

# BMVA News

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

Volume 11 Number 4  
May 2001

**Editor:** Professor Roy Davies  
Department of Physics  
Royal Holloway, University of London  
Egham, Surrey, TW20 0EX  
tel: +44(0)1784 443497  
fax: +44(0)1784 472794  
email: e.r.davies@rhul.ac.uk

<http://www.bmva.ac.uk/>

**BMVA** News<sup>1</sup> 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 30 July 2001.

## Contents

Editorial.....	1
Partnerships for Research and Innovation .....	2
Understanding Visual Behaviour .....	2
Astronomical & Medical Imaging.....	3
BMVC 2003 .....	4
Probabilistic Models in Vision and Signal Processing .....	4
BMVC 2001 .....	6
Book review – Surface Inspection Techniques...	7
Books for review.....	8

<sup>1</sup>The 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

It is interesting to consider what the function of an organ like BMVA News should be. Historically, its role has been largely to inform members what is going on in the world of machine vision and pattern recognition, though naturally, the word ‘British’ in the title perhaps implies that the readers will be primarily of this nationality and will be resident in the UK. It might be useful to question this assumption as it will not be true in many cases, not least if UK members are abroad for a period. However, my main reason for proceeding with this analysis is to consider whether meetings, conferences and courses are or should be the primary interest of this newsletter. In fact, I would suggest that there are other useful roles. One would be to inform readers of new books or relevant tutorial articles. Another would be to inform readers of news about important figures in the world of pattern recognition. Unfortunately, editors such as myself are at the mercy of those who send us articles. If they – or *you* – do not send in articles on these topics they will not be published. So I would put in a plea for news and general snippets of information which others would find interesting. (I could even ask for gossip, though if that became too trivial I would have to exercise my editor’s right of veto!)

To digress for a minute, I feel quite strongly that the RAE<sup>2</sup> scheme seriously undervalues books, including advanced texts and research monographs, both of which will help research to progress nationally and internationally. For this reason I would like to promote the announcement of new books in these columns, and also provide space for substantive book

<sup>2</sup>For those in industry, this is the dreaded Research Assessment Exercise, one of which is at presently being undertaken in UK universities.

reviews. There is perhaps some overlap between advertising and review, but at least review is independent of author and publisher and is hopefully altruistic, though this doesn't mean it isn't subjective. I hope that readers will find the book review in this issue useful and informative. I would like to be able to include more such reviews in coming issues: offers to do the hard work necessary to write reviews will naturally be welcomed!

Roy Davies  
Editor, BMVA News  
email: [e.r.davies@rhul.ac.uk](mailto:e.r.davies@rhul.ac.uk)

## Partnerships for Research and Innovation

Maybe you are an academic expecting industry to fund your research, or maybe you are from industry expecting a university to do free product development for you, or (more likely) maybe you are trying to set up a university/industry link that benefits both sides. Wherever you are on this spectrum a new publication from the CBI will be of use to you. Entitled "Partnerships for Research and Innovation" it gives straightforward practical advice on the mechanisms for enabling collaboration between industry and universities. Despite being a CBI publication, the range of organisations involved in its production is impressive – they range from the CBI and DTI through AURIL to the Research Councils and HEFCE. This means it has a good balance between the needs of the universities and industry.

The guide covers the whole process of forming and managing partnerships. It also provides useful annexes giving an overview of schemes, sources of information, and contact details of the many organisations that can help. Perhaps the limitation of the guide is that much of the advice will seem obvious. Most good advice is obvious and difficult to follow.

As someone who has a passionate belief in the value of partnerships between industry and universities and who works for a research and technology organisation with strong links to both, I strongly recommend that you obtain a copy of the publication and get as many people to read it as possible. It costs just £20 and can be ordered via [//www.cbi.org.uk/](http://www.cbi.org.uk/). The ISBN is 0 85201 553 6.

