




Kenneth Dawson-Howe



A Practical Introduction to
**COMPUTER VISION
WITH OPENCV**

WILEY

**A PRACTICAL
INTRODUCTION TO
COMPUTER VISION
WITH OPENCV**

A PRACTICAL INTRODUCTION TO COMPUTER VISION WITH OPENCV

Kenneth Dawson-Howe

Trinity College Dublin, Ireland

WILEY

This edition first published 2014

© 2014 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.

Limit of Liability/Disclaimer of Warranty: While the publisher and author have used their best efforts in preparing this book, they make no representations or warranties with respect to the accuracy or completeness of the contents of this book and specifically disclaim any implied warranties of merchantability or fitness for a particular purpose. It is sold on the understanding that the publisher is not engaged in rendering professional services and neither the publisher nor the author shall be liable for damages arising herefrom. 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 applied for.

ISBN: 9781118848456

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

I am grateful to many people for their help and support during the writing of this book. The biggest thanks must go to my wife Jane, my children, William and Susie, and my parents, all of whose encouragement has been unstinting.

I must express my thanks to my students for their interest and enthusiasm in this subject. It is always refreshing to hear students discussing how to solve vision problems in tutorials and great to hear their solutions to problems which are often different (and sometimes better) than my own.

I thank my colleagues (in particular Arthur Hughes, Jeremy Jones and Hilary McDonald) for their encouragement and support.

Contents

Preface		xiii
1	Introduction	1
1.1	A Difficult Problem	1
1.2	The Human Vision System	2
1.3	Practical Applications of Computer Vision	3
1.4	The Future of Computer Vision	5
1.5	Material in This Textbook	6
1.6	Going Further with Computer Vision	7
2	Images	9
2.1	Cameras	9
	2.1.1 <i>The Simple Pinhole Camera Model</i>	9
2.2	Images	10
	2.2.1 <i>Sampling</i>	11
	2.2.2 <i>Quantisation</i>	11
2.3	Colour Images	13
	2.3.1 <i>Red–Green–Blue (RGB) Images</i>	14
	2.3.2 <i>Cyan–Magenta–Yellow (CMY) Images</i>	17
	2.3.3 <i>YUV Images</i>	17
	2.3.4 <i>Hue Luminance Saturation (HLS) Images</i>	18
	2.3.5 <i>Other Colour Spaces</i>	20
	2.3.6 <i>Some Colour Applications</i>	20
2.4	Noise	22
	2.4.1 <i>Types of Noise</i>	23
	2.4.2 <i>Noise Models</i>	25
	2.4.3 <i>Noise Generation</i>	26
	2.4.4 <i>Noise Evaluation</i>	26
2.5	Smoothing	27
	2.5.1 <i>Image Averaging</i>	27
	2.5.2 <i>Local Averaging and Gaussian Smoothing</i>	28
	2.5.3 <i>Rotating Mask</i>	30
	2.5.4 <i>Median Filter</i>	31

