

BMVA News

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

Volume 6 Number 3
February 1996

Editor: Dr Paul Rosin
Department of Computer Science & Information Systems
Brunel University
Middlesex UB8 3PH
Tel: (01895) 274000 ext. 3632
Fax: (01895) 251686
Email: Paul.Rosin@brunel.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 March 1996.

Contents

Editorial	1
Robust Vision Meeting	2
IEE/BMVA Colloquium	3
Algorithm Benchmarking Workshop	4
2 nd Asian Conference on Computer Vision	5
BMVA Annual Open Meeting	6
BMVC96 Reminder	7
Whither BMVC97?	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.

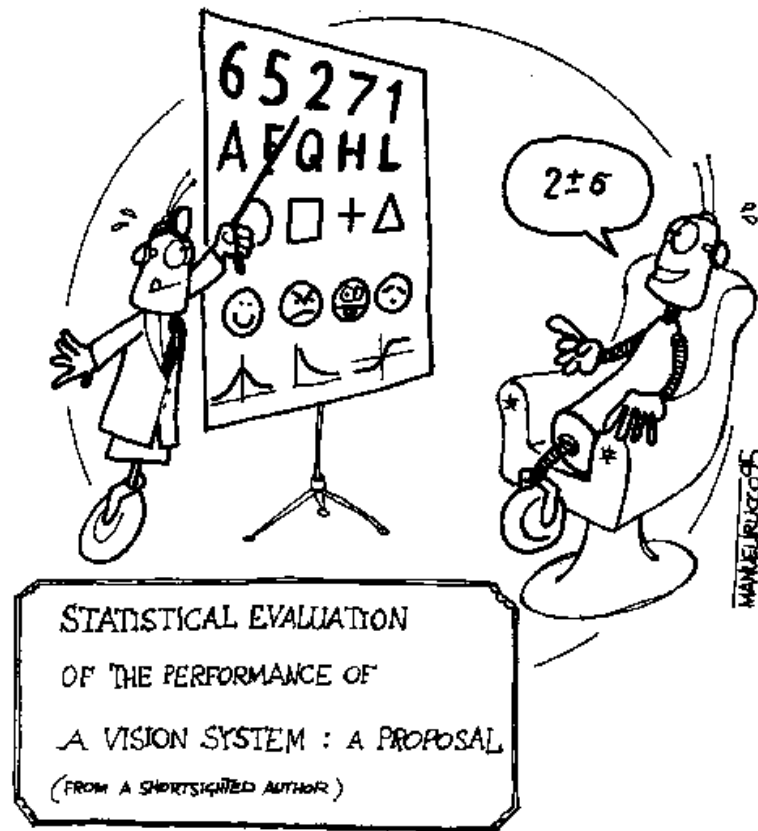
Editorial

As incoming editor I would like to thank the outgoing editor Dr. Phil McLauchlan; in particular for his offer to donate £100 out of *my* pocket to the author of the best letter submitted to BMVA News. Actually, that was an “in” joke amongst editors since the first action of a new editor is to ask the readers to send in any contributions, while the perennial complaint that inevitably follows shortly afterwards is that no contributions have been received. So my £100 is safe!

Actually, there is no shortfall of material for this issue, mainly due to all the reports of meetings, workshops, and conferences that have taken place over the last few months. Anybody actually attending all these events would be kept extremely busy, which of course is a good sign, showing the high level of activity that Machine Vision currently enjoys.

Finally, although I know that any request I make will be swallowed up in the void, I will make one just for good form’s sake. I would like people to send in their favourite quotations from machine vision papers that they have found particularly obscure, incomprehensible, incorrect, or just plain daft. I know that some authors have managed to elevate these styles of paper writing to a virtual artform – and I admit I’m not totally free from blame myself! I will start off with one of my favourites used by James Trevelyan in his paper “Measuring sheep with shaped snakes” presented at the 1992 International Conference on Automation, Robotics, and Computer Vision, pp. CV15.4.1–CV15.4.5:

All sheep have four legs and a head, they are covered with fluffy wool . . .



Robust Vision Meeting

The meeting between the BMVA and Royal Statistical Society on the 6th Dec was perhaps long overdue. In recent years the desire to discard heuristics in favour of sound statistics has penetrated the computer vision community, and it seems only natural to create more links with statisticians. The meeting provided a forum to alert the statisticians to practical problems and to learn from them the recent advances that had been achieved. In both these it was notably successful with each talk raising issues of general interest. Furthermore it was well attended by the academic top brass with three professors giving talks.

