

BMVA News

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

Volume 5 Number 4
March 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 May 1995.

Contents

Editorial	1
DAS'94	1
La vie Française	2
ICIP-94	4
BMVC'95	5
ECCV	6
IPOT95	7
Photogrammetry meets Machine Vision	8

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.

Editorial

The Human Capital and Mobility fund (HCM), for those who don't know about it, allowed researchers from European countries to work in a foreign country for a year or two. If you were really lucky and came from a poor country, for instance Ireland, you could even get an extra year back home. People who came to the UK tended not to mention their salary, largely out of guilt, because it was typically 2 1/2 times bigger than home-grown post-docs. This is because they got their fat £30K+ salary *tax-free*. There are still some of these greedy rogues around, hiding in the corners of our research labs. Shun them, for they are unclean! Oops, I got a bit carried away there. Let's hope that with the forthcoming TMR (Training and Mobility of Researchers), which is replacing HCM there will be a more equitable arrangement. If not, I'll have to pack up my bags and make my own fat salary elsewhere.

As usual we have reports from conferences and meetings in faraway places that we would all love to have visited but didn't have the time or money. We also have details of BMVC'95 and ECCV'96. Bill Triggs has written an invited article about life as a researcher in France, and as is typical of him, has kept to the important issues: food, drink and, er, I think you missed one there Bill. Obviously he doesn't ski. I have achieved the seemingly impossible and put this edition together only a month overdue, and am now preparing for the greatest ordeal of my life: demonstrating a robot vision system to a bunch of kids at the science museum. Of course I will cheat to make it look more impressive, but they won't notice, will they?

DAS'94

IAPR International Workshop on Document Analysis Systems Kaiserslautern, Germany, October 18-20, 1994

DAS'94 was held during the week after the 12th ICPR and the contrast between these two events was quite noticeable. Not only in the climatic and cultural sense but also in the nature of participation. DAS'94 was much smaller (75 participants) and had a single track programme. This and the fact that the theme of the papers was more focused than of those in the ICPR ensured a more enjoyable attendance and more productive working group sessions. There was ample time for meeting with other people working in the field and discussing their work and ideas. I find this particularly useful since by just reading the papers one only learns what people were doing one or two years before the publication date.

The workshop attracted participants from both academia and R & D departments of major companies like AT & T, Fuji Xerox, Daimler-Benz, Siemens, AEG, Hitachi, Toshiba, NEC, Matsushita, IBM, Unisys and Elsag Bailey. It was evident that there is a fast growing interest in the field of document image analysis and understanding that was also underlined by the presence to my surprise of representatives of the U.S. and German governments assessing the state of the art and looking for projects to fund.

The quality of the programme was high with many interesting papers on methods and systems. The structure of the programme incorporated sessions on document segmentation, handwriting and line drawing recognition, data structures and control, classifier combination, system architectures and innovative techniques and applications. There were also demonstrations of systems for the understanding of business letters, processing of form documents, music score recognition, line drawing recognition and bank cheque verification. Finally, several working group sessions were organised in order to draw upon experiences and set goals for the future. One of the tasks was the identification of evaluation criteria for OCR and document analysis systems. Other topics discussed were line drawing and music recognition, handwriting, recognition of multilingual documents and NLP, new directions in document analysis and learning, and forms recognition. The needs of the market and user requirements, as well as possibilities for international collaboration were also explored.

According to the structure of the programme, I

was given the "honour"(!) to present both the first and the last paper of the workshop, the bottom line being that I had no time to relax! Overall the level of organisation of the meeting was impressive and this was also true for the social events. The most notable event was the workshop dinner at Hambacher Slop, a restored castle on top of a hill in the forest. As we were walking uphill towards the castle we were told that there was a surprise waiting for us at the top. It was true. When we reached the gates of the castle we were greeted by the sound of trumpets and drums played by costumed medieval knights. It was a memorable experience as it was dark, the castle was floodlit, and the view of the lights of nearby towns was spectacular.