The guide ends with the following advice: "Don't let the opportunities for research and innovation partnerships slip past just because there may be minor

details and differences of opinion to solve. Instead work at them. The value of partnerships is often greater than you initially expect."

Dr John H Gilby  
The Imaging Faraday Partnerships  
Sira Ltd  
email: [john.gilby@sira.co.uk](mailto:john.gilby@sira.co.uk)

## Meeting Report: Understanding Visual Behaviour

This year's surveillance/tracking/behaviour modelling BMVA technical meeting, on 24 January, promised an exciting line-up of international speakers, and didn't disappoint.

A particularly large turnout saw the show commence with David Hogg's presentation "A Model of Human Interaction" in which several of the Leeds Demos were shown, and the focus of attention moved from the interaction of two people shaking hands, to face tracking and the synthesis of a human face, which interacts with a human, by smiling and acknowledging at appropriate points whilst the human 'speaks' to the computer – we could all use one of these when talking to our mothers!

Hilary Buxton followed with "A Connectionist System for Visually Mediated Interaction" involving a system for face recognition using Radial Basis Function Networks.

INRIA's 'Smart Room', which aims (ultimately) to be fully controlled by the user, is a room which will identify if you have picked up a pen, that you want what you say to be transcribed onto the computer screen, or choosing the eraser, that you want to delete words from your file, etc. Jim Crowley presented "A Probabilistic Sensor for the Perception and Recognition of Activities" which introduced a new use for Gabor Filters, by factoring out the texture of objects, and was about to incorporate into the model some localised knowledge/learning, since in some areas of the room there were activities that were more likely to be performed than others.

The last session of the morning saw "A New Model for Visual Attention" presented by Fred Stentiford of BT, whose work was aimed at image compression with a view to video-style images on Third Generation (G3) mobile phones. They have developed a

novel Visual Attention Estimator, which identifies the keys areas of interest. After some segmentation and illusion examples, we saw the real benefit of such a scheme – region of interest coding – heavily compressing ‘unimportant’ parts of the image, whilst keeping the full detail in the regions of interest.

New to BMVA Technical Meetings was a poster session before lunch. Several exciting ideas were presented, ranging from tracking to automatic gait recognition.

The afternoon session began with Jens Rittscher presenting “Simultaneous Perception and Classification of Visual Motion”, showing the advantages of using filters and condensation over tracking with contours to classify actions within an activity like basketball.

Shaogang Gong spoke about Belief Networks for “Visual Interaction using Gesture and Behaviour” and moved on from the classic example of a face and two hands, to an activity recognition system for the new Queen Mary’s Tuck Shop, with a particularly amusing video showing the detection of a can of pop being stolen.

Tony Jebara demonstrated his wearable computer outfit and spoke about discrimination in generative models in “Action Reaction Learning for Predicting Interactive Behaviour”.

The final talk of the day was by Joel Mitchelson using skin-colour and 2D tracking in 3 images to reconstruct in 3D for “Marker-less Motion Capture for Studio Production”.

Shaogong skipped the summary and suggested we all go down the pub. After a couple of beers, I found myself accepting the task of writing this report.

Full abstracts are on the BMVA website at: <http://www.bmva.ac.uk/meetings/meetings/01/24jan01/>

The meeting was recorded and should be available for viewing from the BMVA website although this is the first time this has been tried.

I’d like to finish by thanking everyone who presented and those who organised the whole event. High quality presentations provided an excellent overview to the area of Visual Behaviour, and the meeting was an overall success.

Chris Needham  
PhD Student  
University of Leeds  
email: [chrisn@comp.leeds.ac.uk](mailto:chrisn@comp.leeds.ac.uk)

## Meeting Report: Astronomical & Medical Imaging

On 18 April 2001, a one day BMVA Technical Meeting in association with the IEE/E4 was held on the topic of Astronomical & Medical Imaging, at the Royal Statistical Society, London.

The aim of the meeting was to explore common areas in Medical Imaging and Astronomy. Although, superficially, there does not seem to be much common ground between these disciplines, the day showed that in fact both can learn a lot from each other.

