

BMVA News

The Newsletter of the British Machine Vision Association and
Society for Pattern Recognition

Volume 6 Number 1
July 1995

Editor: Dr Philip McLauchlan
Department of Engineering Science
University of Oxford
Oxford OX1 3PJ
Tel: (01865) 273127
Fax: (01865) 273908
Email: pm@robots.oxford.ac.uk

BMVA News¹ is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the editor by 31st August 1995.

Contents

Editorial.....	1
Biological Vision.....	1
BMVC'95 Update	3
Representation Meeting.....	3
Teaching Machine Vision.....	4
Summer School.....	5
CVRMed95.....	6
Technology Transfer Initiative.....	9

¹The British Machine Vision Association and Society for Pattern Recognition is a Company limited by guarantee, No. 2543446, registered in England and Wales. Registered Office: Granta Lodge, 71 Graham Road, Malvern, WR14 2JS. The Association is a non-profit-making body and is registered as charity No. 1002307.

ICCV '95.....	9
---------------	---

Editorial

Dear pixel-heads, as I write to you I have just completed a whirlwind tour of the States and Canada, taking in Boston (for ICCV, frisbee-buying and ice cream), Pittsburgh, Albany, LA, San Francisco and Vancouver. Apart from LA and Pittsburgh these places were all so nice that I wonder why anyone would want to live in Britain. They have chairs in bookshops, where you can sit and read. You get free refills of coffee! You can have breakfast in a “Denny’s” diner on the east coast, get on a plane, and next morning have *exactly* the same breakfast on the west coast. Ordering a “large” ice cream/pizza is a mistake because you get impractically enormous amounts of food. What an amazing place. But don’t do what our own Ralph Highnam did and wander drunkenly into Bed-Sty (in New York) in the early hours of the morning, because you’ll get beaten up by people with iron bars, if you’re lucky.

Now I’m back in Oxford, worrying about all the conference deadlines coming up. I’ve got a paper to finish by the end of next week, and I haven’t got any results yet. Wouldn’t life be much easier if you could stay on holiday all the time? We have the usual technical meeting and conference reports from around the globe, including Graeme Sweeney’s much awaited article on the historic first BMVA summer school held this spring at the University of Surrey. Another first is Philips Research Labs becoming the first corporate member of the BMVA. Congratulations to them, and let’s hope they’re not the last!

Biological Vision

Image Analysis for Biologically-based Applications

Twenty-nine scientists met in the splendour of the conference room at Wrest Park, Silsoe to discuss the application of image analysis to biologically based problems. Biological objects present particular difficulties to image analysis because of their variable nature and the unstructured environments in which they are often found. The aim of this meeting was to bring together researchers with an interest in image analysis for biological applications. We hoped to focus attention on the particular problems associated with biological objects. In addition we wanted to publicise, to those working in biological research fields, the strengths and benefits of image analysis.

The meeting was co-sponsored by the Biotechnology and Biological Sciences Research Council, who contributed towards the cost of holding the meeting at Silsoe. Six speakers presented seven talks on scientific research in a variety of subject areas.

John Marchant, from Silsoe Research Institute, started the proceedings by describing ways of using prior knowledge to predispose models to a given shape. His approach incorporated knowledge into snakes by varying the stiffness and connectivity between snake points, in order to give it a predefined initial shape. By using other items of knowledge about the structure of the objects of interest, he succeeded in getting good boundary location. Applied to images of pigs, the technique was used to locate key features such as the ham and shoulder.

Filiberto Pla, from University Jaume I, Valencia, Spain, described ways of assessing colour image data, which were independent of illumination levels. The methods he described gave good discrimination between the colours of ripe oranges, leaves and sky in outdoor scenes. This work was done as part of a larger project, called CITRUS, to automate the process of harvesting ripe citrus fruit. He described a dichromatic model of colour that reduced a colour measurement to just one parameter. He also described a technique for dealing with incomplete data caused by occlusion.

Norval Strachan, from the Central Science Laboratory at Torry in Aberdeen, described a fast, accurate method of measuring the shape, size and colour of different species of fish. His goal was to provide North Sea research vessels with a machine for sorting a catch by species. He was also interested in classifying fish into different sub-stocks to determine which

part of the sea they originated from. A video of the first prototype showed the system working at sea.

