

Fourth Edition


COMPUTER & MACHINE VISION

Theory
Algorithms
Practicalities

E.R. Davies



Computer and
Machine Vision:
Theory, Algorithms,
Practicalities



This book is dedicated to my family.

*To my late mother, Mary Davies, to record her never-failing
love and devotion.*

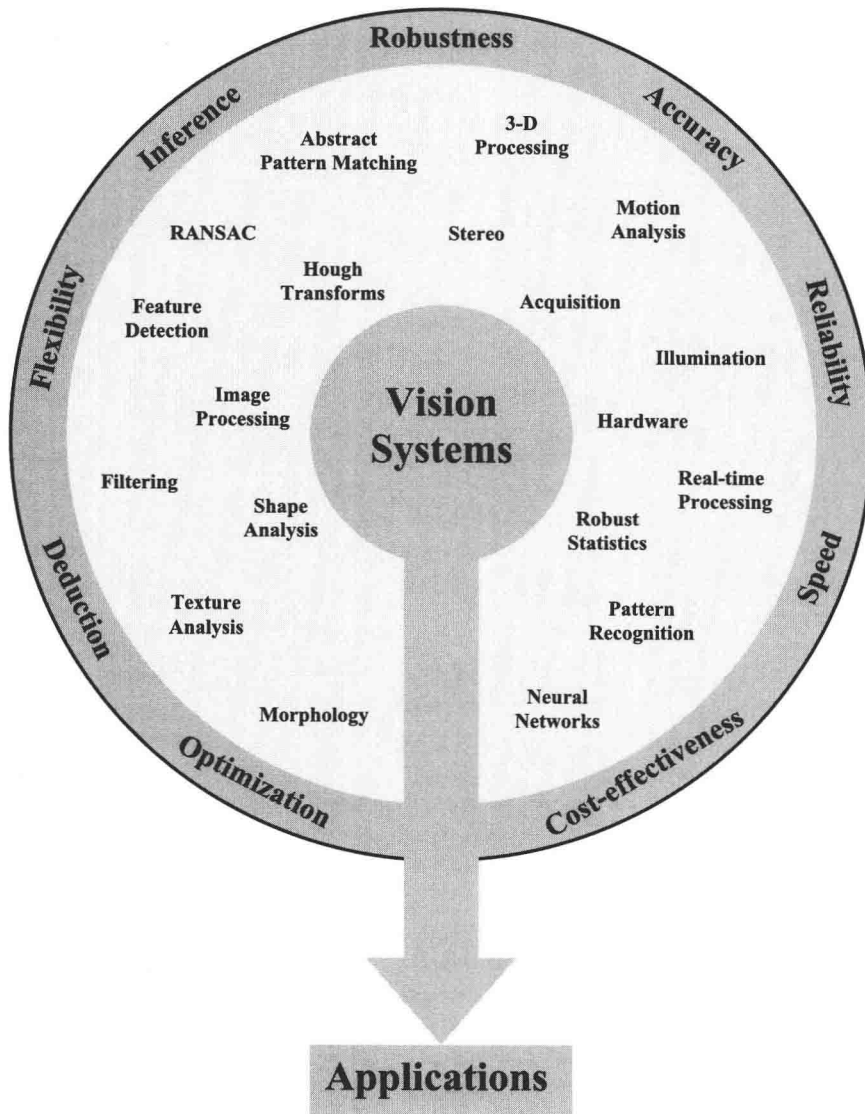
*To my late father, Arthur Granville Davies, who passed on to me
his appreciation of the beauties of mathematics and science.*

To my wife, Joan, for love, patience, support, and inspiration.

To my children, Elizabeth, Sarah, and Marion, the music in my life.

*To my grandson, Jasper, for reminding me of the carefree
joys of youth.*

Influences Impinging upon Integrated Vision System Design



Foreword

Although computer vision is such a relatively young field of study, it has matured immensely over the last 25 years or so—from well-constrained, targeted applications to systems that learn automatically from examples.

Such progress over these 25 years has been spurred not least by mind-boggling advances in vision and computational hardware, making possible simple tasks that could take minutes on small images, now integrated as part of real-time systems that do far more in a fraction of a second on much larger images in a video stream.

