

BMVA News

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

Volume 19 Number 4
June 2009

Editor: Professor Roy Davies
Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX
Tel: +44(0)1784 443429
Fax: +44(0)1784 472794
email: e.r.davies@rhul.ac.uk

<http://www.bmva.org/>

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 7 September 2009.

Contents

Editorial: <i>How Can I Get Even More Value into BMVA News?</i>	1
Report on Facial Analysis and Animation Symposium ..	2
Personal Chair for Reyer Zwiggelaar at Aberystwyth ...	4
Report on Symposium on Vision-Based Automotive Systems	4
Is Conventional Image Processing Missing Something Important?	5
BMVA Thesis Archive	6
Articulated Human Motion – Call for Participation	7
Travel Bursaries	7
BMVC 2009 – Registration Open!	8
Nominations for BMVA Executive Committee	8

Editorial: *How Can I Get Even More Value into BMVA News?*

When editing BMVA News, my first action is invariably to throw away all the formatting supplied by the author and reduce the article to pure text. This is because authors use a variety of programmes or versions of Word, and the result is chaos! In fact, I also have to stick to a certain ‘house-style’, which itself means eliminating arbitrary means of formatting. However, this doesn’t mean that I’m not interested in how authors would like to lay out or format an article. Ideally, I would receive a pdf showing what the author would like, and I would also get a pure text version – but that would be more than most authors would find time to do and hence more than I can request. Unfortunately, the alternative of just receiving the pdf version means that while extraction of text is easy, one gets a whole lot of end of line characters that have to be eliminated, and with several quite long articles, the conversion becomes tedious. Hence having a formatted Word version of an article is the best compromise. In principle, an even simpler solution would be getting a Latex version, but I had to eliminate that possibility some time ago, as I found that for an organ of this length and with the need to include pictures easily, Word provided a much faster solution. (Many Latex aficionados will tell me that that was a misguided decision, but exact placing of pictures using Latex still seems to be an art form which takes time and effort, and frankly lack of time had become too much of a problem.)

Lately, I’ve been asked a number of times how long BMVA News articles can be. Oddly, I’ve never had to send an article back for redrafting, though I have edited some down – more for reasons that there was too much irrelevant verbiage that readers would not want to wade

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.

through. However, I thought it worth looking over some earlier issues to find how long articles had been. In fact, I've been the worst culprit, producing articles in excess of 3 pages, containing galleries of pictures from various conferences (notably BMVC and ICPR). Apart from this worst case, I soon came upon a curious oddity – that most Contents indicated consecutive page numbers, implying that articles are *never* as long as 2.0 pages, most being less than 1 page in length. As articles of up to 2 pages are almost certainly not in the boring category, I feel justified in (so far) never turning down the offer of an article.

I ought to point out that I retain the right to edit the wording of articles 'lightly', to bring them into the house style and to make the grammar correct, and the punctuation exemplary! I guess there is also some risk that my spelling checker will miss-spell some foreign names, for which I must apologise.

Overall, my biggest problem nowadays is not failing to get enough articles, but not receiving enough pictures to go with them and thus to relieve the monotony of the text. Oftentimes, I insert my own pictures (even holiday snaps on some occasions!) to 'lighten' an issue, but an editor should edit, not try to do all the jobs, so I must make a never-ending plea for more pictures. "Have some more tea", said the March Hare. "But I haven't had any yet", said Alice. "How can I have more when I haven't had any?" Sometimes this (paraphrased) conversation echoes in my mind when asking for more pictures. In any case, I have occasionally had some pictures, and for some issues I have had more, but *even more* would be a much better idea!

Free offer

If you want to write an article of more than 2 pages, add more pictures! (The *text* length in BMVA News is normally kept at less than 2 pages.)

For reference, BMVA News is printed using 10 point Times New Roman: this fact should give you enough information to work out how long your article will be.

This month's conundrum

How is it that an article that is formatted in two columns occupies *less* space than one that is formatted in a single column?