3	Histograms	35
3.1	1D Histograms	35
	3.1.1 <i>Histogram Smoothing</i>	36
	3.1.2 <i>Colour Histograms</i>	37
3.2	3D Histograms	39
3.3	Histogram/Image Equalisation	40
3.4	Histogram Comparison	41
3.5	Back-projection	43
3.6	k-means Clustering	44
4	Binary Vision	49
4.1	Thresholding	49
	4.1.1 <i>Thresholding Problems</i>	50
4.2	Threshold Detection Methods	51
	4.2.1 <i>Bimodal Histogram Analysis</i>	52
	4.2.2 <i>Optimal Thresholding</i>	52
	4.2.3 <i>Otsu Thresholding</i>	54
4.3	Variations on Thresholding	56
	4.3.1 <i>Adaptive Thresholding</i>	56
	4.3.2 <i>Band Thresholding</i>	57
	4.3.3 <i>Semi-thresholding</i>	58
	4.3.4 <i>Multispectral Thresholding</i>	58
4.4	Mathematical Morphology	59
	4.4.1 <i>Dilation</i>	60
	4.4.2 <i>Erosion</i>	62
	4.4.3 <i>Opening and Closing</i>	63
	4.4.4 <i>Grey-scale and Colour Morphology</i>	65
4.5	Connectivity	66
	4.5.1 <i>Connectedness: Paradoxes and Solutions</i>	66
	4.5.2 <i>Connected Components Analysis</i>	67
5	Geometric Transformations	71
5.1	Problem Specification and Algorithm	71
5.2	Affine Transformations	73
	5.2.1 <i>Known Affine Transformations</i>	74
	5.2.2 <i>Unknown Affine Transformations</i>	75
5.3	Perspective Transformations	76
5.4	Specification of More Complex Transformations	78
5.5	Interpolation	78
	5.5.1 <i>Nearest Neighbour Interpolation</i>	79
	5.5.2 <i>Bilinear Interpolation</i>	79
	5.5.3 <i>Bi-Cubic Interpolation</i>	80
5.6	Modelling and Removing Distortion from Cameras	80
	5.6.1 <i>Camera Distortions</i>	81
	5.6.2 <i>Camera Calibration and Removing Distortion</i>	82

6	Edges	83
6.1	Edge Detection	83
6.1.1	<i>First Derivative Edge Detectors</i>	85
6.1.2	<i>Second Derivative Edge Detectors</i>	92
6.1.3	<i>Multispectral Edge Detection</i>	97
6.1.4	<i>Image Sharpening</i>	98
6.2	Contour Segmentation	99
6.2.1	<i>Basic Representations of Edge Data</i>	99
6.2.2	<i>Border Detection</i>	102
6.2.3	<i>Extracting Line Segment Representations of Edge Contours</i>	105
6.3	Hough Transform	108
6.3.1	<i>Hough for Lines</i>	109
6.3.2	<i>Hough for Circles</i>	111
6.3.3	<i>Generalised Hough</i>	112
7	Features	115
7.1	Moravec Corner Detection	117
7.2	Harris Corner Detection	118
7.3	FAST Corner Detection	121
7.4	SIFT	122
7.4.1	<i>Scale Space Extrema Detection</i>	123
7.4.2	<i>Accurate Keypoint Location</i>	124
7.4.3	<i>Keypoint Orientation Assignment</i>	126
7.4.4	<i>Keypoint Descriptor</i>	127
7.4.5	<i>Matching Keypoints</i>	127
7.4.6	<i>Recognition</i>	127
7.5	Other Detectors	129
7.5.1	<i>Minimum Eigenvalues</i>	130
7.5.2	<i>SURF</i>	130
8	Recognition	131
8.1	Template Matching	131
8.1.1	<i>Applications</i>	131
8.1.2	<i>Template Matching Algorithm</i>	133
8.1.3	<i>Matching Metrics</i>	134
8.1.4	<i>Finding Local Maxima or Minima</i>	135
8.1.5	<i>Control Strategies for Matching</i>	137
8.2	Chamfer Matching	137
8.2.1	<i>Chamfering Algorithm</i>	137
8.2.2	<i>Chamfer Matching Algorithm</i>	139
8.3	Statistical Pattern Recognition	140
8.3.1	<i>Probability Review</i>	142
8.3.2	<i>Sample Features</i>	143
8.3.3	<i>Statistical Pattern Recognition Technique</i>	149
8.4	Cascade of Haar Classifiers	152
8.4.1	<i>Features</i>	154
8.4.2	<i>Training</i>	156

