BMVA News

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

Volume 29 Number 2
December 2018

Editor: Professor Roy Davies
Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX
tel: +44(0)1784 435064
email: e.r.davies@rhul.ac.uk

http://www.bmva.org/

BMVA News[1] 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 10 March 2019.

Contents

Editorial: *Why Can't Photography be Widely Used in the Surgery?*

The NHS has long found that winter is the time when its situation is close to breaking point, with long queues at GP surgeries and many cancelled operations in hospitals. All this stems from lack of provision of the numbers of qualified medical practitioners that would be needed to cope with an aging population, problems of winter colds and flu, and latterly, negation of the need for foreign and, dare I say, EU doctors and nurses to complement home-grown ones. On the one hand there is the problem of under-investment, and on the other the resulting problems of excessively long queues in surgeries and hospitals. Of course, even in an ideal world it is impossible to reduce queues to negligible proportions, because of the statistical nature of queuing. (Imagine that doctors might sometimes be queueing for patients!)

Now consider a related case, that of the dentist's surgery. In spite of having to pay for a proportion of the cost, one still has to wait for a similar time. On the other hand, over the past 10–15 years I have noticed significant improvements in the quality of dental treatment. Teeth are actually being photographed and modelled in 3D, and pictures are being taken at various stages of the process. Implants have become much more common, albeit expensive, and are certainly not paid for on the NHS – though one can ameliorate this problem by going to foreign parts for cheaper treatment. (But it appears that one only does this at one's peril, as some UK dentists claim to make a living out of correcting shoddy work done by others.) Anyway, my point is that dentistry has moved solidly into the 21st century, with X-rays being used more extensively (because of the lower exposure times needed by modern sensors); with CT scans of the whole upper or lower jaw being provided by special centres; and with digital photography, aided by various 3D metal frames fixed to one's head to make accurate measurements for implant and crown fabrication. Of these, only the CT scans and the fabrications themselves are not normally carried out on the dentist's premises. But the fact is that full use is made of digital photography and visualisation; incidentally, both of these also help the dentist to explain to the patient why his or her treatment is going to cost quite so much!

Returning to the GP's surgery, I can honestly say I have never even seen a digital camera or a GP using one. Of course, within the confines of the allowed 10-minute appointment, even picking up such a camera is going to consume several minutes. In fact, this hasn't prevented me from taking my own pictures and presenting them to the doctor. Actually, I have found this particularly useful, as I have been able to catch the state of skin patches, sores and

[1] 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.

the like and show the doctor the progression. This sort of thing inhibits the doctor from suggesting a mere palliative – by-passing the well-known syndrome that doctors only take any real notice on the *second* visit, the first visit serving only to calm the patient down!

However, the other function of a GP is to refer one to a specialist, such as a dermatologist. In that case, a further queue of some months can easily be involved, by which time the patch may well have changed in size, shape or colour, or developed into multiple patches, or temporarily disappeared. All the more reason to have a reliable photographic record for the specialist, but the GP certainly isn't providing this, so the dermatologist has to start from scratch. But even this dermatologist is not taking pictures, so the next one has to start all over again, from pages of scribbled notes provided by his predecessor. Fortunately, I have got the message that I need to go along with a full photographic record of my own if I want to get the right treatment in the end. While many of us computer vision specialists are well able to do this, the average patient would not even think of it, or be able to take adequate photographs: remember, the human body is an articulated 3D entity and lighting may not be so straightforward to set up if the photos are to give a sufficiently accurate rendition of the problem. Nevertheless, the medical profession does not seem to have any cognisance of all this – though one would have thought that a handful of GPs and specialists would have cottoned on to the state of *current* technology (after all, I bought my first digital camera in 2003 for about £400) and the need for such precision as an aid to diagnosis – especially for complaints that are time-varying. Indeed, over a period of more than a year I have been variously informed that I have ringworm, urticaria, both, neither, or maybe an unknown autoimmune condition (for the first two of which I've received several different treatments, none of which has worked), but no photos have been taken. For a health service not to cost undue amounts of money, it would seem to be advisable to reliably diagnose complaints quickly rather than let them go on and on, with doctor after doctor re-examining them: photography is evidently a simple sufficiently inexpensive way in which this problem could be helped.