For a start, both disciplines need instruments to acquire data; and, as John Fordham noted, the instrumentation demands of the scientist in many biomedical applications overlap those required in astronomy. For example, inexpensive CCD technology can replace the use of photographic film in biology labs. Another case of instrumentation overlap was described by Chris Dainty who spoke about how adaptive optics methods, from the field of astronomy, can be used in ophthalmology. In astronomy, these methods are used to correct the aberration of wavefronts to produce images that are not blurred by the rapidly changing atmosphere. In ophthalmology, the same methods can be used to non-invasively image the retinal cone mosaic, and perform laser eye surgery more accurately and safely. David Hawkes also mentioned gamma cameras and solid-state X-ray detectors, which are technologies common to both disciplines.

There are also many numerical methods shared by both disciplines, the most well-known of which is the Radon Transform originally developed for astronomy but, as David Hawkes noted, now used in image reconstruction in CT scanning and nuclear medicine. Tom Marsh spoke about how by combining tomographic and Doppler methods (popular in ultrasonography) it was possible to image structures surrounding binary stars (even though from Earth they are just a single point of light!). Simon Arridge pointed out that underlying most of these methods is the fully general Radiative Transport Equation, which is applicable to phenomena as diverse as neutron currents within stars and water diffusion in biological tissue.

Another thing that both fields have to deal with is geometrically complex data, and this can be addressed by developing problem-specific solutions or more general methods. Xavier Pennec described general methods for calculating statistics of geometric entities, e.g. how to calculate the mean and variance of a set of 3D rotations. The two fields also have a

lot of image processing concepts in common. Lance Miller spoke about how to distinguish features from random noise on images. Random noise is particularly well understood in astronomy, and this depth of understanding led to the development of a tool to help in the detection of masses in mammograms. Maria Petrou described methods for analyzing 3D texture, the results of which were visualised as 3D closed surfaces. Applied to brain images, the method could distinguish schizophrenics from normal controls in a statistically significant way – who knows what might be detected in images of galactic density fields!

At the end of the meeting Graeme Watt from PPARC made a surprise announcement of a £41 million cross-council “Basic Technology” funding initiative, which was particularly appropriate at this cross-disciplinary meeting. To quote: “Basic technology . . . involves the creation of fundamentally new capabilities in areas such as nanotechnology, imaging, photonics and sensors. Each of which will form the basis of major new industries of the future. Basic technology very often is multidisciplinary, usually involving more than one Research Council community.”

We may already be seeing the benefits of this meeting. Sebastian Oliver may have found a use for the suite of medical image registration techniques presented by David Hawkes, in a survey project on the NASA mission SIRTF. A second meeting is planned for April next year. Contact the organizers for further information: Medical Imaging, Lewis Griffin (Lewis.Griffin@kcl.ac.uk); Astronomy, Seb Oliver (S.Oliver@sussex.ac.uk).

Luiza Larsen  
Radiological Sciences  
King’s College, London

## BMVC 2003

The BMVA Executive Committee looks for expressions of interest by University groups active in research in the field of Machine Vision and Pattern Recognition, to organise BMVC 2003. The Committee will consider all expressions of interest that arrive by 12 June 2001. A short list of offers to organise the conference will be formed by the committee during its meeting on 13 June. The short-listed groups will then be asked to produce full proposals that will be considered by the committee during its meeting in

October. The full proposals should include a full budget and organisational details. Please send your expressions of interest by email to Professor Maria Petrou (m.petrou@eim.surrey.ac.uk).

Professor Maria Petrou  
University of Surrey  
email: m.petrou@eim.surrey.ac.uk

## Meeting Report: Probabilistic Models in Vision and Signal Processing

On 9 May 2001, this one day joint BMVA/RSS Technical Meeting was held at the Royal Statistical Society, London. The Chairpersons were Richard Bowden (Brunel) and Charles C Taylor (Leeds).