8.4.3	<i>Classifiers</i>	156
8.4.4	<i>Recognition</i>	158
8.5	Other Recognition Techniques	158
8.5.1	<i>Support Vector Machines (SVM)</i>	158
8.5.2	<i>Histogram of Oriented Gradients (HoG)</i>	159
8.6	Performance	160
8.6.1	<i>Image and Video Datasets</i>	160
8.6.2	<i>Ground Truth</i>	161
8.6.3	<i>Metrics for Assessing Classification Performance</i>	162
8.6.4	<i>Improving Computation Time</i>	165
9	Video	167
9.1	Moving Object Detection	167
9.1.1	<i>Object of Interest</i>	168
9.1.2	<i>Common Problems</i>	168
9.1.3	<i>Difference Images</i>	169
9.1.4	<i>Background Models</i>	171
9.1.5	<i>Shadow Detection</i>	179
9.2	Tracking	180
9.2.1	<i>Exhaustive Search</i>	181
9.2.2	<i>Mean Shift</i>	181
9.2.3	<i>Dense Optical Flow</i>	182
9.2.4	<i>Feature Based Optical Flow</i>	185
9.3	Performance	186
9.3.1	<i>Video Datasets (and Formats)</i>	186
9.3.2	<i>Metrics for Assessing Video Tracking Performance</i>	187
10	Vision Problems	189
10.1	Baby Food	189
10.2	Labels on Glue	190
10.3	O-rings	191
10.4	Staying in Lane	192
10.5	Reading Notices	193
10.6	Mailboxes	194
10.7	Abandoned and Removed Object Detection	195
10.8	Surveillance	196
10.9	Traffic Lights	197
10.10	Real Time Face Tracking	198
10.11	Playing Pool	199
10.12	Open Windows	200
10.13	Modelling Doors	201
10.14	Determining the Time from Analogue Clocks	202
10.15	Which Page	203
10.16	Nut/Bolt/Washer Classification	204
10.17	Road Sign Recognition	205
10.18	License Plates	206

10.19	Counting Bicycles	207
10.20	Recognise Paintings	208
	References	209
	Index	213

Preface

Perception is essential in order for any entity to interact in a meaningful way with its environment. Humans draw on many senses (such as sight, sound, touch and smell) to perceive the world. Most machines can only receive input through simple input devices, such as keyboards and mice, or through wired and wireless communication channels. However, in recent years, cameras and microphones have been added as standard parts of computers and mobile devices (such as phones and tablets). At the same time, the speed of these devices has increased significantly, making it possible to start to process this data in a meaningful manner. Computer Vision is about how we can automate image or video understanding on machines. It covers the techniques used to automate tasks ranging from industrial inspection (where the image understanding problem is constrained to one which we could easily address 20 years ago) to video understanding in order to guide autonomous robots so that they can interact in a meaningful and safe manner in a world designed for humans.

This book provides a brief introduction to this exciting field, covering the basics of image processing and providing the reader with enough information to solve many practical problems. Computer vision systems are becoming ubiquitous. They are in our homes (in the interfaces of the games consoles which our children use), in our cameras and phones (providing automatic face detection and red eye removal), on our streets (determining the licence plates of vehicles passing through toll gates), in our offices (providing biometric verification of identity), and even more so in our factories, helping to guide robots to manufacture goods (such as cars) and automatically inspecting goods to ensure they look right. Yet it seems that we are only at the beginning of how computer vision can be employed, and we can expect significantly more vision systems to emerge.

For those interested in this field as developers (and that hopefully includes you as you are reading this book) there is very good news as there are a number of high quality systems in which computer vision solutions can be developed, of which two stand out in particular: MATLAB[®] and OpenCV. MATLAB[®] provides an environment that allows relatively rapid prototyping of vision solutions. OpenCV is a high quality library for C and C++, with wrappers for Python and Java (on Windows, Linux, MacOS, FreeBSD, OpenBSD, Android, Maemo and iOS), which provides implementations of many state-of-the-art vision techniques. OpenCV is the platform of choice for many vision developers, is developed collaboratively by the vision community and is available free of charge for educational and commercial use. OpenCV code snippets are provided throughout this book so that readers can easily take the theory and easily create working solutions to vision problems.

