

Newsletter of the
**BRITISH MACHINE VISION ASSOCIATION
AND SOCIETY FOR PATTERN RECOGNITION.**
Vol 3, No 1, January 1992.

Editor: Dr Edwin Hancock, Dept. of Computer Science, University of York, York. Y01 5DD.

Tel: (0904) 433374, Fax: (0904) 432767, email: erh@uk.ac.york.minster

EDITORS NOTE

This month's newsletter begins with an account on the recent ESPRIT Conference and its satellite activities from Bernard Buxton. There are also reports from the recent BMVA Technical meetings on Kalman Filtering from Tim Ellis and Simon Brock-Gunn, and, on Combining Evidence in Computer Vision from Geof Nichols.

Although it is still winter, it is time to cast your thoughts ahead to the autumn and the forthcoming British Machine Vision Conference. This will take place in Leeds from 21-24 September. Conference Chairman Professor David Hogg provides details of this year's arrangements. However, if you are unable to wait so long for a serious vision fix, Giulio Sandini provides a preview of the 2nd European Conference on Computer Vision which takes place in the springtime in Genoa.

Finally, the BMVA Executive Committee has been reviewing its current technical activities. In order to introduce new items into the programme we are keen to receive proposals from anyone interested in organizing one-day topical meetings. We also hope to introduce an annual student papers competition. For more details please read the Technical Activities Section of the newsletter.

The ESPRIT Conference and Exhibition

The ESPRIT Conference and Exhibition, held annually at the Palais des Congrès in Brussels towards the end of November (this year from the 25th-29th), covers the results and achievements of the whole ESPRIT information technology programme. Machine vision thus forms only a small part of the activities at the conference but there are three reasons why, as a vision researcher, it is nevertheless of interest to report briefly on my two and a half day flying visit.

Background: ESPRIT and vision research

First, the scale of the ESPRIT programme is so large that over twenty major projects involving the development of machine vision systems and techniques or their

application have now been funded in the Information Processing Systems, Computer Integrated Manufacturing and Basic Research areas. In addition, a small number of image coding, compression and database projects have been supported in the Advanced Business and Home Systems area to complete the wide variety of vision and image processing projects in the ESPRIT portfolio. Work supported by the CIM and Advanced Business and Home Systems areas has always been strongly applications oriented whilst the IPS area has catered for more fundamental research on vision systems and architectures.

Second, ten of these projects were represented at the exhibition with stands featuring new hardware (P940, DMA), vision systems for the control of robot vehicles (P2502, VOILA), vision incorporated in a surveillance system (P1560, SKIDS), a vision development environment for industrial applications (P2592, VIDIMUS), vision for the control of a robot manipulator (P5227, CIM-SEARCH), and vision systems for inspection in robot manufacturing (P2091, VIMP), finishing (P2640, ICI) and assembly (P2017, TRIOS).

Finally, there was the Information Technology Forum on Thursday, 28 November. This usually features a more political debate about the ESPRIT programme's achievements and future perspectives. This year, the theme of the Forum was information technology in the working environment but the session was obviously going to be lively in view of the perceived crisis in the information technology industry [1] and recent criticism of the effectiveness of the ESPRIT programme. In addition, it was particularly interesting for me, as a vision researcher, to have an opportunity to assess the direction of future activity as fundamental research work on vision techniques, systems and architectures no longer appeared in this autumn's 1991 workprogramme and vision work was already dispersed amongst a number of applications projects in the IPS and other areas.

Vision and its advanced architectures

There were four talks in this specialist session on Tuesday, 26 November. Three of them, by Thomas Skordas (iTMi), Eric Theron (MS2i) and myself (GEC-Marconi) presented work on the DMA, SKIDS and VOILA projects (Theron's talk on advanced architectures actually spanned all of these projects). In the fourth, Patrick Stelmaszyk (again iTMi) described vision research activity in Japan. Patrick's talk was in many ways similar to the one delivered at the BMVA meeting on the development of vision systems for robot vehicles held last May at the GEC-Marconi Hirst Research Centre, but more specifically addressed to the conduct of vision research in Japan. Three important aspects of the organisation of vision research in Japan emerged from this talk: firstly, that industrial applications used (as to be expected) the most simplistic vision techniques; secondly, that advanced research was being carried out in the laboratories of leading large Japanese companies; thirdly, that the Japanese Vision Club plays a crucial role in co-ordinating research activity between the universities and industry. The Club, run by three leading university figures with a restricted industrial membership, seems to owe its success to the sharing of detailed information, not only of members' successes, but also of systems that have failed and the lessons to be learnt from them.