This all means that the focus of research has been in a perpetual state of change, marked by near-exponential advances and achievements, and witnessed by the quality, and often quantity, of outstanding contributions to the field published in key conferences and journals such as ICCV and PAMI. These advances are most clearly reflected by the growing importance of the application areas in which the novel and real-time developments in computer vision have been applied to or developed for. Twenty-five years ago, industrial quality inspection and simple military applications ruled the waves, but the emphasis has since spread its wings, some slowly and some like wildfire, to many more areas, for example, from medical imaging and analysis to surveillance and, inevitably, complex military and space applications.

So how does Roy's book reflect this shift? Naturally, there are many fundamental techniques that remain the same, and this book is a wonderful treasure chest of tools that provides the fundamentals for any researcher and teacher. More modern and state-of-the-art methodologies are also covered in the book, most of them pertinent to the topical application areas currently driving not only the research agenda, but also the market forces. In short, the book is a direct reflection of the progress and key methodologies developed in computer vision over the last 25 years and more.

Indeed, while the third edition of this book was already an excellent, successful, and internationally popular work, this fourth edition is greatly enhanced and updated. All its chapters have been substantially revised and brought up to date by the inclusion of many new references covering advances in the subject made even in the past year. There are now also two entirely new chapters (to reflect the great strides that have been made in the area of video analytics) on surveillance and in-vehicle vision systems. The latter is highly relevant to the coming era of advanced driver assistance systems, and the former's importance and role requires no emphasis in this day and age where so many resources are dedicated to criminal and terrorist activity monitoring and prevention.

The material in the book is written in a way that is both approachable and didactic. It is littered with examples and algorithms. I am sure that this volume will be welcomed by a great many students and workers in computer and machine vision, including practitioners in academia and industry—from beginners who are

starting out in the subject to advanced researchers and workers who need to gain insight into video analytics. I will also welcome it personally, for use by my own undergraduate and postgraduate students, and will value its presence on my bookshelf as an up-to-date reference on this important subject.

Finally, I am very happy to go on record as saying that Roy is the right person to have produced this substantial work. His long experience in the field of computer and machine vision surpasses even the “big bang” in computer vision around 25 years ago in the mid-80s when the Alvey Vision Conference (UK) and CVPR (USA) were only inchoates of what they have become today and reaches back to when ICPR and IAPR began to be dominated by image processing in the late 70s.

Majid Mirmehdi
University of Bristol, UK
September 2011

Preface

PREFACE TO THE FOURTH EDITION

The first edition came out in 1990, and was welcomed by many researchers and practitioners. However, in the subsequent two decades, the subject moved on at a rapidly accelerating rate, and many topics that hardly deserved a mention in the first edition had to be solidly incorporated in subsequent editions. It seemed particularly important to bring in significant amounts of new material on mathematical morphology, 3-D vision, invariance, motion analysis, object tracking, artificial neural networks, texture analysis, X-ray inspection, foreign object detection, and robust statistics. There are thus new chapters or appendices on these topics, and they have been carefully integrated with the existing material. The greater proportion of the new material has been included in Parts 3 and 4. So great has been the growth in work on 3-D vision and its applications that the original single chapter on 3-D vision had to be expanded into the set of *five* chapters on 3-D vision and motion forming Part 3, together with a further *two* chapters on surveillance and in-vehicle vision systems in Part 4. Indeed, these changes have been so radical that the title of the book has had to be modified to reflect them. At this stage, Part 4 encompasses such a range of chapters—covering applications and the components needed for constructing real-time visual pattern recognition systems—that it is difficult to produce a logical ordering for them: notably, the topics interact with each other at a variety of different levels—theory, algorithms, methodologies, practicalities, design constraints, and so on. However, this should not matter in practice, as the reader will be exposed to the essential richness of the subject, and his/her studies should be amply rewarded by increased understanding and capability.

It is worth remarking that, at this point in time, computer vision has attained a level of maturity that has made it substantially more rigorous, reliable, generic, and—in the light of the improved hardware facilities now available for its implementation (not least, FPGA and GPU types of solution)—capable of real-time performance. This means that workers are more than ever before using it in serious applications, and with fewer practical difficulties. It is intended that this edition of the book will reflect this radically new and exciting state of affairs at a fundamental level.

