

BMVA News

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

Volume 11 Number 1
August 2000

Editor: Dr Paul Rosin
Department of Computer Science
Cardiff University
Cardiff, CF24 3XF
Tel: +44(0)29 2087 5585
Fax: +44(0)29 2087 4598
Email: Paul.Rosin@cs.cf.ac.uk

<http://www.bmva.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 October 2000.

Contents

Extracts from the Minutes	1
EPSRC Theme Day	2
BMVA Technical Meeting – Augmented Reality	4
BMVA Technical Meeting – Visual Surveillance	4
ECCV 2000.....	5
Future BMVA Technical Meetings.....	7
Motorola Research Lab Vacancies	7
University of Oxford – Research Posts.....	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.

Extracts from the Minutes

Extracts from recent BMVA Executive Committee Meeting minutes

IPOT

An exhibition stand was mounted at IPOT. A contract has been signed with IPOT for a stand for next year.

Web-site

Dr. Courtney suggested that the web-site should have a page with links to machine vision in the news, and has agreed to do collate the information and set it up.

Foresight Meeting

There were 53 attendees with Foresight panel members present. Richard Brook (SIRA) ran the meeting. The aim was to produce a brief document. This will be used as a basis for an article in IP Magazine.

BMVC Student Bursaries

The rules are:

- Up to 5 per BMVC conference
- Student must be a presenting author
- Student must be a member of BMVA

BMVC Prizes

- Science Prize (£500)
- Industrial Prize (£500) - funded by CRS Ltd
- UKIVA Best Demo Prize (£200)
- Poster Prize (£200)

Summer School

Once again, this went well and engendered positive feedback from the students. Someone else may take it over next year – a decision will be made later in the year.

EPSRC

**Engineering and Physical Sciences
Research Council**

Report on a theme day in Machine Vision and Image Processing held at the Institute of Civil Engineers on Wednesday 7th June 2000.

The EPSRC places increasing importance on evaluation of the research it funds; it is important to learn about the quality of our portfolio, how trends can be identified, what success stories we can recognise and gain knowledge of best practices. Typically, this is performed at the programme level (e.g. IT & Computer Science programme etc.) or the managed programme level.

However, it is clear that some topics cut across a number of EPSRC programme areas yet it is important to evaluate these topics in a co-ordinated manner. The mechanism chosen is the Theme Day, in which an EPSRC portfolio of grants relevant to the chosen theme is evaluated by poster presentation from the investigator(s) to an expert panel. Typically, grants that have finished within the previous year and current grants that have been active for at least 18 months (i.e. have had an opportunity to

make some inroads in to their research goals). Additionally, three international-leading experts were invited to give presentations; their brief was very much to give a forward-looking view of the subject.

Any evaluation process requires some sort of framework to judge/score the projects. The EPSRC has chosen the “QIPE” framework, where

Q = Quality

(i.e. Quality of the research)

I = Impact

(i.e. Potential for Impact on other research)

P = People

(i.e. output of trained people from the project)

E = Exploitability

(i.e. potential for the results to be commercially exploited)

The core output from the Theme Day can be summarised as Observations, Conclusions and Recommendations.

Observations

1. In general, the quality of research supported by EPSRC was found to be high, with the average quality tending towards nationally leading. There was also a high proportion (19%) of proposals rated as internationally leading. For example, 3-D imaging was identified as a UK strength as a result of recent EPSRC grants.
2. The very nature of the subject area ensured that the community has been effective at training postgraduate students and Research Assistants (RAs) who go on to a wide variety of careers, e.g. finance, software industry, academia etc, as well as a range of Vision companies (including start-up businesses). Partly as a result of this effectiveness there does seem to be a recruitment and retention problem in this area of Engineering and Computer Science.
3. There is an enormous opportunity for EPSRC to work with the other Research Councils to identify new areas for transfer of skills in to applications of Vision research (particularly NERC and MRC).
4. The impact of the research was disappointing. The panel sometimes felt that there was a tendency for some researchers to become insular in their outlook and lacked a knowledge and perspective of the wider vision community.

5. Whilst the potential for exploitation of UK research was high, the actual exploitation was disappointingly poor.
6. Some of the apparently poor take up in these application areas may be because the UK does not have the necessary infrastructure to fully exploit the results commercially. In this respect, the UK compares unfavourably to the US.
7. The panel also found investigators often did not consider the possibility of exploitation of the results until very late in the research programme.