The exhibition

According to M Carpentier, the Director General of CEC-DG XIII which runs the ESPRIT programme, this year's Conference and Exhibition were the largest ever with over 2000 people attending the conference and 125 projects exhibiting (25% more than in 1990). In addition, the exhibits were generally of a larger scale and higher quality than in the past and this was also true of the vision systems on display. Thus, four of the five vision projects supported by the Information Processing Systems division (DMA, SKIDS, VIDIMUS and VOILA) **all** had live exhibits whereas in 1990 only VOILA mounted a live exhibit. Similarly, the vision applications in the CIM exhibits were also live.

All four of the above IPS vision exhibits could be highlighted: DMA for the new hardware (FIR filter, NMS, linker, polygonal approximation and stereo fusion special purpose and DSP boards); VOILA for the Genoa DIST robot acting as watchdog over its exhibit and wandering around people's feet; SKIDS for its monitoring of the movement of people around the exhibit; and VIDIMUS for its graphical programming interface which is similar to, but I was told offers more flexibility than the KHOROS system from the USA. In the CIM area, a vision system was controlling a robot "bar-tender" in pouring drinks. Unfortunately, it would have required a combination of this and the SKIDS exhibit to offer the refreshment service required in the rather overheated underground exhibition area. Other CIM exhibits featured high precision, optical calibration systems for robot arms (e.g. for inserting windows in car

doors) and a customised VLSI chip (OPTIC) combining image segmentation, alignment correction and filtering capabilities for real-time object recognition.

The IT Forum

The IT Forum on Thursday, 28 November, opened with presentations by M Carpentier (Director General, CEC-DG XIII) and F M Pandolfi (Vice-President, Commission of the European Communities) both of whom emphasised the difficulties currently facing the European electronics and IT industry. However, both were very "bullish" about the success of ESPRIT and the need for an expanded future programme with an overall target of 6% (equivalent to 5.4 billion ECU) of the EC's financial resources for all R & D activities. Current programmes consume 3.5% of the EC's resources with ESPRIT taking 37% of this *for an industry that, by the year 2000, is predicted to be the largest industrial sector in Europe, with an annual turnover of 300 billion ECU, almost three times that of the automobile industry*. However, Pandolfi stressed that future programmes must be carefully targeted to meeting the objectives on which the IT industry depends, both on the supply (technology) *and* demand side. In particular, he emphasised the need for an indigenous European semiconductor industry linked to the needs of other key technologies and for large-scale targeted projects in microelectronics, software, CIM, display technology and high performance computing.

These opening remarks on IT R & D policy were followed by several talks on the themes of "IT and the citizen", "IT and its users", and "IT and the future", together with highlight presentations on "CD-I technology", "CIM" and "advanced UV lithography". These talks touched upon such diverse matters as the integration of information technology in the aerospace industry, IT and the law, and the importance of meeting users' needs. However, here I shall only briefly mention the talk by S Purvis, Editor in Chief of ITN, and the CD-I presentation.

Purvis described how ITN's news gathering had been almost completely transformed by the introduction of computer technology from the use of an old iron spike for filing incoming news items when he began his career to present systems that could electronically trace and recall news items **and** the decisions taken about them days later. In fact, he stressed that *the systems were only really constrained by the imagination and limitation of the suppliers* (note that ITN own one of the largest suppliers, BASIS in California). Thus there are still many outstanding requirements for easy access and (screen touch) editing of picture sequences that must offer many opportunities for vision and image processing systems developers.