After this and a brief tour of the grounds we went into the originally restored dining hall for a medieval dinner. There, I was in for a second pleasant surprise. After the initial speeches and between loud outbreaks of trumpets and drums playing medieval marches and hunting music, it was announced that the best student paper award was to be given to me. Concluding, I must thank everyone that supported me because it meant a great deal to me. Evidence of international appreciation of one's work is to my mind one of the most important things to any young researcher. I hope that BMVA continues to be supportive to students and also keep an open mind to newer and more applied research areas such as document image analysis.

Apostolos Antonacopoulos
Systems Engineering Group
Dept. of Computation
U.M.I.S.T.
P.O. Box 88
Manchester
M60 1QD
Email: aa@mp.co.umist.ac.uk

La vie Francaise

Your beloved editor Phil is planning his retirement and wants to know what life is really like in Yoorup...

It's 11:59 AM and I just managed to stagger out of bed after a heavy night's debugging, shower myself into qualitative semiconsciousness, eat BREAKFAST (quoi?) and trundle through town to work. 10 people say 'bonjour' and shake hands before I reach



the safety of my office. But to no avail: the sysadmins have characteristically upgraded our main file server so that every printer in the place is broken and my shell dumps core whenever I try to run a command, so I guess it must be time for lunch. This occurs between midday and 2 PM (inclusive) and afflicts the entire French nation very severely indeed.

There is a score of excellent small family restaurants within five minutes walk, but being an untutored barbarian I head for one of the four or five local bread shops to buy a flute for a sondweech. At first sight this may seem eccentric, but the confusion only really sets in when you realize that approximately half the bread shops in France call a flute a pain long and deny the existence of flutes (and vice versa), and that each shop sells at least ten entirely distinct and quite unique varieties of flute. The following encounter is typical:

- PATRONNE: Bonjour m'sieur.
- MOI: Bongdjoor, dje voodrai urn float seel vooz play.
- PATRONNE: (puzzled and slightly insulted) Err. . . nous avons les pains longs, moyennes et courts, farines, tordus, de siegle, complets et au levains, artisanaux et sans sel, M'sieur. Vous desirez exactement quel rapport largeur-longueur, fabrication et materiau?
- MOI: (nonplussed) Umm. . . donnay-moi urn FLOOT seel vouz play.
- PATRONNE: (sighing, caressingly takes a loaf) Certainment m'sieur, 'un pain long'. J'espere que vous serez tres content avec ce pain un peu croustillant et tres savoureux, qui etait fabrique avec la tendresse et les soins de deux cent ans de tradition familiale artisanale. Cinq francs trente s'il vous plait.
- MOI: (glad to escape so lightly) Mursee madamme, aw revwaar.
- PARTONNE: (sadly but proudly) Bonne journee, m'sieur.

Having done my bit to cement Anglo-French relationships² for the day, I take the hard-won flute under my arm in (a personal interpretation of) the traditional manner and scoot back to the bureau. On reflection, it's lucky I didn't want fromage in my sondweech. Not only do French cheese shops really sell 500 different varieties of cheese — in fact each has 500 *different* different varieties — but they are deeply hurt if you fail to enter into an intimate and meaningful personal relationship with each and every one of them. One can easily find oneself saying things like 'the Chevre S^{te} Minabele is jolly fine this season, but it's been a bad year for Rocher S^t Pierre'.

It is also an invariable law of nature that a little old lady will push into the queue just in front of you and proceed to buy 100g each of thirty different cheeses, taking 5 minutes of gentle coaxing to choose each one.

Back in the office, a small but devoted crowd gathers to watch me eat my sandwich. I've only been here 10 months so people are still astonished to see someone eating lunch at his desk. It's lucky they're tolerant of foreign oddities or I'd long since have been lynched for reading at the table — a treasonable offence in France. As it is they mostly just titter 'bon appetit' and shrug sadly. They knew the English were eccentric, but what is to be done with someone who doesn't even take *food* seriously?