Conclusions

1. The UK has a strong community in Machine Vision, producing high quality research. This was particularly true for those working in medical applications.
2. The overall impact of the research did not match the quality with particular weaknesses identified as the insularity of some researchers, and groups working on very specialised application areas.
3. In general there was a lack of focus on people and training by investigators. However, the people trained in this area were seen to be highly employable, to the extent that recruitment and retention of high quality staff was a serious problem (although this problem is not confined to the UK).
4. The major weakness identified by the panel was the lack of exploitation of potentially exploitable research. This was influenced by an inadequate infrastructure to support exploitation and exacerbated by the Research Assessment Exercise, which encouraged the production of academic publications over commercial outputs.
5. Although working from an incomplete data set, the indications are that the IMV programme had only limited success in facilitating the transfer of academic research results into industry.
6. Many basic but extremely difficult problems remain unsolved in Vision research and the UK has the capacity to play a major role in discovering the solutions.

Recommendations

1. The assessment criteria for research proposals should be more closely aligned to the criteria used in the evaluation of completed grants.
2. The EPSRC should investigate what steps it can take to improve the capabilities of UK machine vision researchers to exploit their results, for example, in conjunction with the DTI (e.g. improved involvement of venture capitalists).
3. The EPSRC should make attendance of Theme Days a grant condition.
4. The nucleation of Virtual Reality, Graphics, Video and Vision is predicted to be an important area of research and the EPSRC should encourage work in this area. Despite UK expertise in these areas, currently there is little evidence that the UK is active at this convergence.
5. The EPSRC should encourage medical evaluation to be an integral part of medical imaging proposals.
6. The EPSRC should consider more large initiatives, similar to the recently announced IRCs, to encourage researchers to tackle important, large-scale systems-level problems.
7. The EPSRC should continue to fund grants investigating the complete spectrum from blue skies research through to more applied research in machine vision.

A full report from the Theme Day will be appearing on the EPSRC web site, and hard copies will be available on request. I would be delighted to hear any comments/suggestions you may have on the full report.

Dr Jim Fleming
 APM, Software Technologies
 IT & CS Programme
 EPSRC
 Telephone: (01793) 444 428
 Local Fax: (01793) 444 470
 email: j.fleming@epsrc.ac.uk

BMVA Technical Meeting – Augmented Reality

Date: 3rd May 2000
Location: British Institute of Radiology
Chairpersons: Dr R Bowden (Brunel Uni.)
 Dr A Fitzgibbon (Oxford Uni.)

A joint BMVA and UK Virtual Reality Special Interest Group technical meeting on Augmented Reality (AR) was held in London on 3rd May. It brought together a number of researchers working on a wide range of AR applications. The program, consisting of seven short presentations, covered technical issues of registration, tracking, illumination, and view synthesis in addition to practical industrial and medical applications.

Giles Simon began the day by talking about registration of real and virtual views using points and curve features. His application was in testing the effects of different illumination on a bridge in Paris. The image sequences of the bridge were therefore taken with low light levels and low contrast in order that the virtual light would be dominant. The low contrast made registration difficult, but the detection of free form curves made 3D feature matching possible for registration. Giles' other work involved moving objects in a scene and the modeling of a zoom lens during the image sequences.

Miles Hansard presented a simple procedure for synthesising novel views from two or more basis-images. The method developed is based on the linear relations between images taken by an affine camera with unknown parameters. Miles showed that the method still functions when there is a large difference in orientation between the basis views. It was suggested that this technique would be appropriate for image generation in some AR systems and also for animation.

To finish the morning session I presented work on real time augmented reality for virtual product prototyping. Methods were presented for target based registration and overlay of virtual car interiors for efficient prototyping in the automotive industry.

Philip Edwards showed some gory images from the operating theatre at Guys Hospital to settle us down after lunch. These were during his talk on Microscope-Assisted Guided Interventions – Augmented Reality in the Operating Theatre. The talk encompassed the technical problems involved in the critically accurate registration of pre-operative images during surgery and the visualisation techniques employed to give 3D perception of virtual structures beneath a physical surface.

David Johnston talked about the development of a target system for video positioning. He followed the development of their system from barcodes, circular and spiral targets to a system using region adjacency graphs. The final targets were easy to detect, uniquely labeled, and provided four registration points each. David said that a library for generating and using the targets will be made publicly available later this year for comparative testing.

Shaun Lawson presented an application of AR in remote inspection of hazardous environments. The video returned from a stereo head on a tele-operated robotic vehicle is overlaid with virtual versions of real objects and a virtual cursor that can be used to mark a path for semi-autonomous navigation. The talk focussed on the applications and calibration procedures required to align the overlay.