Finally, the highlight of the whole event was undoubtedly the Philips CD-I technology presentation memorable for the rather dry delivery of G Bastiaens, Director of Philips Consumer Electronics, who led the presentation from the podium and the stark contrast with

his colleague who, with unstoppable enthusiasm, put the technology through its paces from the floor of the chamber. Important points in addition to the technology itself which uses full motion video, *MPEG*, coding were that the market was predicted to be 12 billion ECUs by 1995 and that Philips had made agreements with many Japanese companies and major publishing houses in order to try to avoid a "standards war" as occurred when video tapes were first introduced to the home market.

Bernard Buxton
GEC-Marconi, Hirst Research Centre
Tel: 081 908 9081
email: bfb@uk.co.gec-rl-hrc

BMVC92: Call for Papers

This year the British Machine Vision Conference will be hosted by the University of Leeds and will take place between 21-24 September 1992. This event has established itself as the main UK conference for machine vision and related topics. The emphasis is on UK research being undertaken through national or international collaborative projects, providing a forum for the presentation and discussion of the latest results of investigations. Papers from other nations, especially those collaborating with UK groups, are also very welcome.

Contributions are sought on any novel aspect relating to machine vision and pattern analysis, including:

- image processing and feature extraction
- practical applications
- object recognition and scene analysis
- model based coding
- reconstruction of 3D shape
- architectures
- advanced pattern analysis
- active vision
- computational issues in perception
- motion analysis
- robotic vision and sensor fusion
- neural networks

Full papers not exceeding 10 pages (approx. 5000 words if no figures) should be submitted for review. Papers will be accepted either for oral presentation or for presentation as posters. All papers accepted will appear in the Proceedings. Papers will be reviewed by the BMVA Committee.

Three cash prizes will be given, one for the podium paper and one for the poster paper judged by the programme committee to make the best scientific contribution, and one for the paper judged to have the most commercial potential.

The important dates are as follows

- **SUBMISSION DEADLINE: 1st May 1992**

- Notification of acceptance: 8th June
- Deadline for camera ready copy: 17 July

Papers should be submitted to the Conference Chairman:

Professor David Hogg
School of Computer Studies
University of Leeds
Leeds LS2 9JT
tel: 0532-335765
fax: 0532-335468
e-mail: dch@dcs.leeds.ac.uk

The provisional conference fees for early registration before 17 July 1992 are £220 for members and £235 for non-members; after this date the corresponding fees are £240 and £255. Included in the cost are full residential accommodation, the conference banquet and a printed copy of the Proceedings. For non-members, the fee also provides free BMVA membership for 1993. As in the past, a selection of the best papers will be published separately in a special issue of *Image and Vision Computing Journal*.

Finally, an innovation in this years programme will be a half day pre-conference tutorial to be held on the afternoon of 21 September. The topic will be "Invariance in Computer Vision" and the tutor will be Dr Andrew Zissermann of Oxford University. More details will be circulated when they become available.

Vision Applications of Kalman Filters

The meeting opened with a slight departure from the normal format with a tutorial introduction to the Kalman Filter given by Hugh Durrant-Whyte (Oxford University). The tutorial covered the basic concepts and background to the Kalman Filter and was underpinned by three main examples chosen to demonstrate typical applications - a simple linear system, a simple nonlinear system and a very complex system with more than 60 parameters. Dr Durrant-Whyte used the exemplars to emphasise their value in deriving a thorough understanding of the behaviour of the sensors and the way they are modelled.

Following this tutorial was a presentation from Rachid Deriche (INRIA) who gave an overview of the use of the Kalman Filter in robot tracking and video vision at INRIA. Drawing for examples upon the ESPRIT P940 project and the work of Nicholas Ayache, he emphasised the use of appropriate geometric representations which would be minimal, unambiguous, complete and differentiable. Finishing on an appropriate note, he described his joint work with Olivier Faugeras on a Kalman Filter based token tracking algorithm (from P940) which labels edge based feature data in sequential scenes from a camera undergoing ego motion.