Professor Roy Davies
Editor, BMVA News
email: e.r.davies@rhul.ac.uk

Report on Facial Analysis and Animation Symposium

The BMVA Symposium on Facial Analysis and Animation was held at the Informatics Forum, University of Edinburgh, Scotland, on 10 June 2009. The Informatics Forum was built between 2005 and 2008. The modern spacious design of the building provided a perfect setting for the meeting.

Chairs: Darren Cosker (University of Bath), Gregor Hofer (University of Edinburgh), Michael Berger (University of Edinburgh), Nataliya Nadtoka (University of Surrey).



The four organisers of the BMVA meeting on Facial Analysis and Animation. From left to right: Gregor Hofer, Nataliya Nadtoka, Darren Cosker and Michael Berger.



The coffee breaks and poster/demo session provided much time for discussion.

The conference program was filled with a number of interesting talks and poster presentations on the topics of facial analysis and animation. Many worthy speakers gave talks, both from the UK and abroad. We were delighted to have three excellent keynotes, Tim Weyrich (University College London), Iain Matthews (Disney Research) and Steven Caulkin (Image Metrics). Our workshop attracted a relatively large audience given the fact that it took place outside London. Overall we had 86 attendees.

The first session began with the invited talk by Tim Weyrich (University College London) who in a clear and informative manner presented techniques for acquiring appearance parameters for models of the reflectance properties of human skin. The anatomically inspired skin rendering, based on a detailed model of the physical properties of the skin, produced strikingly realistic results. The next talk by Will Smith (University of York) explored possible improvements to the 3D Morphable Model framework by adding fine surface details as a bump map with the cues inspired by shape-from-shading.

The first session was concluded by an original and refreshing talk by Frank Pollick (University of Glasgow) on the Uncanny Valley effect. He presented the psychological perspective on the realism and likeability of computer animation characters with respect to form and motion. He critically revisited the idea of the uncanny valley and re-examined the plausible psychological causes for its existence.

Iain Matthews (Disney Research), opened the second session of the day with his keynote presentation. He gave an extensive overview of AAM-based tracking techniques. His talk was packed with clear schematic representations of what happens at the different stages of tracking. Iain raised the important issue of extending the AAM-based approach to person-independent non-rigid tracking.

Umar Mohammed (University College London) presented his research findings on the generation of novel facial images given a combination of a local (patch-based non-parametric) model and a global parametric PCA-based model of pixel values. This mosaicing approach produced realistic results and was able to deal with significant pose and illumination changes.

The last talk of the session was presented by Bernard Tiddeman (University of St Andrews). The highlight of his talk was the perceptual evaluation of computer altered images. The computer generated effects of face symmetry and altered sex and skin conditions were used to learn the associated changes in the observed level of facial attractiveness. Another study was performed using image alterations to simulate ageing and rejuvenating effects.

Due to the very rich programme schedule, lunch time was shared with the poster session. The poster session covered a wide range of the topics: sketch-based animation for storyboarding; caricatures for crime suspect recognition; facial detection and tracking; speech synchronized animation; audio-visual emotion recognition; emotion analysis; use of articulatory data

for speech animation; facial animation in non-human computer characters; and animation of facial expressions. In addition, an industry booth was hosted by Dimensional Imaging of Glasgow, Scotland, who provided demonstrations of their passive-stereo 4D shape capture system.

The third session of the day was opened by the invited speaker from industry, Steve Caulkin (Image Metrics, Manchester, UK). He gave an overview of performance-driven animation, and discussed the impressive results achieved in the digital double project "Emily". He raised the issue of achieving consistent animation results, given a large amount of artist work that's typically involved in the final animation product. Image Metrics' performance-driven facial animation technology provides automation of this process, allowing animation to be driven directly from a video sequence.

Guofu Xiang (Robert Gordon University) presented his work on an automated approach to landmark detection. It was followed by an engaging talk by Sajid Farooq (University of Glasgow). He presented the results of his work on fast multi-resolution rendering of dense surface data using a point-based rendering technique. Given the ever-increasing complexity of meshes used for animation, this speed-up is important to achieve real-time rendering and manipulation of 3D data.

The fourth and the last session of the day concentrated on visual speech synthesis, both video-driven, presented by Barry Theobald (University of East Anglia), and 3D model-based, presented by Philip Jackson (University of Surrey). A large body of work on cloning of speaker's articulation, shape and appearance was presented by Pierre Badin (CNRS and University of Grenoble, France).