A typical final-year undergraduate course on vision for electronic engineering or computer science students might include much of the work of Chapters 1–10 and 14, 15, plus a selection of sections from other chapters, according to requirements. For MSc or PhD research students, a suitable lecture course might go on

to cover Part 3 in depth, including several of the chapters in Part 4,¹ with many practical exercises being undertaken on an image analysis system. Here, much will depend on the research program being undertaken by each individual student. At this stage, the text will have to be used more as a handbook for research, and indeed, one of the prime aims of the volume is to act as a handbook for the researcher and practitioner in this important area.

As mentioned in the original Preface, this book leans heavily on experience I have gained from working with postgraduate students: in particular, I would like to express my gratitude to Mark Edmonds, Simon Barker, Daniel Celano, Darrel Greenhill, Derek Charles, Mark Sugrue, and Georgios Mastorakis, all of whom have in their own ways helped to shape my view of the subject. In addition, it is a special pleasure to recall very many rewarding discussions with my colleagues Barry Cook, Zahid Hussain, Ian Hannah, Dev Patel, David Mason, Mark Bateman, Tieying Lu, Adrian Johnstone, and Piers Plummer, the last two named having been particularly prolific in generating hardware systems for implementing my research group's vision algorithms. Next, I am immensely grateful to Majid Mirmehdi for reading much of the manuscript and making insightful comments and valuable suggestions. Finally, I am indebted to Tim Pitts of Elsevier Science for his help and encouragement, without which this fourth edition might never have been completed.

SUPPORTING MATERIALS

Elsevier's website for the book contains resources to help students and other readers using this text. For further information, go to the publisher's website:

[http://www.elsevierdirect.com/companion.jsp?ISBN = 9780123869081](http://www.elsevierdirect.com/companion.jsp?ISBN=9780123869081)

E. R. DAVIES
*Royal Holloway,
University of London*

PREFACE TO THE FIRST EDITION (1990)

Over the past 30 years or so, machine vision has evolved into a mature subject embracing many topics and applications: these range from automatic (robot) assembly to automatic vehicle guidance, from automatic interpretation of documents to verification of signatures, and from analysis of remotely sensed images to checking of fingerprints and human blood cells; currently, automated visual inspection is undergoing very substantial growth, necessary improvements in

¹The importance of the appendix on robust statistics should not be underestimated once one gets onto serious work, although this will probably be outside the restrictive environment of an undergraduate syllabus.

quality, safety and cost-effectiveness being the stimulating factors. With so much ongoing activity, it has become a difficult business for the professional to keep up with the subject and with relevant methodologies: in particular, it is difficult to distinguish accidental developments from genuine advances. It is the purpose of this book to provide background in this area.

The book was shaped over a period of 10–12 years, through material I have given on undergraduate and postgraduate courses at London University, and contributions to various industrial courses and seminars. At the same time, my own investigations coupled with experience gained while supervising PhD and post-doctoral researchers helped to form the state of mind and knowledge that is now set out here. Certainly it is true to say that if I had had this book 8, 6, 4, or even 2 years ago, it would have been of inestimable value to myself for solving practical problems in machine vision. It is therefore my hope that it will now be of use to others in the same way. Of course, it has tended to follow an emphasis that is my own—and in particular one view of one path toward solving automated visual inspection and other problems associated with the application of vision in industry. At the same time, although there is a specialism here, great care has been taken to bring out general principles—including many applying throughout the field of image analysis. The reader will note the universality of topics such as noise suppression, edge detection, principles of illumination, feature recognition, Bayes' theory, and (nowadays) Hough transforms. However, the generalities lie deeper than this. The book has aimed to make some general observations and messages about the limitations, constraints, and tradeoffs to which vision algorithms are subject. Thus, there are themes about the effects of noise, occlusion, distortion and the need for built-in forms of robustness (as distinct from less successful *ad hoc* varieties and those added on as an afterthought); there are also themes about accuracy, systematic design, and the matching of algorithms and architectures. Finally, there are the problems of setting up lighting schemes which must be addressed in complete systems, yet which receive scant attention in most books on image processing and analysis. These remarks will indicate that the text is intended to be read at various levels—a factor that should make it of more lasting value than might initially be supposed from a quick perusal of the Contents.