This text is intended to:

1. Provide a solid academic background to basic computer vision.
2. Provide enough material for a one-semester course in computer vision. Larger, all encompassing, textbooks are very off-putting to new students in this (or any) field.
3. Facilitate practical use of computer vision. The goal is to bridge the gap between the theory and the practical implementation of computer vision and so explanations of how to use the relevant OpenCV library routines are included, accompanied by a full working program including the code snippets from the text in the website mentioned below.
4. Allow students to solve real practical problems by providing images and videos for the 20 application problems in Chapter 10.

Electronic Resources

The electronic resources which accompany this text include:

- the code examples from the text along with images generated from the code to give an idea of the processing done by each section of the code.
- Powerpoint slides for each of the chapters.
- the media (images and videos) for each of the application problems in Chapter 10 of the book.
- links to information on OpenCV.

The resources are available at <https://www.scss.tcd.ie/publications/book-supplements/A-Practical-Introduction-to-Computer-Vision-with-OpenCV> and at (a shorter alternative which redirects to the above page) <https://www.scss.tcd.ie/Kenneth.Dawson-Howe/PracticalVisionBook>

Teaching Computer Vision Using This Text

A computer vision course based on this text would consist of around 22–28 one-hour lectures together with tutorials and labs. Anticipated lecture hours by topic are as follows:

- Introduction: 1–2 hours
- Images (2.1 Cameras – 2.3 Colour images): 2 hours
- Images (2.4 Noise – 2.5 Smoothing): 2 hours
- Histograms (Chapter 3): 2 hours
- Binary Vision (Chapter 4): 2–3 hours
- Geometric Transformations (Chapter 5): 1–2 hours
- Edges (6.1 Edge detection): 2 hours
- Edges (6.2 Contour segmentation): 1–2 hours
- Edges (6.3 Hough transform): 1–2 hours
- Features (7.1 Moravec corner detection – 7.3 FAST corner detection): 1 hour
- Features (7.4 SIFT): 1 hour
- Recognition (8.1 Template matching and 8.2 Chamfer matching): 1 hour

- Recognition (8.3 Statistical pattern recognition): 1 hour
- Recognition (8.4 Cascade of Haar classifiers): 1 hour
- Recognition (8.6 Performance): 1 hour
- Video (Chapter 9): 2–3 hours

For tutorials, it is suggested that the class be broken into groups of three or four students (all in a single large venue) and that the groups should be asked to come up with solutions to some of the vision problems in Chapter 10 (using the vision techniques they have learnt). The intention is that the students discuss how to solve the problems, coming up with ways of combining the techniques that they have learnt in order to solve them. There is more than one solution to all of the problems, so some of the groups should present their solutions to the class, and the class and lecturer should discuss how appropriate the solutions are. For labs and assignments, the same problems can be used, as OpenCV provides the functionality to allow students to prototype solutions to these problems.

1

Introduction

Computer vision is the automatic analysis of images and videos by computers in order to gain some understanding of the world. Computer vision is inspired by the capabilities of the human vision system and, when initially addressed in the 1960s and 1970s, it was thought to be a relatively straightforward problem to solve. However, the reason we think/thought that vision is easy is that we have our own visual system which makes the task seem intuitive to our conscious minds. In fact, the human visual system is very complex and even the estimates of how much of the brain is involved with visual processing vary from 25% up to more than 50%.

1.1 A Difficult Problem

The first challenge facing anyone studying this subject is to convince themselves that the problem is difficult. To try to illustrate the difficulty, we first show three different versions of the same image in Figure 1.1. For a computer, an image is just an array of values, such as the array shown in the left-hand image in Figure 1.1. For us, using our complex vision system, we can perceive this as a face image but only if we are shown it as a grey scale image (top right).