Tim Cootes, from the University of Manchester, spoke about active shape models for the interpretation of biological images. He described a method of learning shape from a set of example images. He then went on to describe methods of fitting the shape models to objects in scenes. His measure of fit used local grey level templates at points on the boundary and a multiresolution search. Animated sequences of Tim's face models, as well as other examples, illustrated the method at work.

Heidi Bloemen, from the Catholic University of Leuven in Belgium, described the use of images from two cameras to obtain topographical data. By projecting a pattern on to the object, she was able to extract depth information, which in turn enabled her to measure the volume of different parts of the object. The objects in question were pigs whose market value was determined by their size and fat distribution. She also described work on tracking *Daphnia* with a view to monitoring water quality.

Derek Sergeant, from the University of Leeds, was also interested in tracking. His aim was to monitor the welfare of poultry. Although the work is at an early stage, some of the difficulties in tracking multiple objects simultaneously were discussed. In particular he spoke of the problems associated with changing backgrounds and maturing chickens.

John Marchant presented the final talk of the day on the use of a robust estimator that gave information about structure from very sparse data. His goal was to extract row information in field crops to provide guidance information for an automatic vehicle. His approach used prior knowledge about the geometry of the vision system and a specialised Hough Transform technique.

The lecture session was followed by a tour of two laboratories at Silsoe. The first was the Robotics and Automation Group laboratory where Tony Hague demonstrated how the techniques described by John Marchant were used to guide an automatic vehicle. The second visit was to the Image Analysis and Control Group laboratory where two pieces of work were demonstrated. Robin Tillet described the use of variable shape models for assessing animal behaviour. Nigel McFarlane talked about determining plant structure with a view to automating the process of pruning.

In addition to the seven lectures and the tour of the laboratories, five posters illustrated some of the other work being done in the Image Analysis and

Control Group at Silsoe. These included work on neural networks, texture analysis, x-ray imaging and colour models.

For further information contact

Christine Onyango,
BBSRC.
Email: Christine.Onyango@BBSRC.AC.UK

BMVC'95 Update

**The University of Birmingham, UK,
11-14 September 1995**

Preparations for this year's conference are well under way. In response to the call for papers there were 102 submissions. The Programme Committee meeting was on 31st May and the authors of papers selected for presentation have been notified.

As usual, the conference is a single-track meeting with both oral and poster presentations. In addition to the contributed papers there will be invited talks and a pre-conference tutorial. There will also be a commercial exhibition of image processing hardware and software.

This year's invited lectures are to be given by Professor Nicholas Ayache from INRIA, Sophia-Antipolis on "Computer Vision, Virtual Reality and Robotics in Medicine"; and by Professor Jitendra Malik from University of California at Berkley on "Smart Cars and Smart Roads". The pre-conference tutorial on "Artificial Neural Networks and Computer Vision" is going to be presented jointly by Dr Roger Boyle and Professor David Hogg from Leeds University.

Registration forms and further information are available from the Conference Chairman:

Mr D. Pycock,
BMVC'95,
School of Electronic and Electrical Engineering,
The University of Birmingham,
Edgbaston,
Birmingham B15 2TT,
Tel: 0121 414 4285
Fax: 0121 414 4291
Email: BMVC95@bham.ac.uk

Information is also available on-line via ftp at: [peipa.essex.ac.uk](ftp://peipa.essex.ac.uk) (in [peipa/info/conferences/bmvc](ftp://peipa.essex.ac.uk/peipa/info/conferences/bmvc)) and on WWW at: <http://peipa.essex.ac.uk>

Representation Meeting

Representation in Natural and Artificial Vision Systems, 24th May 1995

The one day BMVA technical meeting brought people from a wide range of backgrounds together to discuss the issue of representations in both natural and artificial visual systems. Most of the talks focussed on the representation of faces, providing a strong theme through the day. All of the talks were well received by the audience. Each speaker was very happy to answer questions from the audience both during and after their talk, which created a very relaxed and informal atmosphere throughout the meeting.