Patrick Courtney of ITMI Grenoble was first to speak. He raised the issue of algorithm evaluation. As computer vision systems reach maturity and are used to tackle increasingly complicated tasks specification of their performance becomes essential. How are we to sell a system to a user if we cannot say anything about its reliability? Tests on simulated data, although essential can only tell us about the behaviour of an algorithm when confronted with pre-determined noise distributions. As these are in practice unknown and usually intractable the algorithm

must be tested on real data – but how can the results be evaluated? So far most people evaluate by eye (“well... that line appears to fit quite well”), some evaluate against the known ground truth of, typically, a single example. When the ground truth is unknown then there are so far no sound methods for evaluation, and this is an area requiring much more research. Next I spoke about the problem of robustly determining the order of the model in parametric fitting. This problem becomes very complicated when outliers are considered, which render degenerate data (data which fit multiple solutions) non-degenerate. This is a key problem as many computer vision algorithms are based on fitting or estimation procedures. Although it seems almost compulsory to include the word “robust” in a conference paper title, few are genuinely robust to outliers, and fewer still are stable in the face of degenerate data, where there are radically different solutions depending on which points are considered as outlying or inlying. Some attempt at a solution was suggested in his talk.

The last talk of the morning was given by Philip McLauchlan (Doc MacGlock) of Oxford. He launched a bold attack on the indiscriminate use of the Kalman filter in computer vision, considering its use often inappropriate, especially for structure

from motion. His argument centres around the fact that the dynamics of the Kalman filter are often not needed and demonstrated an “infinite dynamics” filter with better performance. An interesting side issue was raised in the subsequent discussion: At the moment there is no computationally efficient robust filter that has a high break point (is not thrown out by a couple of outliers) and works at any sort of speed. The need for such a filter is immense as more vision algorithms feature filtering through time.

The afternoon began with a talk by Professor Adrian Smith (Imperial) on the merits of a hierarchical approach to deformable template fitting. Peter Rockett (Sheffield) then spoke about the use of the Bhattacharyya metric for comparison of models. This metric, developed in 1947 but not fully understood until recently, provides a useful measure to decide the order of the model that should be fit to a set of data, as the variance between the models is normalised to be the same. Professor John Kent (Leeds) presented a pleasing summary of some statistical approaches that have proved successful in a selection of computer vision algorithms. Professor Brian Ripley talked about work from his forthcoming book (for those who didn't see the leaflet it is *Pattern Recognition and Neural Nets*, Cambridge Univ. Press) on the dangers of misclassification. Also on the definition of what was an outlier, he described how in many computer vision algorithms outliers might be structured and represent clutter, but by correct modelling of such clutter these are no longer outliers but part of our model. Despite this there was a high degree of mutual understanding between the two communities with little difference in terminology.

Overall the meeting proved a great success, well attended and well chaired by Neil Thacker (Sheffield) and Jim Kay (Glasgow). There was a call for another joint meeting to be held with the Royal Statistical Society next year, where perhaps some of the statistical problems to be discussed could be sent out to all concerned beforehand, and members selected to give tutorials on pertinent issues.

Phil Torr
Oxford University
email: phst@robots.ox.ac.uk

Document Image Processing and Multimedia Environments

Colloquium organised by Professional Group

E4 (Image processing and vision) and co-sponsored by the British Machine Vision Association held at Savoy Place, London on Thursday 2nd November 1995; co-chairmen: Dr Ransford B. Johnson, University of Bristol and Dr T. Tan, University of Reading

Chairman's Introduction and Welcome

The chairman opened the meeting and introduced the co-chairman, Dr T. Tan, who attended on behalf of Professor Keith Baker, the BMVA representative. Professor Baker could not attend because of pressing duties as a Dean at the University of Reading.

The rest of the chairman's introductory remarks was spent highlighting the need for discussing the issues relating to document image processing (DIP) and multimedia environments. In particular, the problems posed by forms, filled by handwritten text, music scores, and technical diagrams need to be addressed, especially in relation to multimedia (compound) documents.