To complete the day, Simon Gibson presented work on realistic rendering of synthetic objects into real scenes by modelling illumination and shadows. Sphere mapping techniques are used to approximate the illumination onto the synthetic objects, and shadows are built using a number of representative shadow-sources for the scene. The effect of adding shadows was to greatly improve the illusion of a synthetic object sitting on a surface or next to a real object.

In conclusion, the day was well supported by high attendance and a variety of interesting presentations. It was encouraging to see the number of attendees working in such a new research area as Augmented Reality. I would like to thank all the speakers for an interesting day. More information on the talks can be found in the abstracts, made available on the BMVA website at: <http://www.bmva.ac.uk/meetings/meetings/00/3May00/broch.pdf>

Martin Lewin
 Dept. Systems Engineering
 Brunel University
 email: martin.lewin@brunel.ac.uk

BMVA Technical Meeting – Visual Surveillance

The one-day BMVA technical meeting on visual surveillance was held at the usual Portland place in London on 8th March 2000. The special international guest was Professor Carlo Regazzoni from the

University of Genoa, Italy. The venue was organised and directed by Graeme Jones, Kingston University. The meeting was meant to show different flavours of the latest research in visual surveillance. Industrialists could be spotted among the spectators, while Geoff Thiel, from Primary Image Ltd, was the industrialist invited speaker.

From a technical standpoint the meeting had a fair success presenting research results of medium-high level and discussing standard surveillance problems and solutions, which varied from classical to innovative approaches. Professor Regazzoni took the audience through the various generations of visual surveillance systems, stressing the importance of data communication and the constraints imposed by a real-time system designed for visual surveillance applications. Video footage generated by multi-camera systems can be huge, and optimal methods of encoding information and communication architectures are needed to be able to cope with the bandwidth and real-time issues. These issues become crucial when data is acquired with cameras used in day-to-day surveillance, and Geoff Thiel showed us how noisy the video can be, due to weather and lighting conditions.

Multi-camera systems seemed to me to be one of the more interesting and challenging problems raised during the meeting. Tim Ellis was perhaps the main star with his usual very enthusiastic attitude (I recall avoiding him in a Boston restaurant, after a very tiring day of work at a conference, not to be trapped in long technical discussions, I hope he will forgive me for that!). He clearly stated that one of the major troubles with a multi-camera system is that not all camera fields of view overlap, and therefore clever calibration schemes must be devised. He also mentioned a word I love very much: *learning*. Can we devise algorithms able to learn the calibration of a multi-camera system? Do we really need camera calibration?

Another major contribution was given by Jamie Sherrah, who really stunned the audience, well at least me, with extremely well packaged video clips. Jamie presented the work of the vision group led by Shaogang Gong. Their work is innovative and it clearly shows a successful Vision group, which has resources and formidable researchers and students working very hard behind the scenes. Jamie showed what I believe is another cutting edge problem in visual surveillance, i.e. the behavioural analysis of tracked people (or objects in general). If researchers wish to design and develop working, and as automatic as possible, visual surveillance systems, firstly

they have to rely on a robust tracker, which can be used to analyse the dynamics of a scene imaged by one or more cameras. Dr Gong has been working on these problems for a long while, as far as I can remember, and the tracker developed at Queen Mary by Dr Gong and associates is definitely one of the best around. Their more recent work on behaviour analysis is very “trendy” making use of probabilistic models, such as Markov chains, and recent probabilistic filters, such as the CONDENSATION algorithm, developed by Isard and Blake at Oxford labs.

All this was excellent, but we cannot forget that James Orwell presented a good tracker, which perhaps needs some more work but shows good potential, while Steve Maybank stressed once more his love for Information Theory. Steve is a brilliant mathematician, but I have to admit that sometimes he is too much ahead of all of us. What Steve showed appears interesting from a mathematical point of view, but the results he showed were not very encouraging (once he told me he has a very thick skin, and I am sure he will forgive me if I have been blunt!).

Sergio Velastin closed the meeting presenting work on crowd analysis. The presented work was a bit dated, I admit having seen it before, however I am sure that with his new two European grants he will soon have a large amount of new results to show to all of us working in visual surveillance.

Dr Paolo Remagnino
Digital Imaging Research Centre
Kingston University
email: p.remagnino@kingston.ac.uk

ECCV 2000

Why do Computer Vision researchers want to be mathematicians but not statisticians?