The afternoon session was opened by David Castelow from GEC Hirst. He described a complex Kalman Filter used to estimate the ground plane for a mobile robot vehicle from a monocular image of the scene. His talk dwelt on the complexity of the filter (an extended Kalman Filter) to determine the camera pose with respect to the vehicle, and the vehicles position with respect to the ground plane. The final EKF required some 14 state variables.

Tim Atherton described research at Warwick University into the use of the Kalman Filter for a range of image processing problems. His talk related to the use of a hardware architecture which could be used to selectively update a scene flow model. A further example of the use of the Kalman Filter was given by its use in estimating the range of an object (a car) in a dynamic scene.

Roland Marslin (Reading University) described the use of a Kalman Filter to constrain the dynamics of vehicles being tracked around a road traffic roundabout as part of the ESPRIT VIEWS project. By adopting practical constraints on the possible motions of the vehicles, the model tracking was shown to give significantly improved results. This had the marked benefit of reducing the problem of apparent vehicle kangarooing to one of minor side slip.

Following the enthusiastic and optimistic reports from the previous speakers, John Porrill (Sheffield University) ended the day by casting a timely and soberingly pessimistic view of the capability of the Kalman Filter in vision, highlighting the danger of overestimating its ability to solve a wide range of problems without proper regard to its limitations. He presented examples of its use for tracking a vehicle around a curved track. The inability of the filter to reliably follow the vehicle for this rather simple task indicated an underlying limitations in an apparently simple implementation.

The meeting provided a stimulating and varied range of uses of the Kalman Filter. The tutorial session provided a useful opportunity for the uninitiated to become more familiar with the characteristics of the technique. In addition, its diverse use in current vision applications was shown to reflect its popularity as a current “flavour-of-the-month”.

Tim Ellis and Simon Brock-Gunn
City University.

Combining Evidence in Computer Vision

The first BMVA meeting of this year addressed the problem of “Combining Evidence in Computer Vision” and drew on treatments from the fields of AI and mathematical statistics as well as vision itself. Our need for an understanding of how evidence combination should proceed is driven by the presence in image analysis of a hierarchy of object descriptors. Our reasoning leads us to use evidence from different levels of this hierarchy. In his introduction to the meeting Josef Kittler listed some of the issues which would be addressed by a theory

of evidence combination. In particular, a working theory would need a background framework of principles as well as actual methodologies for dealing with a range of symbolic and quantitative evidence types and their uncertainties. As we saw, generalisable methodologies and heuristic principles have already evolved within vision to deal with certain particular applications. The first two speakers described existing matching algorithms.

Josef Kittler described a graph based approach. This is a natural description of a compound object in terms of the attributes of, and relationships between, its parts. It facilitates a heuristic combination of these different pieces of “evidence in one graph-matching support function. Hard constraints introduced by numerical translations of symbolic evidence can be represented by assigning very large penalties through this function. He gave several example applications of this robust though slow matcher including successfully matching the curvature properties and adjacency relations between surfaces segmented from a range image against a model object graph.

Chris Taylor and Tim Cootes (speaker) described an algorithm suitable for recognising objects that need not be of constant shape, such as circuit-board elements. They showed how to summarise a training set in terms of a model with given shape uncertainty. This is done by identifying the few major modes of variation in an eigenvector decomposition of the shape in the much larger space of all boundary point locations. In order to carry out the matching, different filters (edge, corner) were used to detect the various different types of boundary point. The output of each filter was smoothed to create a potential for the existence of that point type over the image. Evidence was combined heuristically, by adding the weighted potentials to produce an overall potential for a boundary to lie at any point on the image. A relaxation algorithm made the match by moving the model template in space, and along its principle modes of shape variability, downhill to the minimum of the boundary potential.

The problem of hierarchical evidence combination is central in bringing together data and context in the presence of prior knowledge of locally consistent label structure. Taking edge labeling as his example, Edwin Hancock described a Bayesian relaxation process for locating consistent interpretations in a resolution hierarchy. The potential combinatorial complexity of the method is lifted by explicitly listing a local symbolic constraints in a so-called dictionary. This provides a means of integrating evidence from different levels of the hierarchy; inter-level consistency constraints are explicitly represented by imposing a tree structure on the dictionary items at different resolutions. The framework was demonstrated for the coarse to fine resolution tracking of edge information under conditions of extreme noise.