Computer vision is quite like understanding the array of values shown in Figure 1.1, but is more complicated as the array is really much bigger (e.g. to be equivalent to the human eye a camera would need around 127 million elements), and more complex (i.e. with each point represented by three values in order to encode colour information). To make the task even more convoluted, the images are constantly changing, providing a stream of 50–60 images per second and, of course, there are two streams of data as we have two eyes/cameras.

Another illustration of the difficulty of vision was provided by psychologist John Wilding considering his own visual system:

As I look out of my window, I see grass and trees, gently swaying in the wind, with a lake beyond . . . An asphalt path leads down through the trees to the lake and two squirrels are chasing each other to and fro across it, ignoring the woman coming up the path . . .

67	67	66	68	66	67	64	65	65	63	63	69	61	64	63	66	61	60
69	68	63	68	65	62	65	61	50	26	32	65	61	67	64	65	66	63
72	71	70	87	67	60	28	21	17	18	13	15	20	59	61	65	66	64
75	73	76	78	67	26	20	19	16	18	16	13	18	21	50	61	69	70
74	75	78	74	39	31	31	30	46	37	69	66	64	43	18	63	69	60
73	75	77	64	41	20	18	22	63	92	99	88	78	73	39	40	59	65
74	75	71	42	19	12	14	28	79	102	107	96	87	79	57	29	68	66
75	75	66	43	12	11	16	62	87	84	84	108	83	84	59	39	70	66
76	74	49	42	37	10	34	78	90	99	68	94	97	51	40	69	72	65
76	63	40	57	123	88	60	83	95	88	80	71	67	69	32	67	73	73
78	50	32	33	90	121	66	86	100	116	87	85	80	74	71	56	58	48
80	40	33	16	63	107	57	86	103	113	113	104	94	86	77	48	47	45
88	41	35	10	15	94	67	96	98	91	86	105	81	77	71	35	45	47
87	51	35	15	15	17	51	92	104	101	72	74	87	100	27	31	44	46
86	42	47	11	13	16	71	76	89	95	116	91	67	87	12	25	43	51
96	67	20	12	17	17	86	89	90	101	96	89	62	13	11	19	40	51
99	88	19	15	15	18	32	107	99	86	95	92	26	13	13	16	49	52
99	77	16	14	14	16	35	115	111	109	91	79	17	16	13	46	48	51




Figure 1.1 Different versions of an image. An array of numbers (left) which are the values of the grey scales in the low resolution image of a face (top right). The task of computer vision is most like understanding the array of numbers

This is the scene I experience, a world of objects with background, acted upon and sometimes acting and interacting in events. I have no problem seeing and hearing and smelling and feeling all these things because they affect my senses directly and they make up the real world.

Or do they? I can look again and notice things I missed before, or see the scene in new ways. There is a white wall framing the window I am looking through and the window in fact fills less of my field of view than the wall, but I did not even notice the wall at first, and my impression was that the scene through the window was a panorama right across in front of me. There are metal bars dividing the window into squares and the glass is obscured with dust and spots but for me the view seems complete and un-obscured. The 'grass' is patches of colour ranging from nearly white in the bright sun to nearly black in the shade but I 'saw' green grass in light and shade. Other changing greenish shapes were for me permanent leafy branches moved by a wind I neither saw nor felt, and two constantly varying grey shapes were squirrels moving with a purpose. Another shape increasing in size and changing in position was an approaching woman. (Wilding, 1983)

1.2 The Human Vision System

If we could duplicate the human visual system then the problem of developing a computer vision system would be solved. So why can't we? The main difficulty is that we do not understand what the human vision system is doing most of the time.

If you consider your eyes, it is probably not clear to you that your colour vision (provided by the 6–7 million cones in the eye) is concentrated in the centre of the visual field of the eye (known as the macula). The rest of your retina is made up of around 120 million rods (cells that are sensitive to visible light of any wavelength/colour). In addition, each eye has a rather large blind spot where the optic nerve attaches to the retina. Somehow, we think we see a continuous image (i.e. no blind spot) with colour everywhere, but even at this lowest level of processing it is unclear as to how this impression occurs within the brain.

