed technical tutorials on projective geometry, tensor calculus, image warping procedures and image noise. A chapter on programming techniques for image processing provides practical hints and advice for persons who wish to develop their own computer vision applications. Methods of object oriented programming, such as design patterns, but also proper organization and verification of the code are discussed. Chapter 14 outlines the software presented in the book and provides the link to the recent version of the code.

Figure 1.1 depicts possible order of reading the book. All chapters can be read in number order or selectively as references to specific topics. There are five main chapters (Chapters 3–7), three auxiliary chapters (Chapters 1, 2 and 8) as well as five technical tutorials (Chapters 9–13). The latter are intended to aid understanding of specific topics and can be read in conjunction with the related main chapters, as indicated by the dashed lines in Figure 1.1.

