

Newsletter of the
**BRITISH MACHINE VISION ASSOCIATION
AND SOCIETY FOR PATTERN RECOGNITION.**
Vol 1, No 2, May 1990.

Editor: Dr John Illingworth, Dept. of Electronics, University of Surrey, Guildford. GU2 5XH.
Tel: (0483) 571281 ext 2299, Fax: (0483) 34139, email: illing@ee.surrey.ac.uk

EDITORS NOTE

I hope that readers enjoyed the first BMVA News. Firstly, reports of a couple of typographical errors which crept into the last issue: Jeff Sullivan should have read **Geoff** Sullivan and M. Sander should have been **M. Sandler**. Apologies for these mistakes. This issue is dominated by reports of meetings which have taken place during the last few months: ECCV90 and the inaugural BMVA technical meeting.

I would also like to announce that there is a change to the newsletter book review editor. Following a period of sterling service, Horace Ip has taken up employment in Hong Kong and therefore resigned the office. However, Dr Maria Petrou of Surrey University has kindly agreed to replace him. The book review editor solicits copies of new books from publishers and then issues them to reviewers who provide reports for inclusion in the newsletter. In return for honestly assessing a book a reviewer keeps the review copy. Maria would like to hear from any people who might like to review books in the vision and pattern recognition areas. Interested individuals should contact her directly on (0483) 571281 ext 2265.

VISION ARCHIVE

The Turing Institute has just set up a vision archive which is a freely available service. Although we only have a couple of hundred images and movies at present, it is intended to acquire as many as possible of the standard scientific images used by the research community. The archive can be accessed over JANET using the following information (Consult your systems administrator if you are not sure). Username is "guest" and the password should be your email address (it must include an @). The path of the vision archive is specified with

"<GUEST>". At your site an example would be: *cpf -b -n "<GUEST>ls-lr"@uk.ac.turing ArchiveListing*. The above would then prompt the user for a username and a password. Note that on your system the command may not be "cpf". Also note that you should set BINARY transfer (as done above with the -b flag, the -n flag merely notifies you of success or failure via E-mail). The file "*ls -lr*" gives a complete listing of the files in the archive and the README provides a short text description of the contents of each directory. Here's another example showing how to access a file lower down the file system hierarchy: *cpf -b -n "<GUEST>/gels/gel11a"@uk.ac.turing gel11a*.

This would give you a hips image of a gel electrophoresis plate containing a stained protein "fingerprint" of blobs.

All the images stored in the archive follow one of the following formats: HIPS, TIFF, GIFF, or no-header (but with associated read-me's to give header information).

The archive is keen to acquire new images that may be freely used in the public domain. Anyone wishing further details about the service or who would like to donate images or movies should contact: Barry Shepherd, Turing Institute, George Street, 36 N. Hanover St, Glasgow, G1 2AD, E-mail: barry@turing.ac.uk, tel: +41-552-6400, fax: +41-552-2985.

Peter Mowforth, Turing Institute

BMVA CHAPTER MEETINGS

Scottish Chapter

The dates and venues for the next three Scottish BMVA meetings are:

Wednesday 28th June at Heriot Watt University

Wednesday 3rd October at Aberdeen University

Wednesday 12th December at Edinburgh University

All meetings will start at around 2.00pm and will involve around five speakers as well as providing demonstrations of vision projects. Further details can be obtained from Peter Mowforth of the Turing Institute.

Northern Chapter

June 14th 1990 two talks:

D.R. Bailes "*Grey level symmetric axis transforms.*"
Dept. of Medical Biophysics, Manchester University.

S.J. Fanson

"*Use of image analysis to measure dimensional changes during paper manufacture.*", UMIST.

July 12th 1990

M. Browne "*Surface Inspection.*"
DIAZ, UMIST

Talks are held in Lecture Theatre 4, Medical School, University of Manchester.

BMVA MEETING REPORT

Mechanisms for the direction of attention in human and machine vision.

This was the first meeting organised by the newly born BMVA, in collaboration with the Applied Vision Association. Its purpose was to bring together people from the two societies whose interests lie in vision, computer on one hand, natural on the other.