Mike Oram (St Andrews) gave an overview of electrophysiological responses of cells found in the temporal lobe of monkeys. Some cells respond selectively to particular profiles of a face ("face cells"), whereas other cells have been found which respond to a human walking in a certain direction ("walk cells", combining form and motion). Mike also reported on studies showing that the neural responses of cell populations can be good predictors of behavioural data: The relative numbers of monkey cells tuned to particular profiles of a head closely matched the relative amount of time spent by human observers examining different profiles of a clay head under experimental conditions.

Hilary Buxton (Sussex) presented the problem of representing faces from the AI perspective. A big problem for face recognition systems is that of coping with invariances: recognising the same face despite different viewing angles, expressions, and so forth. Hilary presented recent work on addressing both static and moving face recognition, using Radial Basis Function networks and recurrent (Elman) networks. Emphasis was placed on classifying a temporal projection through "face space", rather than just trying to recognise individual snapshots from a temporal sequence.

Tim Valentine (Durham) provided a review of the psychological theories of face recognition. Tim then described his model of encoding faces as points in a multidimensional space. The model is theoretical in that the dimensions of the space are not known,

but the model provides a unified account of various psychological phenomena in face recognition. Recent experiments into the "own race bias" effect were discussed, with an interesting result that faces from another race may be recognised, but with a (slight) cost for recognition of "own race" faces.

The emphasis of Ian Craw's (Aberdeen) talk was placed very much on comparison of different coding strategies for discrimination in face recognition systems. The normal eigenfaces of Pentland and Turk were first discussed, before examining alternative processing of the faces. It was found that the decomposition of a face into shape and shape free (morphing of the images to produce a "texture") vectors before applying principal components analysis (PCA) gave the best recognition results. Additionally, it was found that applying a caricature operation (accentuating the atypical features of a face) before PCA improved results.

Charles Heywood (Durham) described the processing of colour information in both monkey and man. P.E.T studies have revealed a region in the human brain that is thought to be the "colour centre". Lesions to this area of the brain induce achromatopsia – a severe, if not complete loss of colour vision. Area V4 of the monkey is thought to be the equivalent monkey colour centre, but lesions to this area did not induce achromatopsia. Instead, recent studies indicate that area TE may be the monkey "colour centre". Charles also discussed the interesting result that although achromatopsic patients lack the ability to discriminate colours, they can use isoluminant colour variations to extract shape information. This provides evidence contrary to the popular notion that one area of the visual cortex performs one function.

The final talk of the day was given by Aaron Sloman (Birmingham). Aaron asked the question of "What's vision for anyway?", demonstrating vividly with many amusing visual examples, that vision is not just a question of "photons in, shape descriptions out." Instead, Aaron argued for a Gibsonian approach to vision, placing it in context of a total system, interacting with a whole host of other sources of information, such as other senses and expectations. In this view, vision cannot be regarded as a module that presents visual descriptions of the world that are then processed by a separate "cognition" module. Instead, vision is fully intertwined with many other processes, and is proving hard to untangle.

The Chairman, Jim Stone (Sussex), concluded the day by reiterating the importance of each of the

many approaches to studying representation in visual systems. In particular, Aaron's talk provided an approach which was very different from the rest of the day's speakers by emphasising that as well as "looking down the microscope" it is also important to "look up through the telescope".

Stephen Eglén,
Cognitive and Computing Sciences,
University of Sussex.
Email: stephene@uk.ac.susx.cogs

Teaching Machine Vision

One day BMVA technical meeting, 12th April 1995

The aim of this meeting was to provide a discussion forum on the teaching of machine vision in the UK. Tim Ellis, chaired the meeting, and presented the findings of THE BMVA Teaching Survey, which were (glibly) "we all work with Khoros or Visilog, and use either Sonka's or Gonzalez' books!" If you try to ask what is the "British Average Syllabus", then you'll find we all spend one and a half hours teaching every topic!