In this same vein of integrating contextual information, John Bridle gave us a pedagogical introduction to Bayesian inference in causal nets, attempting unsuccessfully to avoid taking the example of the hidden Markov model used in speech processing. In cases where context

is modeled by local conditional probabilities, and neighbourhoods build up a tree like structure, as is the case in cause-and-effect processes, knowledge of the state of parts of the tree is carried through the network by the Bayesian algebra in an explicitly calculable way.

The Bayes networks are pure methodology and furnish a useful yet less abstract example of the kind of system espoused by John Fox. He pointed out that the kind of deductive rule based inference from evidence described by several speakers was analogous to an existing AI paradigm which he had worked on in attempting to automate diagnosis in medical applications. In this approach a database consisting of image plus rules is interrogated through a higher level interface. Alternative inferences are ranked by preference with qualitative knowledge reflected in the ranking order and quantitative knowledge used to determine absolute scale. This methodological framework would facilitate any rule based approach to evidence combination.

A problem common to all evidence combining methodologies is that of uncertainty propagation or "error tracking". It is not clear how to do this when the evidences are of qualitatively different type. Bob Fisher presented a system for geometric reasoning, which he represented as a network similar to those John Bridle had described. The network updated bounds on the position of model entities by applying a set of rules to some partial evidence till no further restriction is possible for given evidence. Errors must be tracked through this process; the speaker gave a comparison of interval and distribution based error representations. He repeated the point made by John Fox that interval based errors are more suitable for translating uncertainties from symbolic evidence ("the cup is on the table") into numerical form and added that they are also convenient for representing the hard geometric constraints imposed on model object positions by surfaces. However experimentally it was found that error tracking was both faster and less divergent when one assumed a given family of distributions for the uncertainties as prior knowledge. Not surprisingly, interval based errors, which are not as local as a normal distribution, and deal with the 100% confidence interval, propagate errors in the most conservative fashion.

As far as principled approaches to evidence combination go, the speaker who took the bull by the horns was Jim Baldwin. Taking evidence to be presented in the form of fuzzy sets he defined a mass for each element in a set as the proportion of people who would agree that the element belonged to that set. Each piece of evidence provides a set of candidate interpretations which can be intersected with the sets from other evidences to restrict alternative interpretations. He suggested masses be assigned in the intersection set according to a "principle of least commitment", given consistency with masses in the original sets. Using some simple examples he showed that, while this does not always give the most general mass assignment, it takes the form of a simple algorithm which may be sufficient for practical purposes.

It was apparent from the material presented and the discussion of the 50 or so participants, that while methods for combining evidence are being developed to deal in a heuristic way with vision problems as they arise, researchers have yet to take advantage of the group of theories (such as that raised by Jim Baldwin) from mathematical statistics that provide principled approaches to combining evidence across the hierarchy of image descriptors.

Geof Nichols
Department of Electronic and Electrical Engineering
University of Surrey

ECCV 92

The Second European Conference on Computer Vision (ECCV-92) will be held in Santa Margherita Ligure, a pleasant resort in the Italian Riviera, from May 18 to 23, 1992 in conjunction with 5th Centennial Celebrations of the discovery of America by Cristoforo Colombo.

The main objective of the Conference is to provide a forum to present and discuss recent research advances in the field of Computer Vision.

Building on the success of the First ECCV held in Antibes in April 1990, the Call for Papers has generated 308 proposed contributions from 28 different countries (about one third of the papers were submitted from outside Europe). The requirement set by the Programme Committee to maintain a single-track Conference has limited the number of accepted papers to 16 long papers, 41 short papers and 48 posters, representing 19 different countries. Of these papers one third are coming from non-European countries.