Of course, writing a text such as this presents a great difficulty in that it is necessary to be highly selective: space simply does not allow everything in a subject of this nature and maturity to be dealt with adequately between two covers. One solution might be to dash rapidly through the whole area mentioning everything that comes to mind, but leaving the reader unable to understand anything in detail or to *achieve* anything having read the book. However, in a practical subject of this nature, this seemed to me a rather worthless extreme. It is just possible that the emphasis has now veered too much in the opposite direction, by coming down to practicalities (detailed algorithms, details of lighting schemes, and so on): individual readers will have to judge this for themselves. On the other hand, an author has to be true to himself and my view is that it is better for a reader or

student to have mastered a coherent series of topics than to have a mish-mash of information that he is later unable to recall with any accuracy. This, then, is my justification for presenting this particular material in this particular way and for reluctantly omitting from detailed discussion such important topics as texture analysis, relaxation methods, motion, and optical flow.

As for the organization of the material, I have tried to make the early part of the book lead into the subject gently, giving enough detailed algorithms (especially in Chapters 2 and 6) to provide a sound feel for the subject—including especially vital, and in their own way quite intricate, topics such as connectedness in binary images. Hence, Part 1 provides the lead-in, although it is not always trivial material and indeed some of the latest research ideas have been brought in (e.g., on thresholding techniques and edge detection). Part 2 gives much of the meat of the book. Indeed, the (book) literature of the subject currently has a significant gap in the area of intermediate-level vision; while high-level vision (AI) topics have long caught the researcher's imagination, intermediate-level vision has its own difficulties which are currently being solved with great success (note that the Hough transform, originally developed in 1962, and by many thought to be a very specialist topic of rather esoteric interest, is arguably only now coming into its own). Part 2 and the early chapters of Part 3 aim to make this clear, while Part 4 gives reasons why this particular transform has become so useful. As a whole, Part 3 aims to demonstrate some of the practical applications of the basic work covered earlier in the book, and to discuss some of the principles underlying implementation: it is here that chapters on lighting and hardware systems will be found. As there is a limit to what can be covered in the space available, there is a corresponding emphasis on the theory underpinning practicalities. Probably, this is a vital feature, since there are many applications of vision both in industry and elsewhere, yet listing them and their intricacies risks dwelling on interminable detail, which some might find insipid; furthermore, detail has a tendency to date rather rapidly. Although the book could not cover 3-D vision in full (this topic would easily consume a whole volume in its own right), a careful overview of this complex mathematical and highly important subject seemed vital. It is therefore no accident that Chapter 16 is the longest in the book. Finally, Part 4 asks questions about the limitations and constraints of vision algorithms and answers them by drawing on information and experience from earlier chapters. It is tempting to call the last chapter the Conclusion. However, in such a dynamic subject area, any such temptation has to be resisted, although it has still been possible to draw a good number of lessons on the nature and current state of the subject. Clearly, this chapter presents a personal view but I hope it is one that readers will find interesting and useful.

About the Author



Roy Davies is Emeritus Professor of Machine Vision at Royal Holloway, University of London. He has worked on many aspects of vision, from feature detection and noise suppression to robust pattern matching and real-time implementations of practical vision tasks. His interests include automated visual inspection, surveillance, vehicle guidance, and crime detection. He has published more than 200 papers and three books—*Machine Vision: Theory, Algorithms, Practicalities* (1990), *Electronics, Noise and Signal Recovery* (1993), and *Image Processing for the Food Industry* (2000); the first of these has been widely used internationally for more than 20 years, and is now out in this much enhanced fourth edition. Roy is a Fellow of the IoP and the IET, and a Senior Member of the IEEE. He is on the Editorial Boards of *Real-Time Image Processing*, *Pattern Recognition Letters*, *Imaging Science*, and *IET Image Processing*. He holds a DSc at the University of London: he was awarded BMVA Distinguished Fellow in 2005 and Fellow of the International Association of Pattern Recognition in 2008.