I apologise if I have inadvertently over-used a compilation of my own experiences to illustrate ways in which the health service needs to be improved. That it is practical and ultimately cost-effective to do so is evident from the parallel case of dentistry. Indeed, it is up to all of us vision people to prod away at the NHS and bring it into the 21st century, both with regard to photography and perhaps with better means for automatically taking standardised full colour pictures in the surgery, so that all the GP has to do is to press a button and send the resulting file to a suitable specialist. I can't believe I'm the only basket case who should be concerned about this.

You may have noticed that I haven't gone so far as to recommend providing a fully automated photographically-based diagnostic machine for every GP: while this might in the end be ideal, it could delay the provision of basic photographic recording to encapsulate important visual data, store it, and communicate it reliably to others.

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

The Sullivan Thesis Prize and the BMVA Thesis Archive

Every year, the BMVA awards a prize for the best thesis out of those brought to its attention as having been examined in the previous calendar year. The prize is awarded in the name of Geoff Sullivan, who played a significant role in the early days of the BMVA.

If you are in the final stages of writing up your thesis, please consider submitting it to the BMVA's thesis archive: all the information is on the BMVA website. If you are supervising a PhD student who you think has done particularly well, please consider entering him or her for the Sullivan prize; again, the procedure is on the website.

Dr Adrian Clark
BMVA Chair
email: alien@essex.ac.uk

Travel Bursaries for International Conference Attendance

The current arrangements for BMVA Travel Bursaries will carry forward into 2019. Note that there will be a fixed number of deadlines, as indicated on the bursary link to the BMVA website:

http://www.bmva.org/bursaries

Professor Lourdes Agapito
University College London
email: l.agapito@cs.ucl.ac.uk

BMVC 2019 – Call For Papers

The British Machine Vision Conference (BMVC) is one of the leading international conferences on computer vision and related areas. It is organised by the British Machine Vision Association (BMVA). The 30th BMVC will be held at the University of Cardiff during 9–12 Sept 2019.

Authors are invited to submit full-length high-quality papers in image processing, machine vision and related areas for the 2019 British Machine Vision Conference, bmvc2019.org. Submitted papers will be refereed on their originality, presentation, empirical results, and quality of evaluation.

All papers will be reviewed doubly blind, normally by three members of our international programme committee. Please note that BMVC is a single-track meeting with oral and poster presentations and will include three keynote presentations and two tutorials.

Topics include, but are not limited to:

- Statistics and machine learning for vision
- Stereo, calibration, geometric modelling and processing
- Face and gesture recognition
- Early and biologically inspired vision
- Motion, flow and tracking
- Segmentation and grouping
- Model-based vision
- Image processing techniques and methods
- Texture, shape and colour
- Video analysis
- Document processing and recognition
- Vision for quality assurance, medical diagnosis, etc.
- Vision for visualisation, interaction, and graphics
- Object detection and recognition
- Shape-from-X
- Video analysis for action and event recognition
- Illumination and reflectance
- Deep learning for vision
- 3D computer vision
- RGBD analysis.

Accepted papers will be included in the conference proceedings, published and DOI indexed by BMVA. Past proceedings can be found online here: www.bmva.org/bmvc/. Prospective authors can have a look at the 2018 edition: bmvc2018.org/programmedetail.html, for example.

The paper submission deadline is 29 April 2019 (23:59, Pacific Time). Submission details are available on the BMVA 2019 Web site. Important dates can be found here: bmvc2019.org/dates/.

Note: Because of the anticipated volume of papers for BMVC 2019 – and based on our recent experience – *no extension will be granted to the submission deadline.*

Selected best papers for BMVC 2019 will be invited to a special issue of the International Journal of Computer Vision (IJCV).

Any queries should be sent to the Programme Chairs: programmechairs@bmvc2019.org

The BMVC 2019 Programme Chairs are:

- Professor David Marshall (Cardiff University)
- Professor Majid Mirmehdi (Bristol University)
- Dr Bernie Tiddeman (Aberystwyth University)
- Dr Xianghua Xie (Swansea University)

BMVC 2019 – Call for Workshop Proposals

BMVC is soliciting proposals for workshops to be held to accompany its Conference: see bmvc2019.org. Workshops will take place on the afternoon of 12 September 2019 (i.e., half-day), at Cardiff University at the main conference venue. The purpose of the workshops is to provide a comprehensive forum on topics that would not be fully explored during the main conference, and to encourage in-depth discussion of technical and application issues.