I attended the ECCV this year for the first time in six years. This and ICCV are not conferences that I generally attend, but the offer of a free flight to Dublin to give a tutorial on performance characterisation with Patrick Courtney was quite tempting. In our tutorial (to which only 10 people attended from a total conference registration of 250) we emphasised the need for statistical foundations of both algorithm design and testing. We also tried to summarise various techniques for avoiding the problem of obtaining

ground truth. The response was quite positive (we handed out evaluation forms so we know!), and in the discussion following the presentations I was asked the question which is the title of this report. I have to say that at the time I didn't understand it or why it had been asked. It was my gradual understanding of this issue, over the course of the week, that motivated me to write this article, which isn't intended to be a conference review in the conventional sense.

The following day the conference started and Patrick made a rapid exit. He said he had other things to do.²

From my point of view the conference got off to a relatively good start, Chris Bishop used the word "Bayes" or "Bayesian" several dozen times in the first thirty minutes. I think a higher profile for statistical methodology is a good thing, though this paper fell slightly short of actually providing evidence for the validity of the approach.³ The first morning was generally good and most of the published versions of the presented papers contained at least some form of reasonable evaluation, such as ROC curves, so I was feeling quite optimistic about what was in store. However, the poster sessions and the remainder of the afternoon made me weary, and by the end of the day I was beginning to corner people and force them to endure demonstrations of the software I had assembled on my laptop for the tutorial, particularly if it had similarities with the work they were presenting. Poster presentations were a particularly easy target.⁴

Having not learned from my mistake I pressed on the following morning and had demonstrated software to several people before lunch, as well as attending the morning sessions. By lunchtime there were hopes that things would be looking up again, with a long overdue presentation by Michal Irani which made a useful step towards demonstrating the need to use knowledge of measurement errors in closed form solutions to linear problems. However, by the afternoon my optimism had again faded under a barrage of geometry culminating (for me) in the surprising conclusion by Richard Hartley that we should all use trifocal sensors for calibrating systems of images.⁵ This claim perhaps explains the rather evident but curious

²Patrick Courtney would like to distance himself from any suggestion that this was an intentional slur.

³Never-the-less, the paper went on to win a prize.

⁴I apologise to any of those now reading this article who endured this obviously unenjoyable experience, and any unsuspecting victims in the future.

⁵I was later informed that it had been demonstrated with the aid of empirical data sets by Bill Triggs at the previous conference that this is not the case.

fascination with the geometry of multi-camera systems. However, this overstatement was more than made up for by an excellent piece of methodological evaluation work by the same author the following day, applied to a fingerprint matching system.

At the conference dinner Olivier Faugeras was presented with a classic first edition in geometry for 10 years of contribution to ECCV,⁶ only for him to say that he hoped there would be less geometry at the next conference. Unfortunately, I do not know how many of those present would have agreed with him. Clearly the conference reviewers had a completely different opinion, evidenced by the way that Richard Hartley's evaluation work, as with all of the good attempts at developing evaluation methodology at this conference (including Leclerc's paper on evaluating stereo algorithms and Bernard Buxton's evaluation of segmentation algorithms in a plant husbandry system), were consigned to the poster sessions. Others present commented on the randomness of the reviewing but the trend seemed far too systematic to me.

More general highlights of the presentations included those by Matas, Kutulacos, Hancock and Fitzgibbon. The last of which demonstrated that information from multiple moving objects can be used to eliminate the strong correlation between focus and translation in a zoom camera system. This paper also produced, what was for me, the quote of the conference when referring to the use of a copy of Numerical Recipes as a moving target "this isn't the first time that NR has come to the aid of a desperate vision researcher". I think that comedy should definitely be introduced as a reviewing criteria for vision conferences.⁷

The final day was heavy on rather specific forms of camera calibration but was probably more noticeable for its relative paucity of medical image analysis topics, which must surely by now be the largest application area of computer vision in Europe, and for which there are generally quite well specified performance criteria.⁸

So overall conclusion? I now understand the question asked after the tutorial, though I don't know the answer. The workers in this field certainly have the mathematical skills to understand statistical design methodology and some are applying these techniques already. However, I might suggest that the strong use of statistics and evaluation methodologies⁹ might

⁶Sadly I did not make an effort to remember the title and author.

⁷Though not if it always means Fitzgibbon getting a prize.

⁸Is this just a nasty coincidence?

⁹And I don't just mean being a strong Bayesian.

develop a value system for the research work which would be at odds with the interests of the reviewers in this community. I hope that this is not the case, because, as anyone who thinks about these issues will tell you, the subject will not advance rapidly without an appreciation for the need to be able to compare different theoretical alternatives. The rather fundamental questions raised at this conference by Perona regarding the requirements of a learning vision system will certainly not be answered quickly any other way. I also hope that this is not the reason for the apparently low registration level, despite this years beautiful location (Dublin)¹⁰ and excellent local organisation by David Vernon.