Hassan Ugail (University of Bradford) finished the meeting with an overview of his research on facial modelling. He presented very promising initial results on the application of partial differential equations for modelling and animating human faces.

Overall, this meeting was very successful and unique. It was one of the few BMVA meetings to take place outside London, and the first to focus on this topic area. Given all the positive feedback we received after the symposium, we hope to repeat the event next year, either in Edinburgh or elsewhere.

Nataliya Nadtoka
University of Surrey
email: n.nadtoka@surrey.ac.uk

Personal Chair for Reyer Zwigelaar at Aberystwyth

Reyer Zwigelaar has been promoted to a personal chair at Aberystwyth University, taking up the post on 1 June. He joined Aberystwyth University as a Senior Lecturer in 2004: before that he had faculty positions at Portsmouth and UEA.



Professor Reyer Zwigelaar

His research has concentrated on the interface between medical image understanding and computer vision, with a strong emphasis on mammography (detection of abnormalities, multi-modality and risk assessment) and prostate cancer (segmentation and staging).

Reyer was co-chair for Medical Image Understanding and Analysis in 2002 and 2007, and for the International Workshop on Digital Mammography in 2006. He is currently Associate Editor of IEEE Transactions on Information Technology in Biomedicine. I am sure all members of the BMVA will join with me in congratulating Reyer on his timely and extremely well-earned promotion!

Professor Roy Davies
Editor, BMVA News
email: e.r.davies@rhul.ac.uk

Report on Symposium on Vision-Based Automotive Systems

The recent BMVA symposium, held on 20 May 2009, saw researchers from the UK and across Europe discuss the use of vision-based systems for automotive applications. The use of vision has a number of important roles to play in the automotive industry, and

with several car manufacturers introducing vision-based systems in their products, it continues to be the focus of active research in a number of areas.

The meeting was well attended, with over thirty people from a range of backgrounds arriving at the BCS venue in Covent Garden, London. In total, twelve speakers were invited to present the latest results from their research activities. Whilst many of the speakers were based in the UK, several came from European institutions in France, Spain and Switzerland. This gave attendees the opportunity to see the latest work from international research centres. Despite a busy schedule of speakers, everyone had the opportunity to discuss their own research activities during the breaks between presentations.

Topics ranged across many areas of interest, including detection of road signs, road boundaries and pedestrians, as well as vehicle detection and tracking. Many of the systems being presented supported methods for improving both road safety and the driver experience, such as robust sign detection and interpretation, collision warning systems for rural or night-time environments and road boundary detection.

Several of the projects used fusion of data from a range of sources, including GPS, in-car data, and road features detected from visual sources. Applications for this approach included conventional hazard detection and avoidance, and in one case, a visual navigation system using augmented reality.

In addition, the use of a computer-based system to monitor driver behaviour was presented, and in a related talk, the results from an investigation into the difficulties associated with autonomous vehicle operation were shown. Overall, the range of presentations, which included both in-car systems and external, out-of-vehicle monitoring, was well received by the attendees.

Toby Breckon, meeting chairman, said "We had an interesting day with a lot of input from institutions both in the UK and from wider Europe. With such a strong response to the call, I hope we can perhaps make this a bi-annual topic in the BMVA technical meetings calendar."

Stuart Barnes
Cranfield University
email: s.e.barnes@cranfield.ac.uk

Is Conventional Image Processing Missing Something Important?

During the 1970s there was a substantial flurry of exploration into both the structure/functioning of the human visual system and improved efficiency of image processing. Although, on the face of it, these two subjects are only loosely related, nevertheless on a number of occasions the two areas of research did overlap substantially. Possibly this became best known in the work of David Marr at MIT and in the publication of his book *Vision* [1]. Here it was demonstrated that the early processes of human vision (in the eyeball) could be loosely related to the growing popularity in image analysis circles of Laplacian operators to provide edge detection by sensing the zero crossings of two-dimensional second differences of imagery. Marr showed very elegantly how, in general terms, the substantial blur known to exist in the retinal image of the human eye could be related to the image blur necessary when attempting to analyse sampled images in order to provide adequate anti-aliasing (leading to the DOG – difference of Gaussian – concept). Later an associated team at MIT continued to study the extension of Marr's concepts to optical flow, in particular.