Vision and signal processing necessitate good modelling. This joint BMVA/RSS meeting provided an arena for discussion between a number of researchers devising probabilistic models for their applications. The nine talks varied greatly in the level of abstraction of the problem, but all related to the extraction of information from data consisting of either signals or images.

Gwynfor Jones began by considering a system of two sensors, one visual, with higher resolution, and one infra-red, with better detection characteristics. The application was surveillance; the hot targets of interest contrasted against the cold background but were similar to the background with respect to the visible spectrum. The work aimed to use the infra-red sensor to assist the process of segmenting the visual sensor’s images in order to detect intruders, the segmentation facilitating computational savings and reducing the incidence of false alarms. This was carried out by calibrating the two sensors. A Markov Random Field over the pixels in the visual image was then used with the observations from both sensors to estimate the segmentation.

Next Patrick Wolfe talked about audio feature extraction, the generic task of extracting information from audio signals. At a low level, this might be generation of a time-frequency representation of the data, while at a high level it might be automatic generation of a musical score. An intriguing point was that you could use two related window functions in the FFTs of a sequence of data blocks to be able

to derive an accurate time-frequency representation, though at the expense of phase information. The talk then progressed on to the modelling of the dynamics of the audio signals. TVAR models were used and claimed to be more intuitively modelled in terms of frequencies and amplitudes, rather than the AR coefficients themselves. This results in a non-linearity that necessitated the use of a particle filter. Particle filtering requires a proposal density, which was chosen as the posterior generated by an Unscented Kalman filter. Questions that followed the talk identified a wealth of potential prior information that remained to be incorporated into the approach. The results looked particularly promising given this scope for taking the idea forward.

John Kent began describing his talk on the problem of modelling the evolution of landmark based shapes through time as “tangentially related to vision”. An interesting, if tangentially related, example of this general problem was addressed, namely modelling the aging of animals. Changes in shape are invariant to rotation, position and scale, which might usually be the parameters of interest. The impact was that the modelling of the aging had to be carried out in terms of a shape-space which excluded these other degrees of freedom. The aging was modelled as a deformation of this shape-space over time. This deformation was described in terms of a linear sum of nonlinear functions of the co-ordinates in the shape space and time. The aging of the heads of both rats and young humans could be described using this technique.

The morning ended with a discussion by Arthur Pece of models for clusters and contours in images. For the cluster modelling, the images were abstracted to points in a three-dimensional space comprising image position and intensity. Gaussians then modelled the density of these points in this space. The parameters of the Gaussians were estimated using the EM algorithm. The contour modelling consisted of abstracting edges in the image as edges of three dimensional objects, cars, by using the observation that edges cause discontinuities in the correlations between neighbouring grey levels. These objects can then be tracked using an Extended Kalman Filter. Both models were demonstrated with impressive video sequences.

After lunch, Pakorn Kaewtrakulpong began the afternoon session with a similar topic to the second half of the preceding presentation. This talk differed in that it considered colour. Colour was modelled in terms of points in an RGB space. A reference was made to an observation that the number of colours,

between which one easily discerns the difference, was small - about eleven. The resulting dimensionality reduction was exploited through the observation that the same colour with different illuminations results in points within a cone in the RGB space. The result is an estimate of colour that becomes invariant with respect to shadowing. The additional information relating to the colour of each target was used to assist in the data association process in a Kalman filtering context.

Next, Malcolm Strens gave a talk on the use of grid based filtering methods to track in low SNR image sequences. Separate grids were used for the position, velocity and acceleration spaces. A single target tracker gave good performance until clutter became very significant. In such environments it was found that the tracker could get distracted by persistent clutter near to the target of interest. This problem was circumvented by devising a multi-target tracker that tracked the persistent clutter as well as the target. This tracking was done in such a way that the trackers had to compete for portions of the image. The conclusion was that it was then possible to track targets in the presence of such clutter.

