

BMVC91

BMVC91

Proceedings of the British Machine Vision
Conference, organised for the British
Machine Vision Association by the
Turing Institute

24-26 September 1991
University of Glasgow

Edited by
Peter Mowforth



Springer-Verlag
London Berlin Heidelberg New York
Paris Tokyo Hong Kong
Barcelona Budapest

Peter Mowforth, BSc (Hons), PhD
Director, The Turing Institute
George House
36 North Hanover Street
Glasgow G1 2AD, UK

ISBN 3-540-19715-X Springer-Verlag Berlin Heidelberg New York
ISBN 0-387-19715-X Springer-Verlag New York Berlin Heidelberg

British Library Cataloguing in Publication Data
British Machine Vision Conference (1991: Glasgow)
British Machine Vision Conference: proceedings of the annual conference
of the British Machine Vision Association, 24-26 September 1991, Glasgow.
I. Title II. Mowforth, P.H.
006.3
ISBN 3-540-19715-X

Library of Congress Cataloging-in-Publication Data
British Machine Vision Conference (1991: Glasgow, Scotland)
British Machine Vision Conference: proceedings of the annual conference
of the British Machine Vision Association, 24-26 September 1991,
Glasgow / edited by Peter H. Mowforth
p. cm.
ISBN 0-387-19715-X. – ISBN 3-540-19715-X
1. Computer vision-Congresses. I. Mowforth, P. (Peter) 1953-
II. British Machine Vision Association. III. Title
TA1632.B75 1991
006.3'7-dc20
91-26755
CIP

Apart from any fair dealing for the purposes of research or private study,
or criticism or review, as permitted under the Copyright, Design and
Patents Act 1988, this publication may only be reproduced, stored or
transmitted, in any form or by any means, with the prior permission in
writing of the publishers, or in the case of reprographic reproduction in
accordance with the terms of licences issued by the Copyright Licensing
Agency. Enquiries concerning reproduction outside those terms should be
sent to the publishers.

© Springer-Verlag London Limited 1991
Printed in Great Britain

The use of registered names, trademarks etc. in this publication does not
imply, even in the absence of a specific statement, that such names are
exempt from the relevant laws and regulations and therefore free for
general use.

The publisher makes no representation, express or implied, with regard to
the accuracy of the information contained in this book and cannot accept
any legal responsibility or liability for any errors or omissions that may be
made.

Printed and bound by The Alden Press, Osney Mead, Oxford
34/3830-543210 Printed on acid-free paper

Foreword

Lewis Carroll once wrote a story about a king who wanted a very accurate map of his kingdom. The king had a pathologically fastidious eye for detail and consequently decided that the map was to be produced at a scale of 1:1. The scribes dutifully set to and, in time, the map was made. The map carried details of every tree, every rock and every blade of grass throughout the entire land. The problem occurred when they tried to use it. First of all, the map was extraordinarily difficult to open out and line up with the countryside. Its sheer bulk meant that it took whole armies to carry it and a great host of bureaucrats and technicians to maintain the information. Such was the detail of the map that as soon as the wind blew strongly, whole sections needed to be redrawn. What was worse was that all the farmers protested because the map completely cut out the light from the sun and all the crops died. Eventually the howls of protest became so strong that the king was forced to take action. He did away with the old paper copy and decided to use the kingdom itself as the map. All lived happily ever after.

There are, at least, two morals to this tale. First, you are almost certainly doomed to failure if you do not get the representation of the problem right. Second, it is important that the representation uses a level of precision appropriate for the tasks required of the system – and no more. Nowhere are these requirements more important to get right than in computer vision. The fact that sophisticated, human-style, machine vision seems to be both highly complex and appears to have a seemingly boundless appetite for computing resources means that we must pay great attention to finding good (i.e. simple) representations in our machines. We must also think hard about the tasks that we want vision to solve and so be guided into selecting appropriate data-types and levels of precision that meet the system requirements. We should also try to avoid any unnecessary duplication of information or the computation of symbols that we later discover to be superfluous to our needs. Were we to follow these guidelines when building sophisticated vision systems we might be led to the following conclusions:

1. Do not perform edge detection on the input images. Signal matching problems such as stereo and motion correspondence can be solved without the use of edges and, in any case, you need to keep the gray-level information if you intend to compute shading or colour information. If you only keep the edges you are guilty of throwing away the baby with the bath water whilst if you try to keep the original image and its edge map(s) you are guilty of over representation. When we appreciate that there is no substantive evidence that biological systems perform edge detection, the justification for any such technique looks a bit thin.
2. Foveate your representations. When seeking the holy grail of the minimal representation, there is much we can learn from biological systems. Visual data seems to be log-scaled on almost every dimension. Both intensity and eccentricity are thus scaled so producing the representational appearance of being at the centre of a foveal magnification bubble in which most of our sensitivities and resolutions decrease away from the point of fixation. Many researchers produce high resolution results over the entire field of view of the camera and when asked about foveal/pyramidal representations simply say that this could be "bolted on" afterwards if you wanted the process to go faster. However, such an approach may well miss the point. Those researchers who have explored the use of such data structures have realised that many problems that might originally have needed to be algorithmically coded, now simply drop out of the system as a natural property of the representation. Polar tessellations produce rotation and size invariance as well as a "time-to-impact" computation which is a linear pixel count. Further, resolution pyramids naturally facilitate all the "octave based" algorithms as well as introducing a massive saving in computational effort.
3. Be qualitative. If you can achieve your task using qualitative representations, do not go through all the added difficulty of scaling the data to make it quantitative. For example, if relative gray-values are output from a real-time stereo system then you do not need to scale them to perform hand-eye control, vergence, motion segmentation or tracking providing the stereo system provides signals within the control loop of the system. Such an observation may suggest that whilst a few trinocular systems quantitatively lumber around some research laboratories, nature seems to suggest that two eyes are quite sufficient.

Returning to Lewis Carroll the king's final solution was to do away with the model altogether. Such Brooksonian ideas might go down well in parts of MIT but as soon as any sophisticated problems need to be tackled, the need for models becomes clear. Whilst the

pioneering ideas of Kenneth Craik¹ are as true today as they were when first conceived, today's pressing problem is now to discover the minimal representation to facilitate the process.

Having been granted the opportunity to offload some accumulated prejudices, some things need to be said about BMVC'91. The familiar joke about stage five of a successful project (i.e. "praise and honours for the non-participants"), hangs guiltily over my head. The real work behind organising the conference and collaboration with Springer-Verlag was carried out by Jon Ritchie, Tanya Oliver, Irene Brebner and David Wilson. As a somewhat bemused bystander to the event I appreciate that they, at least, have done much to make this a useful, representative and well organised summary of British machine vision research for 1991.

Peter Mowforth
Director, The Turing Institute
29th July 1991

¹Craik, the Cambridge psychologist, was the first to argue strongly that perception is a model in the brain.

BMVC91

Programme Committee

Mr W.G.L. Adaway
Computer Recognition Systems Ltd
Unit 3, Fishponds Close
Wokingham, Berkshire

Mr M.B. Brown
British Aerospace PLC
Sowerby Research Centre (FPC:267)
P.O. Box 5, Filton
Bristol BS12 7QW

Dr B.F. Buxton
GEC Hirst Research Centre
East Lane
Wembley, Middlesex HA9 7PP

Dr J. Illingworth
Department of Electronic and
Electrical Engineering
Vision, Speech and Signal Processing
Group
University of Surrey
Guildford GU2 5XH

Dr J. Kittler
Department of Electronic and
Electrical Engineering
University of Surrey
Guildford GU2 5XH

Dr J. Knapman
Lead Image Researcher, (MP.137)
IBM UK Ltd., UK Scientific Centre
Athelstan House
St Clement Street
Winchester, Hampshire SO23 9DR

Dr Peter Mowforth
The Turing Institute
George House
36 North Hanover Street
Glasgow G1 2AD

Professor Peter Saraga
Philips Research Laboratories
Cross Oak Lane
Redhill, Surrey RH1 5HA

Mr A.C. Sleigh
Director of Science CIS
Room 1258
MOD Main Building
Whitehall, London SW1A 2HB