Very much in parallel with the foregoing developments, but working completely separately, during the late 1960s and 1970s a group led by me at BAe Filton had been studying multi-parametric detection thresholds for human vision. Considerably before Marr's book appeared I had reached similar *general* conclusions about the actual early processing in the human eyeball. However, as far as the trends in edge detection thresholds were concerned, I was finding two rather important things which required explanation:

1. Whereas the evident Laplacian-like operations were apparently associated with *second* differences, it became clear that the *threshold* trends in human vision as functions of edge contrast were strongly related to *first* differences! How could this be?
2. There was substantial evidence that the local edge thresholds for secondary visual functions such as form recognition, motion detection and stereo perception were of a very high resolution – much higher than one would expect from the retinal receptor spacing. Also, at least in some cases, they had *directionality* built in. In other words, they appeared to be associated with *sub-pixel vectors* rather than *nearest pixel scalars*.

My own book [2] attempted to pull together under one cover the substantial and varied visual threshold performance data collected from many different disciplines, together with a summary of the properties of

the eyeball and a historical record of the various models then existing of visual threshold performance. It was also supported by over 300 classified references. The coverage of the subject was updated by a paper published in 1982 [3].

Then, in 1983, John Canny reported a new edge detection operator starting from a mathematical approach [4]. This operator was immediately shown to have substantial attractions over and above previous operators for edge sensing and rapidly became generally accepted. However, it was (and apparently still is) limited to single image analysis and outputs of nearest pixel scalars.

In the meantime, during the late 1970s and early 1980s, I had been required to extend my studies of visual function to endeavour to model the *details* of the early neural networks, so that it might be possible to simulate a number of the various known (observed) functions of human vision. When the Canny operator was first proposed I was able to study the background behind this in some detail and also to have personal discussions with Canny. After the discussions we generally agreed that our two *basic* resultants were essentially similar, but starting at *opposite* extremes – his from the mathematical needs and mine from an attempt to interpret how human vision was structured.

However, whereas Canny's operator appears to remain purely an edge detector yielding nearest pixel scalar descriptions of edges for single input images, my own processes were able to be developed to encompass the deeper workings of human vision (up to and including at least the Hubel and Wiesel neurones in Area 17 of the Striate Cortex [5]). It has been demonstrated that at this level the available output data (all originating from the same simple imagery in the retinae) are capable of providing a myriad of different outputs, all of them as vectors and at substantially sub-pixel level. These include not only a wide variety of outputs from single input images but also either optical flow maps from *temporal* image pairs or stereo disparity maps from *stereo* input pairs. The background to the development of my own processes, together with descriptions of the algorithms and demonstrations of the various practical capabilities, were published as a second substantial book in 1992 [6]. At that time, of course, personal computers were still somewhat of a luxury and digital cameras barely existed at all! In addition, the Internet was still in its infancy. Hence the whole of the research leading to the material published in both my books was carried out under considerable difficulties (relative to the ease of similar research today). Since that time not only have both personal computers and digital cameras become very powerful, but also it has become commonplace for most people to own them.

With the massive advances in technology, it has been possible for me (with the aid of my son, a computer programmer in his own right) to create a composite suite of software incorporating the majority of concepts reported in [6] under the title “Simulated Human Vision” (SHV). Also, with the massive advances in the Internet, it has been possible during the past 3 years to set up a website dedicated to SHV [7]. This website includes a historical background to the development of the software, demonstrations of its many facets (animated where appropriate) and a section providing a hierarchical description of the human visual tract (with illustrations of the local progression of processing of local scene elements). It also now includes electronic copies of both my books (chapter by chapter as PDFs) plus several relevant recent reports on specialist practical applications. Amongst the topics are:

- Automated brightness and colour matching (conventionally still usually carried out by human visual inspection)
- Fingerprint analysis (a preliminary assessment)
- Rapid, semiautomatic whole-frame stereo fusion
- Examples of possible satellite image analysis
- Automated optical flow analysis (under an earlier title of OEA)
- 3D textile drape and body scanning (using simple, home-constructed equipment)
- A variety of comparisons between use of the Canny operator and equivalent SHV processing.