Neil Thacker followed, slaughtering the holy cow of Bayes’ theorem in forthright manner, his intention being to prompt the audience to consider whether the Bayesian hammer is suitable for every nail. He discussed the limitations of Bayesian inference with respect to the influence of the prior on data that is, in some sense, weak or hard to analyse. The talk also challenged the use of a set of models when it may well be that none are appropriate. Some medical problems were then analysed using non-Bayesian techniques. Discussion between members of the audience then followed on exactly what tools to have in ones statistical toolbox.

The penultimate talk by Maurice Ringer looked at how to use a lack of observations to assist in the tracking of objects. The application considered occlusions in motion capture. The idea is that visible objects cannot be occluded and occluded objects cannot be visible; there are constraints resulting from the set of observations received. An additional term results in the equation for the probability of a given configuration of the body given the observations. It was shown possible to devise a model for these occlusions such that a least squares approach could then be used to nudge the estimates of the configuration of the body into solutions that satisfy the constraints. This was shown to enable the tracking of occluded objects by using the constraints imposed by the occlusion process.



The finale was a talk by Luis Molina Tanco on how to hybridise the motion capture and key frame approaches to the problem of generating synthetic video sequences. A library of motion capture data provided a route from one key frame to another. Certain degrees of freedom were removed from the motion capture data and then a map of valid transitions formed between the frames of motion capture data. The results were impressive; the routine automatically inserted a falling sequence between a walking and crawling key frame.

The presentations all used modelling in one form or another to abstract the problem to a domain more amenable to analysis. The removal of degrees of freedom not pertinent to the inference required for the application was a common theme to a number of the talks. Interestingly, a number of people went from the meeting to the local graveyard, where Reverend Thomas Bayes is buried. Unfortunately, despite a highly peaked prior, the posterior was very flat. The data was weak so we went to the pub.

Simon Maskell & Richard Everitt  
 DERA, St Andrews Road,  
 Malvern, Worcs. WR14 3PS  
 email: s.maskell@signal.dera.gov.uk,  
 r.everitt@signal.dera.gov.uk

## BMVC 2001

**B**MVC 2001 will be held at the University of Manchester, from 10 to 13 September. Professor James Duncan from Yale University, and Professor Pietro Perona from California Institute of Technology have both accepted invitations to speak at the conference, and a healthy 120 submissions have been received, so the conference promises to be very interesting. Presentations will take place in the state-of-the-art lecture theatres at the Medical School, with delegates staying in a pleasant hall of residence 15 minute walk (or 5 minute bus ride) away. The conference dinner will be held at the Lowry Centre (see picture), giving everyone a chance to see the exciting regeneration work taking place in Salford's old dockside area, including views of Liebeskind's new Imperial War Museum.

The conference will start in the afternoon of 10 September, with a tutorial on performance evaluation, given by Patrick Courtney and Neil Thacker. Presentations and posters will take place from 11 to 13 September, and 12 September will be Industry Day. A UKIVA meeting will also be held on 12 September.

Dr Tim Cootes  
 University of Manchester  
 email: t.cootes@man.ac.uk

## Book review – Surface Inspection Techniques

M L Smith *Surface Inspection Techniques*  
Professional Engineering Publishing, London, 2001  
ISBN 1 86058 292 3, hardback £79-00