There were about fifty participants and the meeting was interesting as well as entertaining. The most heroic talk was that by Keith Ruddock who had to put up with a slide projector which had its own ideas as to which slide should be shown next! After the end of the talk the projector with the temperamental personality was promptly replaced by another one which stubbornly refused to show any slides at all, until the kind intervention of a member of the audience! Apart from that, the other highlight of the day was a most entertaining talk by John Mayhew who started by telling us that he hates giving talks, he hates travelling and he hates London! We also learned that the forced mating of the Universities with Industry by Alvey is called necrophilia by some people! (I wonder, who is the dead body in this case?)

Entertainment aside, we heard Vernon Dobson from Oxford discussing associative networks. Apparently, conjunctive nets are most efficient for sparse pattern vectors while Hopfield nets are better for dense vectors. The work reported concerned the study of Bayesian nets which behave like conjunctive nets for sparse vectors and like Hopfield nets in the other extreme.

Keith Ruddock, from Imperial College, reported on some experiments concerning the identification of a target amidst a clutter of reference patterns. It seems that

when the patterns are well defined, either because they are solid or because they are made up from a large number of dots, the time taken to identify the odd one out is independent of the number of reference patterns. On the other hand, when the patterns are defined in a very sketchy way, by a few dots each, the larger the number of reference patterns the longer it takes for the odd one out to be identified. The conclusion was that in the former case our vision system processes the picture in parallel, while in the latter case it processes it in series. Glyn Humphries from Birmingham University, later on that afternoon, reported on similar experimental results but he tried to explain them using a Boltzmann machine-type model of the vision system: The brain requires more time to process the sketchy patterns because there is a greater number of competing possible interpretations and the neural net takes longer to converge to a solution.

John Mayhew described how the "Commode" of Sheffield University works! "Commode", of course, is the name of their mobile robot which is called that because it is a modified wheel chair. In particular, John Mayhew discussed the system of ground plane obstacle detection. He explained how they divided the information the robot has about its environment into a series of levels of increasing awareness, of the form: "There is something there", "Keep an eye on that obstacle", "What exactly is this object here" and so on. The gross information for the first level, which is used to focus attention, is provided by Mallot's algorithm which is about one thousand times faster than the full stereo algorithm used for more detailed tasks.

David Foster of Keele University reported on experiments which show that all aspects of orientation detection in humans are consistent with the presence of two types of filters, one for detecting horizontal and one for detecting vertical lines. He was followed by John Findlay from the University of Durham who presented experimental results designed to determine how the size and position of the saccadic movement of the eye are affected by the size, relative position and contrast and the number of multiple targets in the peripheral vision.

The session after tea was devoted to medical images. We heard Alistair Gale from the University of Nottingham telling us how easily the experts miss important features in radiological images and our hearts sunk at the high percentage of occurring misdiagnoses! However, our worries were alleviated by the next talk, when Sue Astley from Manchester reassured us that automatic analysis of mammograms can be significantly improved by combining evidence from more than one cue used in the process!

On that optimistic note the meeting promptly finished. It offered a good balance of interesting papers from

both sides, machine and natural vision, and the speakers showed respect for the audience by using clear slides and transparencies and by keeping to their times.

Maria Petrou, University of Surrey.

ECCV90

The First European Conference on Computer Vision took place in Juan les Pins on the Cap d'Antibe, France between 23 and 26 April. The main conference was attended by about four hundred participants. This report presents some personal impressions of the main conference.

Opening the conference, Olivier Faugeras pinpointed two unifying themes. The first concerned the topical focus of vision research in the areas of shape and motion interpretation. The second theme was methodological and concerned the increasing mathematical sophistication required for the study of vision. In many ways the second theme was the dominant one. The conference highlighted a number of very nice pieces of theoretical research. The only weakness was probably a lack of results to show what impact this work was likely to have on practical vision systems.

Topically, the conference concentrated on the analysis of shape and motion. Shape description having four sessions and motion analysis six. The remainder of the programme was devoted to image-features, matching and stereo. All of the work was of a high standard, for the most part presented in a lucid and accessible way. The bulk of the programme took the form of a plenary track; only three of the thirteen sessions took place in parallel. Sixty papers were presented orally with an additional twenty allocated to a poster session. Authors were predominantly based in the European Community, Scandinavia and North America.