For examples of past workshops see BMVC 2018 workshops, bmvc2018.org/workshops.html: details from all past conferences can be found here: www.bmva.org/bmvc/.

We also welcome 'Challenge Workshops' that aim to compare new and established methods on common data sets.

BMVC 2019 organisers will collect workshop registrations, provide facilities, and distribute electronic copies of the workshop proceedings.

There will be competition for workshop space, time, and topic coverage. To enable the competitive selection process, proposals must be specific and detailed in justifying relevance and viability. Proposers may be asked to provide additional information, modify aspects of their proposals, or combine their proposal with another one.

Also note that publication deadlines are very tight between the main conference acceptance notification (24 June 2019) and the workshop camera-ready deadline (14 July 2019), so proposers have to be ready to undertake all the work related to soliciting and reviewing submissions and collecting final contributions.

Proposals should be submitted by email to workshops@bmvc2019.org by 17 February 2019. Proposals should be in PDF format and include the following information:

- Workshop title.
- Proposers' names, titles, affiliations, and primary contact email.
- Topics that will be covered.

- Background and experience that makes the proposers well suited for organising the workshop.
- Rough program outline (estimated numbers of orals, posters, and invited talks).
- Names and bios of any invited speakers and indication of whether they have agreed to speak.
- Anticipated target audience as well as expected number of attendees.
- Description of relevance and viability.
- Description of how this proposal relates to previous workshops at BMVC/CVPR/ICCV/ECCV (be as specific as possible).
- Any special space or equipment requests.

Please contact the workshop chairs, Dr Padraig Corcoran, Dr Jing Wu and Dr Dan Gallichan if you have any questions.

BMVA Symposium: Deep Learning in 3D

One-day BMVA symposium in London
Wednesday 20 February 2019
Chairs: Chris Holder, Christopher Willcocks, and Gregoire Payen de La Garanderie

bmva.weebly.com

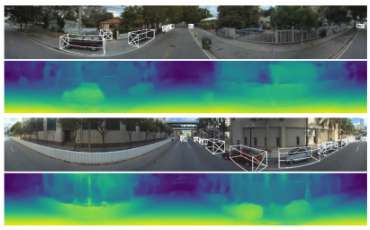

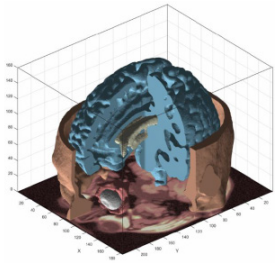

In recent years deep learning has revolutionised the world of artificial intelligence, providing a huge boost to machine learning research as well as to real-world applications of the technology. In this meeting, we aim to explore the key challenges of combining deep learning with 3D vision.

Programme

09:30 Welcome
09:45 Keynote: Deep Learning in 4D, Adrian Hilton
10:45 Coffee
11:00 3D pick and mix: object part blending in joint shape and image manifolds, Adrian Sanchez, Oxford
11.20 Learning to infer the 3D shape of sculptures, Olivia Wiles, University of Oxford
11.40 Volumetric performance capture from minimal camera viewpoints, Andrew Gilbert, University of Surrey
12.00 Clinical evaluation of ML approaches for the classification of 3D gait using models, Adar Pelah, York
12:30 Lunch
13:15 Keynote 2: Alex Kendall
14:15 Spotlights
14:45 Posters and Coffee
- Plug-and-train loss for single view 3D reconstruction, Eduard Ramon Maldonado, UPC
- DL for classification in 3D point clouds of heritage sites, Stamatis Chatzistamatis, University of the Aegean
- Improved object detection with 3D deep neural networks, Justin Le Louëdec, University of Lincoln
- Recovery of superquadric parameters from range images using deep learning, Franc Solina, Ljubljana

15:30 Graph Convolutional Neural Networks for 3D medical images, Michael Edward, Swansea University
15.50 Learning to see the wood for the trees: Deep Laser, Georgi Tinchev, University of Oxford
16.10 Wibergian learning of energy functions, Chris Russell, University of Surrey and Alan Turing Institute
16:30 Discussion Panel

Registration

Book online at bmva.weebly.com
BMVA Members £16, Non-Members £36 (in both cases lunch is included)

Andrew Gilbert
University of Surrey
email: a.gilbert@surrey.ac.uk

Visual Image Interpretation in Humans and Machines: Machines that see like us?