Lunch eaten, I get down to the day's work of breaking up lumps of computational geometry for the man in the EC. Geometry being intractable stuff, this involves a great deal of half vocal swearing similar to that available in any lab in the U.K. Perhaps this is why people say things like 'Gosh, your French must really be improving'. In fact, it's surprising how far one can get with the language skills of a retarded gerbil and an innate talent for cultural insensitivity. However, although one can lead a successful life as a barbaric outcast with only a few hundred words of broken argot, it generally takes foreigners at least half a lifetime of diligent application to construct their first sentence of *good French*, and the other half to remotely begin to understand the intricacies of French social relationships. And even the French don't understand their bureaucracy, whose astounding obfuscatory ability puts even Whitehall to shame.

But wot yer don't know don't hurt yer, and in the mean time there are real advantages to being a (not actually) computer vision researcher in Europe. Although the beer tastes funny and you won't get to see many striped shirts and berets these days, there are excellent cafes and restaurants where you can get good food without being ritually demeaned by the staff; there is a whole new culture to bask in, real mountains, snow, and occasional glimpses of sunlight; and there is even a dim recognition that actually making things might somehow be a more worthwhile activity than asset stripping existing industries, and that research might actually have a valuable part to play in that process.

Bill Triggs,
LIFIA, Grenoble.

²The chief bond being one of mutual puzzlement.

A Snakecharmer in Texas

First IEEE International Conference On Image Processing, November 13-16, 1994, Austin, Texas, USA

Austin, the capital of Texas, is a combination of university town, hi-tech centre, and music venue. Within the city limits are about half a million people, but Austin retains a small-town atmosphere. With over two thousand restaurants and bars to choose from, Austin has a laid-back attitude to life, and a serious attitude to entertainment. The fifty thousand students at the University of Texas make up a regular audience for the music industry which has taken root in the downtown area around Sixth Street. Live bands offer a range of jazz, blues, and rock music every night, and most of the clubs sell surprisingly good beer from Texan breweries. Meanwhile, back at the conference....

ICIP-94 was billed as the inaugural International Conference on theoretical, experimental, and applied Image Processing. The conference was dominated by the latest methods for coding, filtering, enhancing and restoring just about every kind of image. There were also lots of new ideas for generating images using every medium from ultrasound to X-rays, and every apparatus from microscopes to radio telescopes. There were also plenty of applications for all these good ideas, everything from geographical to medical imaging. Medicine made a big impact on the conference, with presentations covering all aspects of biomedical image processing and its applications.

A couple of tutorials filled the first morning, with another two after lunch. I missed Robert Haralick's talk on mathematical morphology to learn something about image compression - Bernd Girod gave a helpful tutorial in which I finally found out what MPEG stands for (Moving Pictures Expert Group). Unfortunately, I didn't survive the afternoon tutorial on maximum likelihood image synthesis by Richard Blahut because it wasn't aimed at mathematically disadvantaged graduate students. Instead I walked over to Sixth Street in search of Mexican food and tequila, thinking that the alternative session on wavelets and sub-band coding would have been even more mysterious.

The high point of the week was a state-of-the-art exhibition of commercial image processing hardware, software, and accessories. One evening the accessories included free food and drink. I assume the idea was to get everyone near the exhibits; it almost worked. The British lunged for the bar; the Amer-

icans lunged for the hors d'oeuvres, presumably because they knew how bad the beer would be. The next day I took my hangover to the Matlab exhibit and they gave me a mouse mat to rest my head on.

The low point of the week was the conference banquet which, for reasons that escape me, cost fifty dollars a ticket. I'm a vegetarian so my expectations about conference food are pretty low, but this time even I was disappointed. When the food arrived it was not just unimaginative, it was almost undetectable! To add insult to injury we were then subjected to an award ceremony during which the wine ran dry. The last straw came with the after-dinner entertainment, a display of traditional Mexican music and dance that made The Three Amigos look competent. I escaped to Sixth Street to drown my sorrows....