There are now so many books on Machine Vision that one might be forgiven for making the remark “Not another book on Machine Vision”. Fortunately, this volume does not even begin to fall into this trap – it has a clear role covering a specific area of automated visual inspection, that of surface inspection. This is a tricky area, as the play of illumination on a curved surface naturally leads to complex intensity patterns which can take quite a bit of analysis, especially if slight surface depressions are to be located and measured. The problem is exacerbated if the surface is textured or subject to highlights, and methods have to be developed to cope with such problems. In addition, this area is strictly application orientated, so it is insufficient to develop methods which work on idealised data in idealised situations: the methods must work day in, day out on all the data that come up on specific product lines, and there must be a high degree of certainty that the methodology will work in all cases that arise. Use of artificial neural networks to make judgements must be considered carefully in the light of the possibility that they will not have been trained sufficiently to cope in all cases. Often studies of inspection situations are apparently criticisable for not being state-of-the-art, i.e. for not using all the latest techniques and gizmos that have been developed in machine or computer vision. However, as indicated above, the application situations are exacting, and methods that have not been tested out on some hundreds if not thousands of images simply cannot be used in industrial environments. This is the backcloth to the techniques described in this book. It has to be evaluated for what it is able to achieve rather than on the basis of how much artificial or ‘neural’ intelligence or fuzzy logic it incorporates. It must also be evaluated in its own terms against the criteria it lays down for itself. From these points of view the book achieves quite a bit.

The main methodology involved in the work is based on the concept of photometric stereo and is largely covered in the first few chapters. Chapter 4 “Photoclinometry” on photometric stereo shows how the local inclinations of surfaces can be mapped; Chapter 5 “Reconstruction of acquired surface detail” shows how inclinations are converted to surface depth maps by integration; and Chapter 7 “Analysis of surface

defects” analyses the situation more thoughtfully and comes up with a significantly more profound approach to the mapping of surface detail, looking at the problem using an enhanced hierarchical gradient space representation, with useful practical implications. In fact, this chapter is key to the applications covered in this volume.

*Surface Inspection Techniques* (with its longer subtitle *Using the Integration of Innovative Machine Vision and Graphical Modelling Techniques*) is an adaptation of the author’s PhD thesis: it is not a textbook and contains no problems or worked solutions. It covers necessary theory, though on the whole its approach is experimental. Of the eleven chapters, the first five together with Chapter 7 provide the general background and theory, the next few cover the experimental work and three industrial case studies, while the last two provide an overview and conclusion. The three industrial case studies are the highlight of the work and demonstrate the value of the methodology. They cover the inspection of decorative ceramic tiles, alphanumeric character recognition for print punched on gas turbine blades, and wood product inspection, and give convincing demonstrations of the value of the techniques. The fact that the applications are quite disparate adds support to the techniques described here being generic and thus of more than passing interest.

At xx + 198 pp., this is by no means a long book, especially as some 33 pages appear in appendices. In fact there are eleven chapters, and these are accompanied by the ten appendices, a glossary, 192 references plus a 10-page bibliography, and a subject index. An unfortunate feature of the book is that its price is quite high, and this will tend to discourage general readers from buying it to widen their knowledge. However, it is not so expensive that academics and *a fortiori* industrialists working in the area will not be able to obtain a copy. Added to this is the fact that the book (in large page format) is nicely laid out with a good many well prepared diagrams and images.

Overall, I have no hesitation in recommending this book to those who will be working in this area. It is easy to read and systematically moves towards the relevant inspection goals, developing its theory en route, and taking the reader with it using helpful diagrams and example images.

Professor Roy Davies  
Royal Holloway, University of London  
email: e.r.davies@rhul.ac.uk

## Books for review

Would you like to have the opportunity to review one of the following books for BMVA News? If you review a book you can of course keep it for your own use! Books will be sent out on a first come–first served basis. The two books that have recently been received for review are:

H.A. Mallot (translated by J.S. Allen) *Computational Vision: Information Processing in Perception and Visual Behavior*, MIT Press, January 2001, ix + 296 pp., hardback £34-50, ISBN 0-262-13381-4

P.F. Whelan and D. Molloy *Machine Vision Algorithms in Java: Techniques and Implementation*, Springer, 2001, xvii + 284 pp., hardback £35-00, ISBN 1-85233-218-2

Roy Davies  
Editor, BMVA News  
email: [e.r.davies@rhul.ac.uk](mailto:e.r.davies@rhul.ac.uk)

**See the BMVA website for further up-to-date information on BMVA activities:**

<http://www.bmva.ac.uk/>