The visual cortex (at the back of the brain) has been studied and found to contain cells that perform a type of edge detection (see Chapter 6), but mostly we know what sections of the brain do based on localised brain damage to individuals. For example, a number of people with damage to a particular section of the brain can no longer recognise faces (a condition known as prosopagnosia). Other people have lost the ability to sense moving objects (a condition known as akinetopsia). These conditions inspire us to develop separate modules to recognise faces (e.g. see Section 8.4) and to detect object motion (e.g. see Chapter 9).

We can also look at the brain using functional MRI, which allows us to see the concentration of electrical activity in different parts of the brain as subjects perform various activities. Again, this may tell us what large parts of the brain are doing, but it cannot provide us with algorithms to solve the problem of interpreting the massive arrays of numbers that video cameras provide.

1.3 Practical Applications of Computer Vision

Computer vision has many applications in industry, particularly allowing the automatic inspection of manufactured goods at any stage in the production line. For example, it has been used to:

- Inspect printed circuits boards to ensure that tracks and components are placed correctly. See Figure 1.2.
- Inspect print quality of labels. See Figure 1.3.
- Inspect bottles to ensure they are properly filled. See Figure 1.3.




Figure 1.2 PCB inspection of pads (left) and images of some detected flaws in the surface mounting of components (right). Reproduced by permission of James Mahon



Figure 1.3 Checking print quality of best-before dates (right), and monitoring level to which bottles are filled (right). Reproduced by permission of Omron Electronics LLC

- Inspect apples to determine if there is any bruising.
- Locate chocolates on a production line so that a robot arm can pick them up and place them in the correct locations in the box.
- Guide robots when manufacturing complex products such as cars.

On the factory floor, the problem is a little simpler than in the real world as the lighting can be constrained and the possible variations of what we can see are quite limited. Computer vision is now solving problems outside the factory. Computer vision applications outside the factory include:

- The automatic reading of license plates as they pass through tollgates on major roads.
- Augmenting sports broadcasts by determining distances for penalties, along with a range of other statistics (such as how far each player has travelled during the game).
- Biometric security checks in airports using images of faces and images of fingerprints. See Figure 1.4.
- Augmenting movies by the insertion of virtual objects into video sequences, so that they appear as though they belong (e.g. the candles in the Great Hall in the Harry Potter movies).

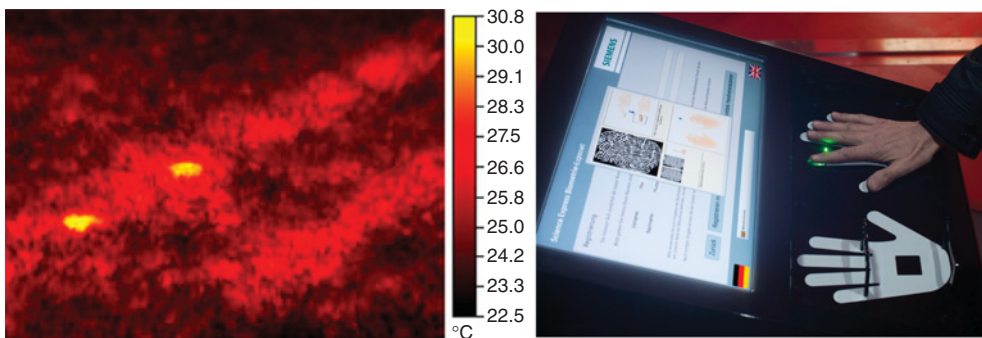


Figure 1.4 Buried landmines in an infrared image (left). Reproduced by permission of Zouheir Fawaz, Handprint recognition system (right). Reproduced by permission of Siemens AG

- Assisting drivers by warning them when they are drifting out of lane.
- Creating 3D models of a destroyed building from multiple old photographs.
- Advanced interfaces for computer games allowing the real time detection of players or their hand-held controllers.
- Classification of plant types and anticipated yields based on multispectral satellite images.
- Detecting buried landmines in infrared images. See Figure 1.4.