Dave Cooper of Manchester gave two brief presentations on the theme of Software Environments for Machine Vision. The first was an overview of the object oriented approach used at Manchester by a wide range of students, within the Wolfson Image Analysis Unit (Chris Taylor's mob). Their experience suggested that given sufficient help, most could fairly quickly pick up the nuances of class libraries and object oriented ways. He then gave a rather briefer description of the US megalith known as IUE (Image Understanding Environment), an object-oriented approach to populating the world with image processing and computer vision algorithms and data structures. The scale and complexity of the system suggested that all the well known problems associated with large multi-library, object-oriented packages, the compiler and linking requirements were likely to restrict its use.

After lunch, Simon Perkins of Edinburgh University gave an account of HIPR: a hypermedia image processing reference package that had been JISC funded. The motivation for this work was that textbooks have plenty of information, but only few, or low quality, images. Image processing software, on

the other hand, has little in the way of real tutorial help, but does have plenty of interaction. He wanted to get the best of both worlds, and has developed a system that has a library of some 800 images, developed from about 200 raw images, using some 50 image processing operations. It does seem an interesting teaching tool, with distribution on CD-ROM imminent. As Simon said, "wait a week or so for notices!"

Adrian Clark from Essex spoke about Network facilities and PEIPA, the Pilot European Image Processing Archive with URL

<http://peipa.essex.ac.uk/ftp/ipa/README.html>

It seems our Adrian is a bit of a squirrel who's hoarding away as much software, source code, test images, and reference material as he can. But he can only do it with your help. His aim, is to stop you re-inventing the wheel, to provide a good set of well tested algorithms in the language of your choice, and a good set of test images too. From their README screen you can see how to get your paws on his hoard of goodies, and how you can let others see the results of your labours. He would like too, to get sets of teaching notes for all to use, and a common set of example projects which can be tested on the image set. How about this? The BMVA would like a series of lectures notes on a variety of some 20 topics. The house style would be for each to have 5 pages of introductory material, 10 core material, and 5 pages that would lead into research topics. Could you help? If so, get in touch with the committee and offer a topic.

Lastly John Anderson of British Aerospace gave his industrial perspective on what he would like our graduates to know. He went through the regular set of general skills on segmentation, model based vision, and data fusion, through maths, geometry and statistics, and on to sensors. He gave a set of general requirements. He wanted innovative recruits, with good organisational skills, patient, tolerant, communicative people, who could sell their ideas, extrapolate the needs of production If that list looks long, then it is actually a cut down version of his overhead, but it all seems to boil down to people with good inter-personal skills. Now everyone seems to be looking for those recruits, not just us vision mob.

The day sparked off more than the normal amount of questions and comments, even though the attendance, by BMVA standards, was low at about 15-20 people. The "cry" came that we should make more use of PEIPA, and that perhaps we (the BMVA) could get funding through the HEFCE to produce high quality teaching material for all to use.

Geoff Dowling,
City University.

Summer School

Visions of a vacation in Surrey

"Hi, would you like to come on a fun filled educational vacation school at the University of Surrey, all expenses paid (except travelling costs) to learn about Computer Vision ?" ... of course ... who wouldn't ?

The BMVA proudly held their first computer vision school at the University of Surrey, sponsored by the EPSRC. Thirty-five of the brightest, finest or maybe just available 1st and 2nd year doctoral students, attended ready to receive the enlightenment from the BMVA's best and brightest, in a series of lectures and tutorials aimed at providing the students with an understanding of all that computer vision can and should entail.

Having arrived half an hour before the start and finding my way to the reception point, I bumped into the Leeds Mafia something I recommend only to the strongly willed and those game for staying up late. First event of the day was the "this is what I'm researching speech", an amazing 4 minutes to summarise your research hosted by the Gamesmaster - John Illingworth. It's amazing what some people actually research ... from gerbil brains to knee replacements. We finished 3 and 1/2 hours later and it was off to find the plush accommodation kindly laid on for us by Surrey before the first feeding frenzy of the week. Having heartily tucked into all that was being offered, we returned eagerly for the evening session at 20:00 (I kid you not) where we gathered for the literature review exercise. Groups of strangers openly discussed, ravaged and argued over a set of pre-distributed papers before compromising so that we could get to the Student Union (SU) on time. After closing time we headed back to the Leeds flat for a serious discussion and the odd game of cards (or was it vice versa ?).