Acknowledgements

The author would like to credit the following sources for permission to reproduce tables, figures and extracts of text from earlier publications:

Elsevier

For permission to reprint portions of the following papers from *Image and Vision Computing* as text in Chapters 5 and 14; as Tables 5.1–5.5; and as Figures 3.29, 5.2, 14.1, 14.2, 14.6:

Davies (1984b, 1987c)

Davies, E.R. (1991). *Image and Vision Computing* **9**, 252–261

For permission to reprint portions of the following paper from *Pattern Recognition* as text in Chapter 9; and as Figure 9.11:

Davies and Plummer (1981)

For permission to reprint portions of the following papers from *Pattern Recognition Letters* as text in Chapters 3, 5, 11–14, 21, 24; as Tables 3.2; 12.3; 13.1; and as Figures 3.6, 3.8, 3.10; 5.1, 5.3; 11.1, 11.2a, 11.3b; 12.4, 12.5, 12.6, 12.7–12.10; 13.1, 13.3–13.11; 21.3, 21.6:

Davies (1986a,b; 1987a,e,f; 1988b,c,e,f; 1989a)

Davies et al. (2003a)

For permission to reprint portions of the following paper from *Signal Processing* as text in Chapter 3; and as Figures 3.15–3.20:

Davies (1989b)

For permission to reprint portions of the following paper from *Advances in Imaging and Electron Physics* as text in Chapter 3:

Davies (2003c)

For permission to reprint portions of the following article from *Encyclopedia of Physical Science and Technology* as Figures 9.9, 9.12, 10.1, 10.4:

Davies, E.R. (1987). Visual inspection, automatic (robotics). In: Meyers, R.A. (ed.) *Encyclopedia of Physical Science and Technology, Vol. 14*. Academic Press, San Diego, pp. 360–377

The Committee of the Alvey Vision Club

For permission to reprint portions of the following paper as text in Chapter 14; and as Figures 14.1, 14.2, 14.6:

Davies, E.R. (1988). An alternative to graph matching for locating objects from their salient features. *Proc. 4th Alvey Vision Conf., Manchester (31 August–2 September)*, pp. 281–286

CEP Consultants Ltd (Edinburgh)

For permission to reprint portions of the following paper as text in Chapter 20:

Davies, E.R. (1987). Methods for the rapid inspection of food products and small parts. In: McGeough, J.A. (ed.) *Proc. 2nd Int. Conf. on Computer-Aided Production Engineering, Edinburgh (13–15 April)*, pp. 105–110

EURASIP

For permission to reprint portions of the following papers as text in Chapter 21; and as Figures 21.5, 21.7–21.11:

Davies (1998)
Davies et al. (1998b)

These papers were first published in the *Proceedings of the 9th European Signal Processing Conference (EUSIPCO-1998)* in 1998, published by EURASIP.

IEEE

For permission to reprint portions of the following paper as text in Chapter 3; and as Figures 3.4, 3.5, 3.7, 3.11:

Davies (1984a)

IET

For permission to reprint portions of the following papers from the *IET Proceedings and Colloquium Digests* as text in Chapters 3, 4, 6, 7, 13, 20–23, 25, 26; as Tables 3.3; 26.4, 26.5; and as Figures 3.21, 3.22, 3.24; 4.8–4.12; 6.5–6.9, 6.12; 7.6–7.8; 13.12; 20.2–20.4; 21.1, 21.2, 21.12; 22.16–22.18; 23.1, 23.3, 23.4; 25.4–25.8:

Davies (1985; 1988a; 1997b; 1999f; 2000b,c; 2005; 2008b)
Davies, E.R. (1997). Algorithms for inspection: constraints, tradeoffs and the design process. IEE Digest no. 1997/041, Colloquium on *Industrial Inspection*, IEE (10 Feb.), pp. 6/1–5
Sugrue and Davies (2007)
Mastorakis and Davies (2011)
Davies et al. (1998a)
Davies and Johnstone (1989)

IFS Publications Ltd

For permission to reprint portions of the following paper as text in Chapters 12, 20; and as Figures 12.1, 12.2, 20.5:

Davies (1984c)