Work on 3D shape description and recovery covered both range-volumetric and contour-based approaches. Some noteworthy ideas included the use of deformable part-models, charting surface structure using smoothness constraints, the inference of projective invariants, and, the tracking of extremal shape contours using snakes. Steve Zucker deserves a mention for the dexterity with which he animated the parts and protrusions approach using playdough.

Motion and tracking are by no means my area of interest. Even so I found some highly interesting talks in these sessions. John Aloimonos argued for the 'purposeful' paradigm in active vision. This favours the evolution of goal-directed vision processes over more generalised ones. The approach was nicely demonstrated on a real-time tracking problem. In a talk that spanned both shape and motion, Olivier Faugeras presented a method for tracking 3D curves using geometric features of the spatio-temporal surface.

Work on matching and recognition was represented at a variety of levels. At the lower-level there were several papers on accumulator-based model matching. Higher level approaches included work on consistent labelling and the application of ATMS for feature-labelling.

The edge-industry showed no sign of recession and dominated the sessions on image-features. Work reported included multiscale contour linking, a comparison of different optimality criteria and work on 3D edge detection. One idea gaining in popularity was the use of phase rather than amplitude as the basis for characterising image features. This concept was evoked in connection with scale-space representation and also in stereo feature matching.

Faugeras's second theme was exemplified by ideas surfacing in different topical areas. One example was the extent to which speakers were using methods from differential geometry for both shape and motion analysis. Evidence for the interest in this area could be found at the conference book-shelf where Jan Koenderinks "Solid Shape" began to look decidedly dog-eared. Another such theme was the increased interest in Bayesian approaches. One paper argued for their use as an alternative to symbolic strategies in high-level vision. Mean field-theory approaches presented an interesting way of reconciling stochastic and mechanistic approaches to scene reconstruction.

The organisers are to be congratulated on putting together a high-quality technical programme and providing Europe with a vision forum of international excellence. My only suggestion for improvement would be for the addition of invited review talks to the programme. Notwithstanding this, the good news is that there is to be a second conference in the series. It will be held in Genoa in 1992. Hopefully the precedent of a Mediterranean venue will be strictly maintained throughout the series.

Edwin Hancock, Rutherford Appleton Laboratory

BRA WORKSHOP REPORT

This article reports on a workshop held on the day following the first European Conference on Computer Vision (ECCV). It was one of three workshops arranged by members the Working Group on Vision, an ESPRIT II/Basic Research Actions (BRA) supported amalgamation of researchers throughout Europe involved in computer vision, and was principally organised and chaired by Jim Crowley from LIFIA (IMAG) Grenoble.

The format of the workshop was arranged to provide a mix of both European projects and current progress from some of the US laboratories. A total of eight talks were scheduled, each allocated a 30 minute presentation period, followed by 15 minutes of discussion. This

proved to be a very successful format, given a judicious curtailing of some of the presentations by the chairman, and provoked some very stimulating discussions on the current state-of-the-art.

The chairmans' opening remarks provided an overview on some of the research at LIFIA into active vision. In particular, he briefly outlined the details of their contribution to a BRA project (BRA 3038), Vision as Process (VAP). VAP is a project to develop a modular vision system, exploiting mechanisms such as focus of attention in spatial, temporal and model spaces, and to look at the control of discrete and modular computer vision algorithms in a goal-directed fashion. Crowley described a software communications architecture which would exist between different modules (robot, matching mechanisms, model base etc.) based on UNIX sockets.

The first scheduled presentation was given by John Aloimonos, from the Univ. of Maryland. He developed an argument for inferring object shape and structure from motion, without solving the general problem in its entirety, widely appreciated as a non-trivial task. He argued for "purposive and qualitative vision", based on the ability to perform complex visual tasks without reconstructing the world. Perhaps his most controversial conjecture (partly supported by Ed Riseman in the later discussion) criticised the poor performance of the majority of optical flow algorithms. This was intriguing, in light of the fact that some 10 papers on or using optical flow algorithms were presented at ECCV earlier in the week.