Tuesday arrived and the first lecture was given by Tim Ellis who discussed the image formation issue, backwards from the notes. This was followed by Bernard Buxton who discussed at great pace the issues in image modelling and what ensembles of images could help you do, apparently everything would

be ok and his notes would follow soon (thankfully they have). After the lunch recess came Adrian Clark, bounding about the virtues of good software writing and management, and the odd advert for Unix and the Psion3a. We followed this by presenting our group work from the previous night, in a question and discussion session. We were so enthusiastic that we delayed our evening meal by half an hour just to fit all the questions! "The what the hell are wavelets?" still stands out as the best question and this was posed by the reviewer! Having nourished ourselves, we set off on an adventure to find the lab building where Adrian Clark would demonstrate the merits of Khoros. This must have been the most eagerly awaited session since it gave us the first chance to read our e-mail in the two days we'd been at Surrey! It was rumoured some even tried a quick Usenet session, not I. Having successfully crashed Khoros three times, my partner, I and the majority readily retired to the SU to lament on the days activities. This lamenting ended at about 02:30 after learning several new card games and the Oxford cork trick.

Maria Petrou woke us up on Wednesday to the problems of doing things ad hocery and not applying theory. Whether we actually believed her, time will tell. Neil Thackery proceeded to take us into the deep waters of Statistics and Estimation followed by a good explanation of Neural Nets and some of the difficulties they entailed. Time for lunch. A change in viewpoint was presented by Geoff Sullivan who discussed how we ourselves "see" in his Biological Vision lecture. The day ended with the school dinner where we, the chosen, gathered round and said our good words about the school and those who organised it. A fine impromptu speech was given by the Gamesmaster with a retort coming from Geoff Sullivan before once again we tackled the Surrey SU. Unfortunately this was the last we saw of the Gamesmaster who's organisational skills seem to have taken the most from him as he retired ill to bed (never to be seen again!).

For myself, Thursday proved to be the most interesting day with a superb presentation from Andrew Zisserman on Stereo Calibration and Object recognition. This was preceded by Roberto Cippola who discussed the use of Differential Geometry in Computer Vision. Joseph Kittler had begun the day with his topic of Mid-level vision. Two interesting presentations from Chris Taylor and Rob Blisset enlightened us on the transition from "Research Lab to Industry" and what actually was required in the dreaded

real world. The evening entertainment was a visit to the local dancehall (minus the Leeds Mafia - why?) where we proud students displayed our prowess for drinking and dancing and how these have to be mutually exclusive events. Return to the campus we debated till 04:00 as to whether it was worth going to bed. Having decided it was, the vast majority of the debaters failed to make it to the first lecture on Friday - David Murray's Motion and Active Vision, primarily due to a desired lack of activity and motion to be found in one's bed. Those who did attend spoke meritorious wonderment. Our curtain closer was performed by Josef Kittler who discussed some valuable points on (self) assessment and reporting research results. From this we headed to our final lunch where tearful goodbyes and the discussion of who, from our motley crew would lead the "British" in the field of Computer Vision. No doubt, the answer will come in time.

I have to say that I found a lot of the course applicable to my area of research and made contact with people in similar areas which has given me at least one new avenue to follow. The course should not be looked on as introduction to Computer Vision by the students but as an insight into some of the intermediate and advanced levels of research currently being undertaken. An expansion of the image formation and reconstruction issue, which I felt was lacking especially for those people doing medical imaging, should be considered for future schools.

So if by chance you are visited by a request to attend one of these schools, then say yes and watch out for the Leeds Mafia - they know some devilish card games.

Graeme J. Sweeney,
Human Computer Systems,
University of Essex,
Wivenhoe Park,
Colchester C04 3SQ.
Email irnbru@essex.ac.uk
WWW: <http://hcslx1.essex.ac.uk/~irnbru>

CVRMed95

In early April 1995 the First International Conference on Computer Vision, Virtual Reality and Medical Robotics (CVRMed95) was held in Nice,

France. The conference chair was Nicholas Ayache (INRIA).