First Session

The first paper was presented by the chairman on the topic 'Image Processing for Compound Documents', in which topics such as document analysis, character recognition, analysis of technical drawings, standard document representations and a DIP application – transformation of scanned schematics to CAD format, were discussed. There was some digression from the written text.

In total four papers were presented in this opening session. Due to the quality of the papers presented and the amount of information conveyed and the stimulating discussion that followed each presentation, the scheduled duration was not strictly adhered to and the session ended some minutes late.

Second Session

This session was chaired by Dr T. Tan. Unfortunately, one of the four presenters was unable to attend. However, the time slot provided sufficient time for the meeting to get back on schedule. Up to and including the first paper in this session, presentations were made from universities. The remaining two papers were industrial presentations.

Poster Session

After a lunch (provided for the speakers courtesy of the IEE) was the poster session, including eight posters. This gave the presenters who did not make oral presentation the opportunity to discuss their work amicably with the rest of the audience.

Fourth Session

This final session comprised of three industrial presentations. It was chaired by Dr Ransford Johnson. The final paper in particular, dealt with standards in relation to document image processing and multimedia environments. This presentation was very detailed and the procedure and need for obtaining an internationally recognised standard was presented.

Discussion and Close

This final part of the meeting commenced with matters arising from the meeting. After discussing a few points, the chairman then raised the point of investigating the possibility of formulating a DIP standard that could be submitted for consideration to the International Organisation for Standardisation (ISO). This could take the form of an initial study of current research, standards and products relating to DIP and multimedia. A questionnaire was circulated for possible contributors and interested persons. The study would also complement work previously carried out by the chairman on Multimedia Expertise and Projects Survey. The presenter of the final paper offered to contribute to the procedure of formulating the ISO standard since he is well experienced in the processes involved.

The chairman finally thanked all those who attended and expressed the hope that he could rely on the audience and presenters for future contributions when other related colloquia are staged.

On behalf of the IEE Professional Group E4 (Image processing and vision) Professor Andy C. Downton of Essex University also thanked the audience for their attendance and Dr Ransford Johnson for suggesting the topic of the colloquium and chairing the meeting.

The meeting closed at approximately 4.45pm.

Ransford Johnson
University of Bristol
email: R.B.Johnson@bristol.ac.uk

Algorithm Benchmarking Workshop

On Monday 10th July 1995 an ECVnet workshop was held in Paris on the subject of Algorithm Benchmarking. Contrary to what the title of this meeting might imply those in attendance were not there to discuss the ways in which vision algorithms could be timed on different processing platforms. The main aim of the meeting was to discuss a more deeply rooted problem in the very ethos of algorithm research which can be summarised as follows.

There has been much good work in the past few decades in the development of algorithms for the extraction of various types of information from images. This work has generally concentrated on the assumptions that must be made regarding the data and the numerical form of possible solutions. Less work has been published on the systematic evaluation of the resulting algorithms. This may in part be due to the fact that a rigorous evaluation is a large amount of extra work and has often not been perceived as publishable in the same way as a novel piece of mathematics or the demonstration of a new application. Many algorithms in the literature are tested on only a very small number of images. It is generally agreed that algorithms need to be tested with much larger statistics if any meaningful measure of performance is to be obtained. This is particularly true for evaluating 'robust' algorithms. However, these tests are rarely performed and in our opinion this is normally due to two reasons, firstly the scale of the testing problem and secondly the difficulty of selection of test images. A meaningful methodology for algorithmic evaluation is needed for at least two reasons; to demonstrate the capabilities of an algorithm to potential users, and to provide a systematic method for evaluating (perhaps incremental) changes to algorithms. Without such a methodology the outlook on continued development of computer vision as a scientific subject can only look bleak. Currently there is little agreement even on what things should be measured for even the most common algorithmic solutions. This makes it difficult to expect such problems to be addressed in the published literature.