Dana Ballard from Univ. of Rochester provided a refresher on his work into the development of an anthropomorphic eye/head sensor. This was illustrated with a number of video demonstrations, one indicating the value of vergence control in such stereo imaging systems. Interest was expressed in the motor drive capability of the 'eyes', quoted as up to 300 degrees per second, which allows the system to match the very fast movements of the human visual system. A further demonstration showed the use of a crude colour tracking algorithm which characterised the overall colour distribution of model objects as 2D colour histograms, and then matched these 2D histograms over the image, in order to detect the presence of possible colour matches. The algorithm, developed by a graduate student, was not robust and Ballard emphasized that they had not solved the colour consistency problem, but that never-the-less, under certain constraints, it might prove to be a valuable cue generator in searching colour images.

The third and fourth presentations described further aspects of the VAP project, with Eric Granum from Univ. of Aalborg providing a more detailed overview of the project, expanding on the principal features and aims. Jon Eklundh from KTH, Sweden, provided further emphasis of Ballard's research by describing the

development of a sophisticated eye/head system along the lines of the Rochester Groups' work, but including some mechanisms for independently controlling the focus of attention of the two 'eyes', allowing for both vergence and independent eye movement. In addition, a zoom facility was provided for each 'eye', a feature that might be used to emulate the high-resolution capability of the fovea.

After lunch, Ernst Dickmanns from Univ. D. Bundeswehr Munchen described a major theme of his research with a presentation entitled "Control of Perception Guided by High Level Spatio-Temporal Models". The talk provided details of his groups' 4-D approach to dynamic vision, treating temporal events as simply an extension to the three spatial co-ordinates and emphasized the use of selecting small sets of 'optimal' features which could be tracked using high-level predictive models. Accompanying videos ably demonstrated the efficacy of this approach, and particularly impressively, the ability to control a vehicle driving on a road at up to 90 kph.

Giulio Sandini (Univ. of Genoa) presented research into a novel image sensor based on the mammalian eye: a circular, space-variant sensor with a high-resolution centre (the fovea) and decreasing towards the periphery. Of interest was the character of the resultant signal, given a particular mapping into a 2-d co-ordinate system. The talk emphasized an important aspect of current computer vision research: that although the vast majority of systems are based on the use of an orthogonal 2 or 3-d co-ordinate system, this is unlikely to match the mapping system employed by the mammalian eye/brain system, and that the structure of this system must fundamentally influence the mechanisms that are used for visual perception.

The penultimate presentation, by Schmuël Peleg, visiting the SRI-Sarnoff Research Centre at Princeton, described a hierarchical (image pyramid) approach to estimating motions. The technique is intended to cope with a variety of different motions in the scene, initially estimating the most prominent motion in the scene at the lowest resolution, and iteratively refining this estimate within finer resolution images. This algorithm is then 'recursively' applied to estimate other motions in the same fashion once the most prominent has been characterised. Again, a most convincing video demonstrated the operation of the algorithm, allowing accurate estimates of ego motion, translational and rotational motions in an animated scene.

The workshop closed with a presentation by Roger Mohr from LIFIA, who described a simple and elegant technique (based on projective geometry) for determining 3-d scene geometry from a single image, without complex calibration, using only a small number of defined reference points within the image.

The workshop provided a well-focused theme and an excellent selection of presentations. It was enthusiastically received by all who attended.

Tim Ellis, City University, London

MACHINE VISION CONFERENCE

A Machine Vision Conference is being organised by Blenheim Online in association with Image Processing Magazine and NEL (National Engineering Laboratory). It will be held on 7th and 8th June at the Queen Elizabeth II Conference Centre in London. The conference aims to be a primer to machine vision systems and how they can be used in industry. The emphasis is on current capabilities and problems - it is **not** a highly technical conference aimed at academics. The conference cost is £626.75 (excluding accomodation costs). A special rate is available for academics. Further details can be ob-

tained Blenheim Online Ltd. Tel: 081-868 4466 or Fax 081-868 9933.

DIARY

22 May 1990 IEE Colloquim on "Electronic images and image processing in security and forensic science", Savoy Place, London.

22 May 1990 BMVA Meeting on "Markov Models", 1 Victoria St, London.

17-21 June 1990 10th International Conference on Pattern Recognition, Atlantic City, U.S.A.

11-13 July 1990 SERC/DTI Transputer Applications Meeting, Southampton.

29 July - 18 August 1990 SUSSP on Pattern Recognition and Image Processing in Physics, St Andrews, Scotland.

25-27 September 1990 BMVA Conference, Oxford.