Many important aspects of Computer Vision will be discussed at the Conference, including: low-level features, texture, color, active vision and binocular heads, stereo and motion estimation, reconstruction and representation of curved surfaces, recognition. The programme has been organized such as to allow reasonable space for informal discussions. Three VHS multistandard video equipments and 3 Sun workstations (one SparcStation IPC, one IPX and one SparcStation 2) will be accessible by all participants willing to present the results of their research activity. Participants interested in participating in the experimental sessions are advised to contact the secretariat or Dr. Massimo Tistarelli at the following e-mail address: tista@dist.dist.unige.it.

The proceedings of the Conference will be published by Springer - Verlag in the "Lecture Notes in Computer Science" series and will be distributed to all the participants along with a copy of the Video Proceedings. The scientific programme will be articulated in three parts:

- May 18th: an "Esprit Day" will be organized to present ESPRIT projects active in the field of Computer Vision and its applications. Participation in the ESPRIT Day will be free of charge and will be open to everyone interested in vision including participants from outside the European Community. Organizations wishing

to present their project or to attend the meeting should contact:

Dr. Patrick Van Hove,
CEC DG XIII A4, Office BRE9-217,
200 Rue de la Loi, B-1049 Brussels, Belgium.
Fax: +32 2 236.8364
e-mail: pvh@dg13.cec.be

- May 19 to 22: ECCV-92 Conference. Besides the official programme the participants will be given ample space for informal activities such as the presentation of videos and the participation in the "experimental session" (demonstrating the results of software programs on Sun-workstations available at the Conference site).

- May 23rd: three workshops will be organized by the members of the ESPRIT Basic Research Working Group on Vision to discuss some specific topics of Computer Vision. Participation at these workshops is limited in number and requires a separate registration procedure. For additional information, contact:

Prof. James Crowley,
LIFIA-IMAG, 46 Avenue Felix Viallet,
38031 Grenoble, France.

The Conference has been organized by DIST - University of Genova, in collaboration with the "Consorzio Genova Ricerche" and INRIA and with scientific support of the "European Vision Society" and the Special Project on Robotics of the Italian National Council of Research. The support of Elsas Bailey, DEC Italy, Sun Microsystems Italy, and Sincon-Fase is also gratefully acknowledged.

For further information please contact
Prof. Giulio Sandini
DIST - University of Genova
Via Opera Pia 11a
16145 Genova
e-mail: giulio@dist.unige.it

Future BMVA Technical Activities.

The BMVA Committee has recently been reviewing its programme of technical activities. We would particularly welcome proposals from anyone interested in organizing a one day meeting of topical interest to our membership. The Association can provide assistance

with distributing publicity and arranging the venue for the meeting. Proposals should be made to the Meetings Coordinator Dr Tim Ellis whose address is given below.

An innovation in next years programme will be the introduction of a one-day graduate-student papers competition. This will take place in the Spring. Second and final year Ph.D. students will be encouraged to submit abstracts for 30 minute talks in the broad areas of computer vision, image processing and pattern recognition. Prizes will be awarded for the most original technical contribution and for the best presentation. This meeting will be co-ordinated by Dr Adrian Clark of Essex University; fuller details including a call for participation will be issued in a forthcoming newsletter.

The following meetings are currently scheduled to take place in 1992.

- 11 March 1992 "*Intelligent Image Databases*"
- 15 April 1992 "*Speech Processing*"
- 27 May 1992 "*Shape*"
- 8 July 1992 "*Colour and Texture*"
- 22-24 September 1992 "*BMVC 92 - Leeds*"
- 21 October 1992 "*Machine and Natural Vision*"
- 21 December 1992 "*Geographic Information Systems*"

The majority of meetings will be held at the Institute of Radiologists which is located at 36 Portland Place in Central London - just five minutes walk from Oxford Circus underground station. No registration fee is payable by BMVA members although a charge of £10 will be levied for non-members (unless they are members of a co-sponsoring organisation). Arranged lunches will be available on prior registration but a small fee may be payable for these. Fuller details including final programme will be mailed to members nearer the time of each meeting.

For further information about the BMVA meetings programme contact

Tim Ellis
Centre for Information Engineering
City University
Northampton Square
London EC1V 0HB
email: t.j.ellis@uk.ac.city