I am aware, from browsing the BMVA website, that there are now very many and varied groups with potential machine vision and image processing needs. The foregoing list of reports covers a varied subset of these, but I feel certain that many, if not all, of the other fields represented could also possibly benefit from exploration of some of these alternative techniques. I am particularly conscious of the new explosion of ‘Smart’ cameras with inbuilt software and I would venture to suggest that some of the tricks used by human vision may well enable such smart cameras to provide enhanced performance more simply (or cheaply).

In the 1980s it took many minutes to carry out any sort of complex image processing beyond simple edge finding, followed by a tedious procedure for interrogating the results of such processing. Now, with the aid of SHV, it is possible to run entire processes even for large, very high resolution input images in a matter of seconds. Also, with the aid of modern graphics programs (supplemented, on occasion, by specialist additional routines provided as a support to the main SHV software), any form of results can be viewed and assessed very quickly. In fact, with

specialist supplementary software, it has even been possible to process 640×480 flow sequences fully at up to 30 frames/second. In order to permit potential users ‘hands on’ experience using what, to most, will be a very new approach to image processing, the SHV website offers (for free download) a limited version of the software suite (under the title ‘SHV_taster’). This limited version, although operating from old fashioned DOS command files, nevertheless allows experimentation with substantial images similar to the best that could be explored and reported in days when [6] was being prepared. The main programs of the *full* SHV software are designed to be much more Windows friendly, running directly from a Windows front end.

For easy hyperlink access to the more important *practical* demonstrations on the website, plus the general literature and the ‘SHV_taster’ software, please go to:

<http://www.simulatedvision.co.uk/page192/>

References

- [1] Marr D. (1982) *Vision*, W.H. Freeman and Co., San Francisco.
- [2] Overington I. (1976) *Vision and Acquisition*, Pentech Press, London.
- [3] Overington I. (1982), “Towards a complete model of photopic visual threshold performance”, *Opt. Eng.*, **21**, 2.
- [4] Canny J.F. (1983), “Finding lines and edges in images”, MIT Tech. Report No. 720, Cambridge, Mass., USA.
- [5] Hubel D.H., Wiesel T.N. and LeVay S. (1977), “Plasticity of ocular dominance columns in monkey striate cortex”, *Phil. Trans. R. Soc. London*, B **278**, 377.
- [6] Overington I. (1992) *Computer Vision; a Unified, Biologically-Inspired Approach*, Elsevier North Holland.
- [7] <http://www.simulatedvision.co.uk/>.

Dr I. Overington
email: ian.overington@btinternet.com

BMVA Thesis Archive

In order to promote and improve access to the large base of high quality PhD research undertaken in Computer Vision in the UK, the British Machine Vision Association last year launched a new online repository. This will act as a single-source archive of all past, current and future PhD work undertaken in this area in

UK academic institutions. The service will allow students to quickly and easily share the results of their work with the Computer Vision community, nationally and internationally, and it will be a tremendously useful database for searching and reviewing previous PhD research work undertaken in the UK.

The real value of this service can only be realised if the UK community support the effort and so the BMVC would like to encourage all members to use and, where possible, contribute material to the repository. Contributions are required to be in PDF format and supplementary material such as videos and images are welcome.

The PhD repository can be accessed through the main BMVA website (www.bmva.org). If you have any problems submitting your thesis to the repository please contact Dr Aphrodite Galata.

Dr Aphrodite Galata
Manchester University
email: agalata@cs.man.ac.uk

Articulated Human Motion – Call for Participation

One-day BMVA technical meeting in London, UK, on 23 October 2009

Chairs: Dimitrios Makris (Kingston University), Aphrodite Galata (University of Manchester)

Research in video-based solutions for estimating and interpreting Articulated Human Motion attracts the interest of a growing number of researchers in the field of Computer Vision, Computer Graphics and Biomechanics. Recent research has focused on the development of markerless non-intrusive systems, in contrast to the Motion Capture systems that require the attachment of markers or sensors on the human body.