When I wasn't eating and drinking I found time to concentrate on more difficult tasks like image segmentation and object recognition. With ten sessions running in parallel for three days I couldn't see all of the work. Instead I concentrated on what seemed most interesting at first glance. The poster sessions were particularly useful because, in a busy morning, I could visit all the stands and still have time to discuss the most interesting work with the authors. Even so, I missed quite a few talks due to over-indulgence in the free chocolate-chip cookies that appeared from nowhere. I did manage to see two of the three plenary sessions that took place first thing each morning. Peter Burt gave a talk on multi-resolution methods in computer vision which I remember vividly because I understood some of it. The next day Richard Baker gave an entertaining talk about digital video. Unfortunately, I was busy recovering from the banquet on the final morning when George Swenson got out of bed to talk about radio astronomical imaging.

Overall, ICIP-94 was both fun and educational. There were more than 900 participants, and if you believe the organising committee the conference succeeded in every way. If you get the chance take a look at the proceedings - three volumes, each of about 1000 pages. That's 20 pounds of paper - half my baggage allowance - covered with test images of Lena! Extra copies are for sale (IEEE members \$150; non-members \$200) from icip@pine.ece.utexas.edu. ICIP-95 will be held in Washington, DC between October 22 and 26; for advance information send an email message to icip95@ccr-p.ida.org.

Jim Ivins,
AIVRU, University of Sheffield.

BMVC'95

Sixth British Machine Vision Conference

11–14 September 1995 The University of Birmingham, UK.

The British Machine Vision Conference is the main UK conference for machine vision and related topics. High quality contributed papers are sought describing recent and novel research in the areas of computer vision, image analysis and processing or pattern recognition. Papers describing research being undertaken throughout national or international collaborative projects are particularly welcome. The conference is a single-track meeting with both oral and poster presentations. In addition to the contributed papers, there will be talks by invited speakers and a pre-conference tutorial programme (free to registered students).

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 sensory fusion
- neural networks

Four copies of full papers not exceeding 10 pages (approx. 5000 words if no figures) should be submitted for review. It may be possible to purchase additional pages and or colour prints. Two title pages should be included, one of which should be anonymous, containing only the title of the paper and up to 5 keywords. Papers will be reviewed by the elected BMVA Committee with a small number of specially coopted members. By submitting a paper to BMVC, the author(s) warrant that it, and any related paper with

essentially the same technical content, has not been and will not be submitted to any other conference during the BMVC review period. Papers will be accepted either for oral presentation or for presentation as posters.

A printed copy of the Proceedings will be available to delegates at the conference and a selection of the best papers will be published separately in a special issue of 'Image and Vision Computing'.

Deadlines and format

Deadline for paper submission	24 April 1995
Notification of acceptance	12 June 1995
Deadline for Camera-Ready copy	7 July 1995

Guidelines for the format of papers can be found on the PEIPA at:

ftp: [peipa.essex.ac.uk](ftp://peipa.essex.ac.uk) in `ipa/bmvc`
(login name: anonymous)

WWW: <http://peipa.essex.ac.uk/index.html>

Papers should be submitted to the Conference Chairman:

David Pycock
BMVC'95
School of Electronic and Electrical Engineering
The University of Birmingham
Edgbaston
Birmingham B15 2TT
Tel: +44 21 414 4285
Fax: +44 21 414 4291
Email: BMVC95@bham.ac.uk

FIRST ANNOUNCEMENT

ECCV '96, 14–18 April 1996, University of Cambridge

Following the highly successful conferences held in Antibes, Santa Margherita Ligure and Stockholm, the Fourth European Conference on Computer Vision will be held from 14–18 April 1996 in Cambridge, England. The Conference is to be held under the auspices of the European Vision Society (EVS) and the British Machine Vision Association (BMVA).

The programme will consist of a single track of the highest quality, previously unpublished, contributed

papers, delivered either orally or as a poster. Contributions are sought on new research on any aspect of computer vision.