One-day BMVA symposium in London on Wednesday 10 April 2019
Chair: Andrew Schofield (Aston University)

bmva.weebly.com

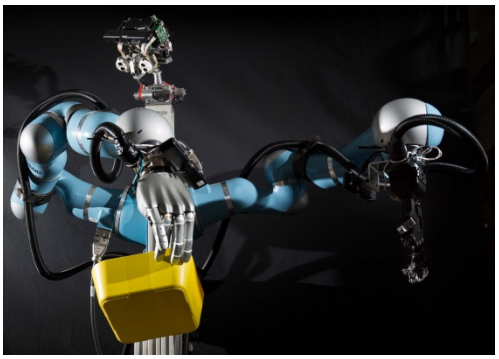

Image © University of Birmingham: reproduced with permission

Call for Papers

The object recognition and the game playing performance of deep convolutional neural networks now equal or surpass those of humans. Deep neural networks share some features in common with the human visual system including multiple layers of processing with the early layers being convolutional in nature. Moreover, techniques such as

representational similarity analysis show that appropriately trained neural networks develop representation spaces similar to that of inferior temporal cortex which is known to support object recognition in humans. These results and the superficial similarity in network structure between artificial and biological neural networks lead some to conclude that the former is a good functional model for the latter.

In contrast, others note that deep neural networks are easily fooled by image manipulations that are barely noticeable to humans or by specially constructed image elements that trick artificial networks but are seen but ignored by humans. There are also stark differences between artificial and biological networks with mean features of the latter omitted from artificial systems. There are differences too in the style and rates of learning and generalisation between humans and machines. Such findings suggest that models of human vision should be quite different from deep neural networks.

This one-day meeting will consider the issues in human and machine vision, discussing how artificial neural networks might be augmented with more biologically plausible features with the aim of making them more robust, and alternatives to neural network models and how their performance compares to the state of the art and human vision.

Submission Deadline: All those interested in presenting at this meeting are invited to submit an abstract of their talk at bmva.weebly.com by 6 February (firm deadline).

Andrew Gilbert
University of Surrey
email: a.gilbert@surrey.ac.uk

High-Performance Computing for Computer Vision

One-day BMVA symposium: BCS London on 22 May 2019
Chairs: Giuseppe Ciaccio and Nicoletta Noceti – Università degli Studi di Genova (IT)

Keynote Speakers

- Tae-Kyun Kim, Imperial College London
- Marta M Betcke, UCL
- Juan Gómez Luna, ETH Zürich

Call for Papers

The last decades have witnessed amazing progress in the field of Computer Vision, where state-of-the-art methodologies provide astonishing performance. At the root of this growth is the conjunction between advances in computer architecture and availability of large amounts of data: these have triggered new lines of research while shedding renewed light on classical Computer Vision tasks, considered intractable from a practical standpoint until a few years ago. New high-performance computing platforms and paradigms (many-core processors, FPGA accelerators, large-scale clusters, cloud computing) are gaining momentum and becoming mainstream.

The interplay between algorithms/methods for Computer Vision and high-performance computing architecture/ programming thus continues to provide interesting challenges to researchers in the field. The goal of this one-day meeting is to provide a view on trends and issues in the use of modern high-performance and large-scale computing platforms for the challenges of Computer Vision.

Topics of interest include, but are not limited to:

- high-performance computing in image processing and computer vision
- computer vision on large-scale/distributed/cloud platforms
- machine/deep learning for visual data GPUs and FPGAs for visual processing of real-time visual tasks and energy-efficient computer vision.

We encourage submissions from both academics and industries. The work can be in progress or recently published, or it may describe novel or prospective research. Please submit a short abstract summarizing your contribution: this may include links to illustrations, demonstration material or papers giving more details.

Submission deadline: All those interested in presenting at this meeting are invited to submit an abstract of their talk at bmva.weebly.com by 6 March (firm deadline).