Some examples of existing computer vision systems in the outside world are shown in Figure 1.4.

1.4 The Future of Computer Vision

The community of vision developers is constantly pushing the boundaries of what we can achieve. While we can produce autonomous vehicles, which drive themselves on a highway, we would have difficulties producing a reliable vehicle to work on minor roads, particularly if the road marking were poor. Even in the highway environment, though, we have a legal issue, as who is to blame if the vehicle crashes? Clearly, those developing the technology do not think it should be them and would rather that the driver should still be responsible should anything go wrong. This issue of liability is a difficult one and arises with many vision applications in the real world. Taking another example, if we develop a medical imaging system to diagnose cancer, what will happen when it mistakenly does not diagnose a condition? Even though the system might be more reliable than any individual radiologist, we enter a legal minefield. Therefore, for now, the simplest solution is either to address only non-critical problems or to develop systems, which are assistants to, rather than replacements for, the current human experts.

Another problem exists with the deployment of computer vision systems. In some countries the installation and use of video cameras is considered an infringement of our basic right to privacy. This varies hugely from country to country, from company to company, and even from individual to individual. While most people involved with technology see the potential benefits of camera systems, many people are inherently distrustful of video cameras and what the videos *could* be used for. Among other things, they fear (perhaps justifiably) a Big Brother scenario, where our movements and actions are constantly monitored. Despite this, the number of cameras is growing very rapidly, as there are cameras on virtually every new computer, every new phone, every new games console, and so on.

Moving forwards, we expect to see computer vision addressing progressively harder problems; that is problems in more complex environments with fewer constraints. We expect computer vision to start to be able to recognise more objects of different types and to begin to extract more reliable and robust descriptions of the world in which they operate. For example, we expect computer vision to

- become an integral part of general computer interfaces;
- provide increased levels of security through biometric analysis;
- provide reliable diagnoses of medical conditions from medical images and medical records;
- allow vehicles to be driven autonomously;
- automatically determine the identity of criminals through the forensic analysis of video.




Figure 1.5 The ASIMO humanoid robot which has two cameras in its ‘head’ which allow ASIMO to determine how far away things are, recognise familiar faces, etc. Reproduced by permission of Honda Motor Co. Inc

Ultimately, computer vision is aiming to emulate the capabilities of human vision, and to provide these abilities to humanoid (and other) robotic devices, such as ASIMO (see Figure 1.5). This is part of what makes this field exciting, and surprising, as we all have our own (human) vision systems which work remarkably well, yet when we try to automate any computer vision task it proves very difficult to do reliably.

1.5 Material in this Textbook

This textbook is intended to provide an illustrated introduction to the area of computer vision. It provides roughly the amount of material which can be covered in a one-semester, year four or five university course. While this text covers the theory behind basic computer vision, it also provides a bridge from the theory to practical implementation using the industry standard OpenCV libraries (by explaining how the operations can be invoked in OpenCV).

In Chapter 2, we consider the basics of cameras and images, along with consideration of the noise that is exhibited by many images and the techniques through which this noise can be removed or attenuated.

In Chapter 3, we consider how image information can be summarised in the form of a histogram and how those histograms can be used in order to enhance images or to extract information from the images.

Chapter 4 looks at the most commonly used technique for industrial vision – that of binary vision, where we simplify images so that every point is either black or white. This approach makes the processing much easier (assuming that the binary image can be obtained correctly).

Chapter 5 looks at how we model and remove distortion from images and cameras (which is often introduced by the camera/lens system).