All reviewing will be double blind by a Programme Committee of leading international researchers selected by the Conference Board and, as usual for the European Conference, the proceedings will be published by Springer-Verlag.

Deadline for submission of papers	1/10/95
Notification of acceptance	10/12/95
Deadline for camera-ready manuscripts	20/1/96
Conference begins	14/4/96

Conference Chair: Prof. Bernard Buxton (University College London)

Local Arrangements Chair: Dr. Roberto Cipolla (University of Cambridge)

Conference Board

N.Ayache	INRIA, Sophia Antipolis
M.Brady	University of Oxford
H.Burkhardt	University of Hamburg-Harburg
J.Crowley	INPG, Grenoble
E.Dickmanns	Universität der Bundeswehr, Munich
J-O.Eklundh	Royal Institute of Technology, Stockholm
O.Faugeras	INRIA, Sophia Antipolis
G.Granlund	Linköping University
D.Hogg	Leeds University
J.Koenderink	Utrecht State University
J.Mayhew	Sheffield University
R.Mohr	INPG, Grenoble
H-H.Nagel	IITB, Karlsruhe
B.Neumann	Hamburg University
G.Sandini	DIST, University of Genova
V.Torre	University of Genova

A call for papers will be issued in the Spring. For further details of the call, conference programme, workshop facilities, and registration and local information please contact the Conference Secretariat:

ECCV'96 Conference Contact,
42 Devonshire Road,
Cambridge CB1 2BL,
England.
Tel: 01223 323437
Fax: 01223 460396
Email: cc@confcon.demon.co.uk

Information is also available on the World Wide Web at <http://svr-www.eng.cam.ac.uk/Research/Vision/ECCV>.

Reduced registration fees will be available to students and to members of the EVS and national societies affiliated to the EVS.

In addition, as at previous Conferences, it is planned to hold a small number of specialised workshops on Friday 19 April immediately following the main meeting. Anyone interested in organizing a workshop is invited to contact the Conference Chairman (b.buxton@cs.ucl.ac.uk).

IPOT95

If you had any doubts about how much interest the industrial community has in image processing and machine vision then a visit to this years IPOT (the Image Processing and Optical Technology exhibition) at the NEC in Birmingham would have soon set you straight.

The exhibition resembled a sort of hi-tech market place where, if you hesitated by a stand for even a moment, someone would pounce on you and threaten you with sharpened patten until you bought their all-singing-all-dancing vision system. That is, unless you told them you were a student in which case they'd either offer you a job or just smile politely and disappear.

Because, as representatives of the BMVA, we had nothing to sell I initially felt a bit like a fish out of water but that feeling quickly changed as our stand attracted considerable interest. Neil Thacker and myself had expected a quiet morning, chatting to the infrequent visitor and drinking cups of coffee, but the reality was somewhat different. From the exhibition opening at 9.30am we didn't have time to stop and think, never mind drink coffee, so we were only too pleased to hand the reins over to Phil McLauchlan for a few hours during lunch.

The stand was visited by both researchers, who were at the exhibition to catch up on the latest technology, and industrialists who wanted to know what sort of work was going on in the research institutes. Both parties were enthusiastic about the activities of the BMVA (by the end of the two days of the exhibition around 100 BMVA information packs had been given out) so we can expect to see a quite few more faces at this years conference.

It wasn't until the afternoon of the second day that I got the chance to visit the other stands. As I had anticipated, all manner of paraphernalia was on show from cameras and frame-grabbers to the latest vision systems. All of this was pretty interesting but

for people involved in machine vision research, like myself, the visual inspection systems on show were probably the biggest attraction. Roughly these systems could be divided into two categories.

First there were the dedicated inspection systems which the vendor would assemble and program up for the particular task you had specified. Many of these could be described as real-time, providing 100% coverage of a production line but only performing very simple tasks such as template matching or blob analysis to ensure components were present and positioned correctly.