This workshop was organised by Patrick Courtney and attended by 11 delegates:

Adrian Clark	University of Essex, UK
Neil Thacker	University of Sheffield, UK
Jan Nielsen	AITEK S.r.l., Italy
Claus B. Madsen	Royal Institute of Technology, Sweden

Henrik I. Christensen	Aalborg University Denmark
Wolfgang Foerstner	University Bonn Germany
Patrick Stelmaszyk	ITMI-Aptor, France
Patrick Courtney	ITMI-Aptor, France
Jim Crowley	LIFIA-IMAG, France
Christophe Guizard	Cemagref, France
Pascal Brand	LIFIA, France

Following a series of presentations a healthy discussion was held in which we tried to reach some basic conclusions regarding the general problems facing people wishing to do systematic algorithm evaluation and how they may be addressed. The main (consensus) conclusions were as follows:

- Benchmarking is POSSIBLE. There is often a large parameter space to be searched but this is feasible as it may be reduced or searched by running tests over long periods of time. The large space suggests that initial trials, to build confidence and trust, should be run over smaller spaces.

- Benchmarking is WORTHWHILE. It can help us to gain knowledge of the algorithms we are studying; develop the field of computer vision; add value to our existing algorithms (since an algorithm which has been subjected to many tests and is still used is to be trusted more than a novel one); and helps us to build complex multi-module systems.

- We have ANSWERS to the classic criticisms against benchmarking.

- Agreement on some terms of VOCABULARY seem to be missing. The terms signal, noise, reliability, robustness seem to be used in different senses by different workers depending on their background. It would be useful to prepare a list of definitions. Some progress has already been made in this area for certain terms by other groups.

- There are several possible evaluation METHODOLOGIES, all rooted firmly in the techniques of probability and statistics.

- There is a need to change to cultural attitude towards benchmarking by:

1. Promoting techniques at PhD research level.
2. Rewarding benchmarking activities. Many conference seek “novel” approaches and developing such novelty was felt to be easier and less useful than developing and verifying existing ones.
3. Using common databases, properly described; involving end-users (associations, medicine,

photogrammetry, defense who have a real need for usable machine vision) and ask them to promote or co-promote their results.

The meeting concluded with a set of proposals for actions:

1. Ask the EC to organise a full peer-reviewed workshop at ECCV96, covering methodology, testing and replication of results.
2. Create a special group on “performance characterisation” (a term preferred to benchmarking) within ECVnet.
3. Create a group of people interested by these methodologies with: each participant to send a list of interested people to ITMI; ITMI to send minutes of the meeting to these people.
4. Start discussion on a glossary of terms.

More information regarding talks at the workshop and future meetings can be obtained from Patrick Courtney: pik@itmi.cgs.fr

Neil Thacker
University of Sheffield
email: N.Thacker@sheffield.ac.uk

2nd Asian Conference on Computer Vision

Concorde Hotel, Singapore.

The 2nd ACCV (really the 3rd if you include the conference held in Korea) was held in Singapore from the 6th. December for three days. Having studied the proceedings for the 1st. ACCV, I was expecting a small low key affair. How wrong could I be! The proceedings came in three volumes which is typical of many of today’s conferences. Roll on CD Roms! The increase in the number of papers accepted reflects the enormous increase in the interest in Computer Vision in the Asian region in recent years. One nice aspect of the proceedings was the inclusion of statistics on the subjects and countries of origin of the papers. The most popular subject area by far was Motion Analysis and Tracking which attracted 44 papers. Next most popular was Two/Three Dimensional Vision with 30 papers. Overall the classic research areas of Computer Vision were well supported. Of note

is the increased interest in multi-media which was reflected by papers on Image and Database Retrieval (23) and the applications of the other areas. Considering the countries, Japan was best represented with 121 papers, followed by Singapore with 43. Australia fared well with 29 papers. Overall 37 countries were represented with 450 papers accepted.

The number of papers and the length of the conference inevitably meant there were organisational problems. Five parallel sessions were required and each paper was only allowed 15 minutes for presentation. However a poster session replaced one of the sessions for some of the time so there was less of a clash. A plenary session at the beginning had talks by Shirai and Anil Jain who both covered the latest issues in computer vision from a Japanese and United States perspective respectively. The hospitality of the host country was excellent with both speakers presented with gifts for their talks.

The use of five parallel sessions meant that each session was designed to hold 100 attendees or so. On the face of it this would have been ok except that the first parallel session of the conference included a session called "Future Directions in Computer Vision". This proved very popular and had to be moved to the main conference room to accommodate an estimated 50% of the attendees. However the organisation coped well with this hiccup and the speakers were soon entertaining us with their research areas and statements of future issues. Three of these papers (by R. Picard, R. Bolle and A. Bobrick) were concerned with indexing and describing image sequences. The other three were on coding (T. S. Huang), deformable medical models (J. Aggarwal) and space applications (C. H. Chien). Looking at the contents of the sessions revealed some problems in scheduling. At times there were a number of sessions that were related held at the same time which meant it was impossible to hear all papers of interest.