The idea of organising the conference was conceived following the success of a similar, but smaller scale AAAI symposium held in Stanford, CA, during March 1994. At that meeting it was felt that there was a need for a conference to report on new directions in medical robotics and medical image understanding. Why yet another vision conference you may ask? The simple answer is that there is not currently a conference that provides a forum to disseminate new results in computer-assisted medical image interpretation and that brings together the robotics and clinical communities. At the one extreme, medical imaging conferences general focus only on clinical issues and the results of extensive clinical trials rather than theory, methodology and application of image formation, measurement and understanding. At the other extreme, established computer vision and robotics conferences generally only have a small session on medical applications at best and have not encouraged publication of high quality experimental work. It is quite standard practice at vision conferences to see application papers which include experimental results and conclusions based on an analysis of 1-6 (specially chosen) images. If medical image understanding and robotics is to be taken seriously by clinicians *and* used in real applications the vision community not only has to take the task of vision algorithm verification more seriously but also has to work more closely with clinicians and hospitals to understand medical issues and clinical requirements and constraints. CVRMed95 was created to provide a step towards achieving this goal.

The CVRMed95 programme focused on three major areas; medical image understanding, registration and therapy planning, simulation and control. The 48 long and regular papers were organised in a single track session. In addition there were two poster sessions and a panel debate. As noted above, one of the novel features of the event was that it aimed to provide an international forum for industrialists, academics and clinicians to exchange ideas. A total of 213 people representing 16 countries attended the 3 day conference. Of these, approximately 30 attendees were industrialists from 15 different companies.

There was naturally a strong European presence at the meeting with most of the leading European medical institutions and university research groups represented. The presentations spanned a range of topics broadly organised into the categories of feature extraction and segmentation, shape analysis, atlases, registration, and therapy planning. The emphasis

was very much on volumetric image interpretation using MR and CT images. One reason for this is clearly that many of the useful vision applications in the near future are going to be applications such as image-guided therapy which involve, for example, registering segmented 3D CT/MRI data of a patients anatomy onto video images of the patient in the OR. Thus, many of the presentations discussed progress in developing techniques to support this direction of work including results on automatic image segmentation for delineation of organs and quantification of lesions, image-to-patient and multi-modality registration methods, and generating computationally meaningful and concise shape representations of anatomical structure (atlases). Much of this work is still in the feasibility and early validation stages of development. This is primarily due to the fact that the field is still very much in its infancy. For one thing it is only in the last 3-5 years that the quality and resolution of CT and MR data has been adequate for quantitative image interpretation. Hence, much of the current effort is focusing on establishing how well generic computer vision techniques work on medical data and has not reached the point of understanding how specific image interpretation techniques perform at meeting clinical goals. In addition to the volumetric imaging work there were a few papers describing progress on image understanding based on 2D medical imaging (X-ray and ultrasound). This included reports on progress in X-ray image enhancement for breast cancer screening, and 3D reconstruction of vascular structures using digital angiographic (X-ray) and intravascular (ultrasound) data. Here a current new promising trend is to make use of modelling the image formation process for image enhancement and to help improve the overall quality of segmentation of subtle features.

The forum provided an opportunity for some of the larger medical image analysis/robotics labs to present a review of some of their work and to predict the future directions of the field. Russ Taylor gave an excellent presentation on medical robotic surgery research undertaken at IBM during the past 7-8 years. The University of Hamburg's work in a wide variety of areas was cleverly presented through a video originally put together as a historical review of medical imaging as a celebration of the 100th anniversary of Roentgen X-ray. Takeo Kanade from CMU chose to look into the future rather than to review the past and described a number of futuristic ideas about potential directions in microsurgery and the development of materials for biomedical applications.

The conference was a clear success. Planning is already under way to hold the next CVRMed95 in

1997. The possibility of making the conference a biennial event alternating with MRCAS (Medical Robotics and Computer Assisted Surgery) is also being considered. For those of you interested in learning more about the technical content of the meeting the conference proceedings is published by Springer-Verlag under their Lecture Notes in Computer Science series (volume 905).

Alison Noble,
Oxford University.
Email: noble@robots.oxford.ac.uk

Technology Transfer Initiative

Presentation of university generated vision applications to a wide industrial audience

Brunel University is one of the 8 UK teams sponsored by EPSRC (CDP) under the Technology Management Initiative. Our project addresses problems of the transfer of technology between universities and industry. We are currently identifying a series of case studies in which selected research products will be presented to an industrial audience.