Mr G.D. Sullivan
Department of Computer Science
School of Engineering and
Information Sciences
University of Reading, P.O. Box 220
Whiteknights
Reading RG6 2AY

Professor C.J. Taylor
Department of Medical Biophysics
University of Manchester
Stopford Building
Oxford Road
Manchester M13 9PT

Local Organisers

Peter Mowforth (*Conference Chair*)

Tanya Oliver

Irene Brebner

Jon Ritchie

Jim Richmond

David Wilson

The Turing Institute

George House

36 North Hanover Street

Glasgow G1 2AD

Contents

Image Motion Analysis Made Simple and Fast, One Component at a Time <i>Peter J. Burt</i>	1
Visual Modelling <i>Demetri Terzopoulos</i>	9
Distributed Dynamic Processing for Edge Detection <i>H. Tunley</i>	13
Boundary Detection Using Bayesian Nets <i>N. Bryson and C.J. Taylor</i>	22
Parallel Implementation of Lagrangian Dynamics for Real-time Snakes <i>R.M. Curwen, A. Blake and R. Cipolla</i>	29
Supervised Segmentation Using a Multi-resolution Data Representation <i>Isaac Ng, J. Kittler and J. Illingworth</i>	36
3D Grouping by Viewpoint Consistency Ascent <i>L. Du, G.D. Sullivan and K.D. Baker</i>	45
A Trainable Method of Parametric Shape Description <i>T.F. Cootes, D.H. Cooper, C.J. Taylor and J. Graham</i>	54
Using Projective Invariants for Constant Time Library Indexing in Model Based Vision <i>C.A. Rothwell, A. Zisserman, D.A. Forsyth and J.L. Mundy</i>	62
Invariants of a Pair of Conics Revisited <i>Long Quan, Patrick Gros and Roger Mohr</i>	71
A Modal Approach to Feature-based Correspondence <i>Larry S. Shapiro and J. Michael Brady</i>	78

A Method of Obtaining the Relative Positions of 4 Points from 3 Perspective Projections <i>H. Christopher Longuet-Higgins</i>	86
Properties of Local Geometric Constraints <i>Alistair J. Bray</i>	95
Texture Boundary Detection – A Structural Approach <i>Wen Wen and Richard J. Fryer</i>	104
The Inference of Structure in Images Using Multi-local Quadrature Filters <i>K. Langley and T.J. Atherton</i>	111
Low-level Grouping of Straight Line Segments <i>A. Etemadi, J-P. Schmidt, G. Matas, J. Illingworth and J. Kittler</i>	118
Connective Hough Transform <i>Shiu Yin K. Yuen</i>	127
Ellipse Detection and Matching with Uncertainty <i>Tim Ellis, Ahmed Abbood and Beatrice Brillault</i>	136
Cooperating Motion Processes <i>P.H.S. Torr, T. Wong, D.W. Murray and A. Zisserman</i>	145
Tracking Curved Objects by Perspective Inversion <i>Alistair J. Bray</i>	151
Optimal Surface Fusion <i>Peter R.J. North</i>	160
Recursive Updating of Planar Motion <i>D.W. Murray and D.M. Pickup</i>	169
A Fractal Shape Signature <i>R.I. Taylor and P.H. Lewis</i>	178
Locating Overlapping Flexible Shapes Using Geometrical Constraints <i>David H. Cooper, Christopher J. Taylor, Jim Graham and Tim F. Cootes</i>	185
Gaze Control for a Two-Eyed Robot Head <i>Fenglei Du, Michael Brady and David Murray</i>	193
Visual Evidence Accumulation in Radiograph Inspection <i>Horace H.S. Ip</i>	202