A wide range of industrial applications will benefit from research in this area. For example, virtual replaying and analysis of athletes' motion patterns in sports analysis, gait analysis, behaviour understanding and event detection in visual surveillance, body motion-based interface in Human Computer Interaction, biomechanical analysis in medical biomechanical analysis in orthopaedics, and so on.

The high-dimensional space of human poses, the non-linearity of articulated motion and the projection of 3D

human bodies onto 2D image planes are some of the challenges in this domain.

The aim of this meeting is to bring researchers together to present and discuss their work in this interdisciplinary research domain. Submissions are invited in the following areas:

- 2D/3D pose recovery from still images
- 2D/3D articulated tracking
- Dimensionality reduction of human pose space
- Multiple-camera fusion
- Behaviour modelling and recognition
- Gait Analysis
- Human body modelling
- Evaluation datasets and metrics.

Other topics within the broadly applicable domain of modelling, tracking and understanding of articulated human motion will also be considered for inclusion.

Please submit an extended summary about one A4 page in length (no longer than 2 pages), PDF preferred. Send contributions by email attachment (1Mb max please) to Dimitrios Makris (d.makris@kingston.ac.uk) by 18 September 2009.

Dr Dimitrios Makris
Kingston University
email: d.makris@kingston.ac.uk

Travel Bursaries

The BMVA awards a number of travel bursaries each year to help student members present their research at significant international conferences. Students must apply after their paper has been accepted but before they travel, and are required to write a conference report for BMVA News when they return.


Applications should be made by email to bursaries@bmva.org. For further details see the BMVA website:

<http://www.bmva.org/>.

Dr Adrian F Clark
University of Essex
email: alien@essex.ac.uk

BMVC 2009 – Registration Open!

There were 325 submissions to BMVC this year – an increase of 10% on the submissions to last year’s conference in Leeds. Papers were submitted from all around the world: only 24% of them came from the UK, proving once again that BMVC is truly international! This is well illustrated by the graphic below.



More than 100 dedicated reviewers and 25 area chairs considered the papers. In the final area chair meeting on Tuesday 16 June, 30 papers were accepted for oral presentations and 95 for poster presentations. This gives BMVC 2009 an overall acceptance rate of 38%, or 8% for oral presentations.

By the time this issue of BMVA News is published, registration will be open. Final camera ready submissions are due on Tuesday 21 July.

UK Student Workshop

A 1-day student workshop is to be co-located with BMVC on Friday 11 September, immediately following the conference. For submission and other details see the workshop website at:

http://web4.cs.ucl.ac.uk/research/vis/bmvc2009/index.php?option=com_content&view=article&id=57&Itemid=64

Dr Simon J.D. Prince
University College London
email: s.prince@cs.ucl.ac.uk

Nominations for BMVA Executive Committee

Nominations are requested for the forthcoming election of BMVA Executive Committee members.

We are looking for new members with the ideas and time needed to help enhance our research community.

Nominees must be paid-up members of the Association and agree to serve for a period of two years.

A member of the Committee is expected to be in a position to take an active role in BMVA activities, and must participate in the bimonthly committee meetings taking place in a location deemed mutually most convenient to committee members.

Completed nomination forms should be sent to Dr Thacker at the address below and must be received by 31 July 2009. This must be signed by the individual standing and by one other BMVA member. The nomination must include a brief biographical statement, for distribution during the electoral process.

The BMVA committee has ten elected members, five of whom are elected each year. Details can be found at:

<http://www.bmva.ac.uk/admin/exco.html>

The members elected in 2007 who will stand down this year are:

- Dr A. Fitzgibbon,
- Dr P. Hall,
- Dr D. Makris,
- Professor R. Davies,
- Professor M. Mirmehdi.

If more than 5 nominations are received then a postal ballot will be held, with voting papers returned by 1 September.

The results will be announced at BMVC 2009 in London, and in the BMVA News.

Dr N.A. Thacker
BMVA Secretary
email: neil.thacker@manchester.ac.uk