The second category comprised what you might call visual inspection tools where the user is provided with a range of image processing utilities to extract the information they want from their image data. These systems tended to be aimed at scientific and engineering applications where fairly detailed measurements need to be taken. A consistent framework has been adopted where the user is presented with a graphical environment for prototyping their program by clicking on tool buttons and specifying areas of interest with the mouse pointer. These actions can then be recorded in a script file which may be replayed at any time to repeat the list of actions on new images. Finally the script file can be replaced by dedicated code which calls the relevant library functions to produce a stand-alone application. I saw a nice improvement on this idea where the script language was C and could be replayed through a C interpreter. The final, stand-alone code was simply produced by compiling!

With the increasing availability of processing power at affordable prices it doesn't take a lot of imagination to see how some of the more sophisticated techniques, which academic research has developed, might be integrated into these systems in the near future. The scope of visual inspection would benefit greatly from results in the areas of texture analysis, feature extraction, object recognition, tracking and so on. It is this observation which highlights the need for groups, like the BMVA, to encourage collaborative projects between industry and academia. Given the level of interest generated at the show I, for one, am optimistic that we shall see a reasonable amount of work filtering from the research labs. into industry in the not too distant future..

A.P. Ashbrook
The University of Sheffield

Photogrammetry meets Machine Vision

One is often struck by how insulated different disciplines are from each other, even in the 90's. This BMVA one day meeting attempted to bridge one such divide, and was particularly well attended by photogrammetrists.

The meeting was opened by Mike Cooper of City University who gave a tutorial introduction to photogrammetric methods. It is a matter of routine for photogrammetrists to compute positions of points from several image views using a huge minimization algorithm called a "fibre bundle adjustment" [or iterative least squares to the rest of us]. The minimization simultaneously adjusts internal camera parameters and world coordinates to fit [already matched!] image coordinates. The procedure is well understood and returns a wealth of information including estimates of the variances.

For the average sloppy vision researcher the numbers game is impressive. Numbers like 900 Meg images, 0.01 pixel accuracy, and 2 days cpu time were all bandied about.

Faustin Banda from UCL gave a talk on facial shape recovery from stereo which would easily have been presented at BMVC. His topic was echoed later by Colin Urquhart from the Turing Institute, whose project has been presented at BMVC! This was followed by John Turner from Camera Alive who introduced the very relevant topic of making money. He made the distinction between applications that needed speed, and others that needed accuracy. The application that grabbed my attention was the reconstruction of a CAD model of an entire gas refinery from a couple of rolls of film, and a Hasselblad. [OK, and some software!]

Peter Muller of UCL gave an overview of the wide ranging work done in his group, including aerial imaging for maps and other remote sensing topics. He presented a table of comparative performance measures for several stereo algorithms.

Sometimes at a vision seminar I drift off to sleep, and awake wondering where I am. If the speaker is showing a slide with mammograms, satellite images of Greece, or fine art wood panels from medieval times then I am not much the wiser. Stuart Robinson of City went on to explain how he used photogrammetry to obtain precise measurements of small shifts in the shape of fine art wood panels.

Finally Tim Clark gave a talk to warm the heart of any camera calibration fundi. How to eke the last

0.01 pixel accuracy from your camera. Does your calibration algorithm consider camera temperature? Is your spot location sensitive to laser speckle? By how much can 12 bit imagery improve spot location compared to an 8 bit image? [Homework!]

Considering that stereo is at the heart of the discipline of photogrammetry, it was surprising that epipolar line received its first (apologetic) mention of the day at 15:50.

Finally I was struck by the dynamic nature of the field of photogrammetry. Although it is over 100 years old it remains active and industrially relevant. Since they are vulnerable to competition from other position sensing technologies, such as laser altimetry or GPS, they have a natural interest in ‘value added’ processing of their images, which leads naturally on to image interpretation. Many of the photogrammetrists present would in no way consider themselves “researchers in machine intelligence”, but it seems that their most exciting direction of growth will bring them to use more and more techniques that we might have thought the preserve of computer vision.

Andrew Stoddart,
University of Surrey.
Email: a.stoddart@ee.surrey.ac.uk