A New Approach to Active Illumination <i>J.P. McDonald, J.P. Siebert and R.J. Fryer</i>	210
A Comparative Analysis of Algorithms for Determining the Peak Position of a Stripe to Sub-pixel Accuracy <i>D.K. Naidu and R.B. Fisher</i>	217
Synthetic Images of Faces – An Approach to Model-Based Face Recognition <i>Andrew C. Aitchison and Ian Craw</i>	226
Finding Image Features Using Deformable Templates and Detailed Prior Statistical Knowledge <i>Alan Bennett and Ian Craw</i>	233
Relational Model Construction and 3D Object Recognition from Single 2D Monochromatic Image <i>S. Zhang, G.D. Sullivan and K.D. Baker</i>	240
Recognising Cortical Sulci and Gyri in MR Images <i>C.I. Attwood, G.D. Sullivan and K.D. Baker</i>	249
Classification of Breast Tissue by Texture Analysis <i>Peter Miller and Sue Astley</i>	258
Model-Based Image Interpretation Using Genetic Algorithms <i>A. Hill and C.J. Taylor</i>	266
Automated Analysis of Retinal Images <i>Piotr Jasiobedzki and Chris J. Taylor</i>	275
Segmentation of MR Images Using Neural Nets <i>D.L. Toulson and J.F. Boyce</i>	284
Detecting and Classifying Intruders in Image Sequences <i>Paul L. Rosin</i>	293
Structure from Constrained Motion Using Point Correspondences <i>T.N. Tan, G.D. Sullivan and K.D. Baker</i>	301
Model-Based Tracking <i>A.D. Worrall, R.F. Marlin, G.D. Sullivan and K.D. Baker</i>	310
Local Method for Curved Edges and Corners <i>R.T. Shann and J.P. Oakley</i>	319
The Kinematics and Eye Movements for a Two-Eyed Robot Head <i>Fenglei Du and Michael Brady</i>	323

Colour and Texture Analysis for Automated Sorting of Eviscera <i>Richard J. Gibbons and David J. Williams</i>	327
Image Coding Based on Contour Models <i>S. Marshall</i>	331
An Automated Approach to Stereo Matching Seasat Imagery <i>Mia Denos</i>	335
Data Fusion Using an MLP <i>D.M. Booth, N.A. Thacker, J.E.W. Mayhew and M.K. Pidcock</i>	339
Passive Estimation of Range to Objects from Image Sequences <i>T.J. Atherton, D.J. Kerbyson and G.R. Nudd</i>	343
Imaging Polarimetry for Industrial Inspection <i>R.J. Fryer and J. Miller</i>	347
Computing with Uncertainty: Intervals versus Probabilities <i>Mark J.L. Orr, Robert B. Fisher and John Hallam</i>	351
Recognizing Parameterized Objects Using 3D Edges <i>Ian Reid and J. Michael Brady</i>	355
Optic Disk Boundary Detection <i>Simon Lee and J. Michael Brady</i>	359
Computation of Smoothed Local Symmetries on a MIMD Architecture <i>M.W. Wright</i>	363
Parameterising Images for Recognition and Reconstruction <i>Ian Craw and Peter Cameron</i>	367
Kalman Filters in Constrained Model Based Tracking <i>R.F. Marslin, G.D. Sullivan and K.D. Baker</i>	371
A Novel Approach to Motion Segmentation <i>D. Sinclair, Andrew Blake, Paul Beardsley and David Murray</i>	375
An Efficient and Robust Local Boundary Operator <i>Mark J. Robinson, John P. Oakley and Richard T. Shann</i>	379
The Active Stereo Probe: Dynamic Video Feedback <i>J.P. Siebert, C.W. Urquhart, D.F. Wilson, J.P. McDonald, P.H. Mowforth and R.J. Fryer</i>	383
Design of an Anthropomorphic Robot Head <i>K. Undbekken, D.F. Wilson, P.H. Mowforth, S. Solbakken and B. Eiklid</i>	387

A Monocular Ground Plane Estimation System <i>D.A. Castelov and A.J. Rérolle</i>	392
Recognition with Second-Order Topographic Surface Features <i>Robert B. Fisher</i>	396
Heuristically Guided Polygon Finding <i>K.C. Wong, J. Kittler and J. Illingworth</i>	400
The Amplification of Textural Differences <i>D.M. Booth, J.E.W. Mayhew and M.K. Pidcock</i>	408
Edge Labelling by Fusion of Intensity and Range Data <i>G. Zhang and A.M. Wallace</i>	412
Author Index	417