The conference was held in the Concorde Hotel and most of the attendees were accommodated here. The organisers supplied ample tea, coffee and snacks but not lunch. Those who didn't want to, or couldn't afford to, eat at the hotel, ventured outside to find small eating houses where a reasonable lunch could be had for about five dollars. In my opinion lunch should be supplied at conferences, possibly as a buffet to allow easy interaction of the attendees. This would not put the cost of the conference up but would lead to greater interaction. Apart from this little gripe the hotel was a good venue for the conference and I liked the luxury of just wandering out of my room to attend the sessions.

The banquet was a more traditional Singapore affair with a lion dance and many small courses. We were horrified to find that alcohol was not part of the \$50 cost of the meal! We made do with a couple of glasses of beer which, being Singapore, was expensive.

On the evening of the first day of the conference a meeting was held with one or two representatives of each countries' vision communities. Australia was well represented by Prof. T. Caelli and Prof. R. Jarvis. Prof. Caelli was appointed as the permanent Australian representative on the Advisory Board. It was pleasing to note that Australia was fully welcomed into the Asian vision community. The most important issue to be decided by the Advisory Board was the location of the 1997 conference. Hong Kong put up the best proposal and was accepted with little fuss. This is a good choice given the importance of 1997 to the Hong Kong community. Profs Caelli and Jarvis voiced an interest in holding the 1999 conference in Australia and this was received and noted.

On the evening of the second day a meeting was held to discuss the setting up of ACVnet. This will be on the lines of ECVnet - the European network. Funding will be supplied to support a web page, mailing lists etc.

When we look at Computer Vision, we see there are a few specialised conferences of note. ICCV is the premier international conference (organised by the PAMI technical committee). There is also ECCV which is specific to the European vision community. ACCV naturally fills the role of being the premier vision conference for the Asian Pacific region and should be regarded as such. Overall ACCV is set to become the most important regional conference for the Asian Pacific region and Australia should and will have an important role in the success of the conference and vision research in the future.

Geoff West
Curtin University
Australia
email: geoff@cs.curtin.edu.au

BMVA Annual Open Meeting

Below is a report from the BMVA Annual Open Meeting held on Tuesday 12th September that was received too late to be included with the other reports that were printed in the last issue of the newsletter.

Meetings Secretary Report

A total of 8 BMVA Technical Meetings have been held on the intervening period since BMVC95. In keeping with the aims of holding collaborative meetings with other Societies and Associations, three of the meetings were co-sponsored. In addition the BMVA has co-sponsored a member of vision-related IEE meetings.

One of the benefits of co-sponsoring meetings is to provide members the opportunity to attend meetings at the associated membership rate. BMVA members can hence attend IEE meetings at the IEE members rate (and IEE numbers attend ours free). In addition, the mutual benefit gained from these joint meetings springs from the wider and more diverse audience that can and do attend the meetings. As an example, the joint meeting held with the Photogrammetric Society drew an audience which was substantially composed of their members (in fact, it would have been nice to see a few more BMVA members at the meeting!)

Attendance at this years meetings was almost exactly the same as for 1994 (225 vs 224). The meetings have continued to be held at the British Institute of Radiology which has again upgraded its facilities by installing a new pair of slide projectors. In addition, as a result of a number of problems with using the projectors, they now provide a projectionist to ensure the smooth operation of the projection equipment.

In the first meeting of its kind for some time, a meeting was held away from the BIR, in the leafy surroundings of the Silsoe Research Institute of Bedfordshire. The event offered a range of technical presentations, as well as an opportunity to visit the laboratories of the Institute anticipated to observe the vision work in progress. It is anticipated that one or two technical meetings per year might be arranged at members institutions.

Finally, the Technical Meetings have now been approved under the IEE CPD (Continuing Professional Development) scheme. Attendees are eligible to claim $2\frac{1}{2}$ points (pdu's) for each full one day meeting attended.