I have to finish writing here as England look like they may finally beat the West Indies at cricket. I have a particular interest as I have a ticket to see them when they come to Manchester.¹¹

Notes from our evaluation tutorial will shortly be available from the BMVA web pages.

Neil Thacker
ISBE
University of Manchester
email: nat@smb.man.ac.uk

Future BMVA Technical Meetings

The currently scheduled meetings for 2000/2001 are as follows:

11–14 Sept 2000 *Eleventh British Machine Vision Conference*, Majid Mirmehdi, University of Bristol, see <http://www.cs.bris.ac.uk/Events/BMVC2000>.

1 Nov 2000 *3D Surface and Volume Texture*, Maria Petrou, Surrey University.

24 Jan 2001 *Understanding Human Gestures and Behaviour*, Jamie Sherrah, QMW.

March 2001 *Medical Imaging and Astronomy*, Lewis Griffin, KCL.

May 2001 *Adaway's unsolved problems in Computer Vision*.

As usual we are constantly on the look out for new meetings. So if you feel there is a subject you would

¹⁰Thanks go to Bernard Buxton, John Barron and Paul Whelan for taking every opportunity to make me drink more of various Irish tipples than was good for me.

¹¹I may have to sell it if the West Indies aren't playing well enough to beat England.

like to chair a meeting on or feel there is a subject which we have recently neglected please contact richard.bowden@brunel.ac.uk with your suggestions.

Richard Bowden
Vision and VR Group
Dept. Systems Engineering
Brunel University

Image Processing and Understanding Course

DERA

27th - 29th September 2000,
DERA Malvern, Worcestershire, UK

A comprehensive seminar on the key technical aspects of image processing and image understanding techniques

Learn how to decide what processes and techniques to use to tackle important applications and problems and how to use and apply the technology

Up to date course including presentations by key experts from top research groups:

Prof. J. M. Brady - University of Oxford
Dr R. Cipolla - University of Cambridge
Dr M. J. Varga - Defence Evaluation & Research Agency

Three days of practical learning and in-depth teaching through formal lectures will be reinforced by examples of existing systems and approaches.

To see the latest details of the course programme, visit the DERA web site at: <http://www.dera.gov.uk> (click on News, then Events)

Course Fees: *Academic - £875 plus VAT*
 Industrial - £1100 plus VAT

For a copy of the terms and conditions, or to make a booking please contact:

Sarah Heeks, E707, DERA, St. Andrews Road, Malvern, Worcs, UK, WR14 3PS.
Tel: +44(0)1684 896925, Fax: +44(0)1684 896419, email: heels@signal.dera.gov.uk

Motorola Research Lab Vacancies

The Motorola UK Research Lab is currently expanding, and there are 8 posts available this year. It conducts leading-edge research into the technologies and applications that will underpin the mobile information revolution. As well as developing ground-breaking technologies, we participate in setting European and global standards, maintain strong relationships with universities, and collaborate in

European funded programs. Our talented staff members are drawn from across the world, and we are always looking for talented people to join our team working in Multimedia Applications, Speech Technologies and Software and Systems Engineering.

If you want to know more about working for Motorola Labs in the UK, or worldwide, please contact:

Andrew Aftelak
Motorola UK Research Lab
tel: +44 1256 484496
fax: +44 1256 471383
Andrew.Aftelak@motorola.com

University of Oxford – Research Posts

The Department of Engineering Science seeks to appoint a postdoctoral Research Assistant and a Postgraduate Research Student, to work within the Robotics Research Group (<http://www.robots.ox.ac.uk/overview.html>) on a project funded by the European Union Framework V programme. The Research Assistant's post will be for 2.5 years, while the Studentship is for 3 years.

The posts are to study the problem of creating novel, 'virtual' views of dynamic scenes which have been captured by multiple (typically fixed) cameras, with a target application of creating virtual replays for football matches.

For both posts, further particulars may be obtained from Mr C.J. Scotcher, The Senior Administrator, University of Oxford, Department of Engineering Science, Parks Road, Oxford, OX1 3PJ; to whom written applications should be made, enclosing a curriculum vitae and the names and addresses of two referees.

Please quote 'RA/IDR/DF/00/58' in all correspondence.

The closing date for applications is 18th August, 2000.

The University is an Equal Opportunity Employer.