The Council of the Institution of Mechanical Engineers

For permission to reprint portions of the following paper as text in Chapter 26; and as Tables 26.1, 26.2:

Davies and Johnstone (1986)

MCB University Press (Emerald Group)

For permission to reprint portions of the following paper as Figure 20.6:

Patel et al. (1995)

The Royal Photographic Society

For permission to reprint portions of the following papers¹ as text in Chapter 3; as Table 3.4; and as Figures 3.12, 3.13, 3.25–3.28:

Davies (2000f)

Charles and Davies (2004)

Springer-Verlag

For permission to reprint portions of the following papers as text in Chapters 6, 21; and as Figures 6.2, 6.4:

Davies (1988d), Figs. 1–3

Davies, E.R. (2003). Design of object location algorithms and their use for food and cereals inspection. Chapter 15 in Graves, M. and Batchelor, B.G. (eds.). *Machine Vision Techniques for Inspecting Natural Products*. Springer-Verlag, pp. 393–420

Peter Stevens Photography

For permission to reprint a photograph as Figure 3.12(a).

F.H. Sumner

For permission to reprint portions of the following article from *State of the Art Report: Supercomputer Systems Technology* as text in Chapter 9; and as Figure 9.4:

Davies, E.R. (1982). Image processing. In: Sumner, F.H. (ed.) *State of the Art Report: Supercomputer Systems Technology*. Pergamon Infotech, Maidenhead, pp. 223–244

¹See also the Maney website: www.maney.co.uk/journals/ims.

Unicom Seminars Ltd

For permission to reprint portions of the following paper as text in Chapter 20; and as Figures 11.2b, 11.3a, 12.3, 12.13, 13.12c, 14.8:

Davies, E.R. (1988). Efficient image analysis techniques for automated visual inspection. *Proc. Unicom Seminar on Computer Vision and Image Processing, London (29 November–1 December)*, pp. 1–19

World Scientific

For permission to reprint portions of the following book as text in Chapters 7, 21, 22, 23, 26; and as Figures 7.1–7.4, 21.4, 22.20, 23.15, 23.16, 26.3:

Davies (2000a)

Royal Holloway, University of London

For permission to reprint extracts from the following examination questions, originally written by E.R. Davies:

EL385/97/2; EL333/98/2; EL333/99/2, 3, 5, 6; EL333/01/2, 4–6;
PH5330/98/3, 5; PH5330/03/1–5; PH4760/04/1–5

University of London

For permission to reprint extracts from the following examination questions, originally written by E.R. Davies:

PH385/92/2, 3; PH385/93/1–3; PH385/94/1–4; PH385/95/4; PH385/96/3, 6;
PH433/94/3, 5; PH433/96/2, 5

Glossary of Acronyms and Abbreviations

1-D	one dimension/one-dimensional
2-D	two dimensions/two-dimensional
3-D	three dimensions/three-dimensional
3DPO	3-D part orientation system
ACM	Association for Computing Machinery (USA)
ADAS	advanced driver assistance system
ADC	analog to digital converter
AI	artificial intelligence
ANN	artificial neural network
APF	auxiliary particle filter
ASCII	American Standard Code for Information Interchange
ASIC	application-specific integrated circuit
ATM	automated teller machine
AUC	area under curve
AVI	audio video interleave
BCVM	between-class variance method
BetaSAC	beta [distribution] sampling consensus
BMVA	British Machine Vision Association
BRAM	block of RAM
BRDF	bidirectional reflectance distribution function
CAD	computed-aided design
CAM	computer-aided manufacture
CCD	charge-coupled device
CCTV	closed-circuit television
CDF	cumulative distribution function
CIM	computer integrated manufacture
CLIP	cellular logic image processor
CPU	central processing unit
DCSM	distinct class based splitting measure
DET	Beaudet determinant operator
DEXA	dual-emission X-ray absorptiometry
DG	differential gradient
DN	Dreschler–Nagel corner detector
DoF	degree of freedom
DoG	difference of Gaussians
DSP	digital signal processor
EM	expectation maximization
EURASIP	European Association for Signal Processing
FAST	features from accelerated segment test
FFT	fast Fourier transform