1994/94 Programme and attendance

16th Nov	Adaway's Unsolved Problems	
	<i>Cancelled</i>	
7th Dec	Medical Applications	48
18th Jan	Image Coding Workshop	30
1st March	Photogrammetric Methods	39
31st March	Student Papers	

12th April	Teaching Machine Vision	16
26th April	Analysis for Biologically Based Applications	28
24th May	Representation in Natural and Artificial Vision Systems	34
12th July	Industrial Vision	30

1995/96 Proposed programme

18th Oct	Active Vision Phil McLaughlan	
2nd Nov	Multimedia Keith Baker / IEE	
6th Dec	Robust Statistics and Robust Vision Neil Thacker / RSS	
Jan	Machine Vision and HCI Roberto Cippola	
March	Colour Perception Tim Ellis / Colour Society	
May	High-order Statistics Josef Kittler / RSS	

[Note that the proposed programme is out of date, and a new one will be included as soon as possible - Editor]

BMVC96 Reminder

Seventh British Machine Vision Conference

9-12 September 1996

The University of Edinburgh

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).

Up-to-date information on BMVC96 can be found on the WWW at http://www.dai.ed.ac.uk/staff/personal_pages/rbf/BMVC96/bmvc96.htm while more information and guidelines for the format of papers can be found on the PEIPA at:

FTP: `peipa.essex.ac.uk` in
`ipa/info/conferences/bmvc`
(login: anonymous)
WWW: `http://peipa.essex.ac.uk`
under `bmva/index.html`

Papers should be submitted to the Conference Secretary:

Mrs Judith Gordon
Department of Artificial Intelligence
The University of Edinburgh
5 Forrest Hill
Edinburgh EH1 2QL
Tel: +44 131 650.3094
Fax: +44 131 650.6899
Email: `BMVC96@aifh.ed.ac.uk`

IMPORTANT DATES

Deadline for paper submission	29 April 1996
Notification of acceptance	10 June 1996
Deadline for camera-ready copy	8 July 1996

Whither BMVC97?

English is a funny language! The dropping of a single “aitch” from the title of this article could lead the reader to wrongly infer a wish that the annual conference of the Association should expire, cast off its mortal coil, curl up and die! Nothing can be further from the truth. The BMVCs have firmly established themselves as one of the most lively aspect of the Association’s activities and the real intent of the title is to extend this reputation by establishing where the 1997 BMVC might be held.

The organisation of BMVC96 is already well advanced in the capable hands of Bob Fisher and his colleagues in Edinburgh. But as with the painting of the Forth Bridge, as soon as one job has been completed then the cycle must start off all over again. Thus, in March of this year the Executive Committee will discuss venues and bids for hosting the 1997 event.

In looking for a possible venue the Committee will be mindful of many factors including:

- geographical location and the desirability of evenly spreading the event throughout the UK,
- key personnel and other institutional support for organising the event,

- availability and standards of local arrangements including lecture theatres, exhibition space, catering and accommodation.

The general shape of the conference is well known: a 3-4 day event of single track talks with two international keynote speakers and a couple of poster sessions. Usually a small commercial exhibition and book display is also organised. The conference is preceded by a half day tutorial and includes a (usually memorable) conference banquet. Accommodation is typically in University study bedrooms and there is a general philosophy to provide an inexpensive conference which enables the majority of the UK community to participate and benefit, both technically and socially.

Historically the conference has bounced its way up and down the country in a fairly even-handed way. In chronological order from 1990 it will have visited Oxford (Midlands), Glasgow (Scotland), Leeds (North), Guildford (South), York (North), Birmingham (Midlands) and this year Edinburgh (Scotland). It would be an interesting exercise in spatial reasoning to predict its next venue! (the under-represented South perhaps? or Wales? or Northern Ireland? or the West of England?).

To help resolve the matter, the Committee would be happy to hear from anyone who feels that they would like to undertake the formidable challenge of organising the event. This involves considerable work and commitment but is not without its own satisfactions both in terms of helping the UK research community and as a corollary to establishing a site’s claim to be a major UK centre of vision research. So, if you and your colleagues feel that you would like to see the workings of a BMVC from the inside (some would say: not a pretty sight!) then I would like to hear from you. Drop me an email or give me a call and you never know what wonders the future holds ...

John Illingworth, BMVA Chairman
Vision, Speech and Signal Processing
Group
University of Surrey
Guildford
GU2 5XH
Tel: 01483 259835
email: `J.Illingworth@ee.surrey.ac.uk`