We are inviting researchers, from UK HEIs, with a "product" in the areas of vision, image processing or pattern recognition which might have reached the stage of transfer to industry, to contribute to our industrial panels. The amount of time required from the contributors is minimal.

Our search is focused on manufacturing products and (in order to exploit the skills of our team) we are particularly interested in identifying case studies in robotics and machine vision. All the information disclosed will be subject to a confidentiality agreement.

The research project is interdisciplinary, involving two academic departments at Brunel, as well as the Brunel Research Services Bureau. Research staff are drawn from the Centre for Research into Innovation Culture and Technology (CRICT), in the Social Sciences Department, and from the Department of Manufacturing and Engineering Systems (M&ES). The research team includes Professor Steve Woolgar, Dr Bob Grieve, Mr Peter Race, Mr Jean-Noel Ezingard, Dr John Kirkland, Ms Janet Vaux and Dr Paula Gomes. Industrial collaborators include North

West London TEC, West London Leadership, West London TEC, Human Centred Systems.



For more information, please contact:

Dr Paula Gomes,
Research Fellow,
Brunel University,
Dept. of Manufacturing and Eng. Systems,
Uxbridge, Middlesex UB8 3PH,
Tel: 01895 274000 ext. 2935,
Fax: 01895 812556,
Email: paula.gomes@brunel.ac.uk



REALLY-USEFUL COMPUTER VISION

OR: MANDEL TRUCCO'S UNSOLVED PROBLEMS - MUCH MORE FUN THAN ADWAY'S!

PART 1 OF 12,435 559,814: BIOLOGICAL APPLICATIONS



MANUEL TRUCCO '95



ICCV '95

The Fifth International Conference on Computer Vision was held in Boston, USA this year. The conference ran from Tuesday June 20 until June 23, and the was held on the MIT campus, where most people had decided to stay. There was an early scramble for electric fans as the rooms were not air-conditioned and we had arrived in the middle of a heat wave.

The first paper of the conference was a call to standardization. The authors are attempting to create an "Image Understanding Environment" that will enable researchers to move images and data between different formats. Although this type of work is important, it is difficult to write an interesting paper on the subject. The rest of the morning was taken up by recognition, calibration, navigation and shape recovery.

The first part of each afternoon would be taken up with poster sessions. The ICCV had accepted twice as many posters as oral presentations. Unfortunately, the amount of room available for poster presentations was quite restricted. The resulting scrum to reach the best posters was more reminiscent of the January sales than an academic conference. The final session concentrated on shape recovery.

Tuesday evening saw a reception for the conference. The beer flowed and the food was plentiful.

Face and gesture recognition on Wednesday morning was an interesting session, maybe because the video presentations were so good. This was followed by papers on curve matching, pose and correspondence. After lunch, there was another scrum for more posters, amongst them many more on face and human gesture recognition. A session of deformable models finished the day.

The conference banquet was held in the evening. We decamped to the local marine centre where lobster was on the menu, which we ate surrounded by jellyfish and penguins. The Marr Prize was shared between two papers — the first on improving scanned images from books by Wada *et al.*, and the second on a theory of Specular Surface Geometry by Oren and Nayar. It was announced that the next ICCV conference would be held in India in December 1997, and that the venue for the 1999 conference would also be decided during the week.

Thursday morning saw sessions on colour, motion and learning. Demetri Terzopoulos presented a controversial paper on "Animat Vision" in this session. Undoubtedly the best video of the week, but many people doubt that his work is relevant to computer vision.

After the final poster session, there was a session on representations and geometry. Philip Torr must get a mention for a presentation so laid-back it was horizontal. He got a cheer from the sizeable French contingent for announcing the French win over England in the world cup, and a larger round of applause from the British contingent for announcing John Major's resignation.

There were two sessions on motion and stereo on Friday morning and the conference was over. The conference in 1999 will be held in Greece. The biggest rumour of the week has it that the ICCV will separate from the IEEE. This might make future conferences somewhat less expensive to attend.

Dave Reynard,
Oxford University.