Registration

Book online at bmva.weebly.com
BMVA Members £16, Non-Members £36 (in both cases lunch is included)

Andrew Gilbert
University of Surrey
email: a.gilbert@surrey.ac.uk

MIUA 2019 – Call for Papers

MIUA 2019 will be held at the University of Liverpool on 24–26 July 2018. The submission deadline is 18 February 2019.

MIUA 2019 is the 23rd conference of the Medical Image Understanding and Analysis series organised in the United Kingdom for communicating research progress in biomedical image analysis. Its goals are the dissemination and discussion of research in medical image processing and analysis. All researchers in medical image analysis are encouraged to attend, including mathematicians, computer scientists, bioinformaticians, clinicians, engineers and bioscientists. Together, we aim to encourage growth and raise the profile of this multi-disciplinary field. The conference features keynote speakers, tutorials, workshops, and oral and poster presentations.

Paper submissions:

Authors are invited to submit full papers of length between 8 and 12 pages (1 column – LNCS Springer format) showing original research contributions in medical image analysis and processing. The conference proceedings will be published in the Springer CCIS – Communications in Computer and Information Science and there will be a special issue in the Journal of Imaging for selected papers.

Important dates

Special Session Proposals due:	27 January
Paper Submission deadline:	18 February
Author Notification:	1 April
Camera-ready papers due:	12 April
Early-bird registration due:	28 April

Programme:

We will precede MIUA with a two-day Workshop on Image Processing Techniques and Applications, which can be booked separately or alongside the conference. MIUA will take place Wednesday–Friday; and, on Saturday, we are planning a social/networking event at the Snowdonia National Park.

22–23 July:	IPTA 2019
24–26 July:	MIUA 2019
27 July:	Social/Networking Event.

We have five exciting keynote talks planned for MIUA 2019 from leaders in the field of medical imaging from the UK, Germany and the USA:

- Dianggang Shen, Jeffrey Houpt Distinguished Investigator and Professor in the Department of Radiology and BRIC at UNC-Chapel Hill. He is a Fellow of IEEE and Fellow of The American Institute for Medical and Biological Engineering (AIMBE).
- Carola-Bibiane Schoenlieb, Professor of Applied Mathematics and head of the Cambridge Image Analysis (CIA) group at DAMTP, University of Cambridge. She is Co-Director of the EPSRC Centre for Mathematical and Statistical Analysis of Multimodal Clinical Imaging.
- Olaf Ronneberger, Professor of Computer Science at the Institut für Informatik at Albert-Ludwigs-Universität Freiburg and Chair of Pattern Recognition and Image Processing. Olaf is also a senior research scientist at DeepMind Health.
- Sebastien Ourselin, Head of the School of Biomedical Engineering & Imaging Sciences at King's College London. He is Director of the EPSRC Image-Guided Therapies UK Network+ and co-founder of Brainminer.
- Alejandro Frangi, Diamond Jubilee Chair in Computational Medicine at the University of Leeds. He leads the Centre for Computational Imaging and Simulation Technologies in Biomedicine and serves on the Scientific Advisory Board of EIBIR.

You can find more information on our website at https://www.miua2019.com

UNIVERSITY OF LIVERPOOL

Dr Bryan M Williams
University of Liverpool
email: bryan@liverpool.ac.uk

Report on AVSS 2018

AVSS 2018 took place Auckland, New Zealand on 27–30 November 2018.

AVSS is the premier annual international conference in the field of Video and Signal-Based Surveillance. It is a forum of reference for the field which offers the opportunity to meet and foster collaboration, with a tradition of participation from the worlds of research, industry, and relevant government agencies. Its primary aim is to promote the advancement of signal processing, image and video processing, audio processing, pattern recognition, and computer vision in the context of surveillance.

AVSS takes place every year and has been hosted across Europe, America and Asia. This year, it was hosted for the first time in New Zealand by the Auckland University of Technology (AUT). The first day was a day of workshops and tutorials covering such topics as traffic monitoring, soft biometrics, human behaviour analysis and digital forensics. The tutorial was given by Professor Ioannis Pitas (Aristotle University of Thessaloniki, Greece) on deep learning and computer vision techniques in the context of multi-drone vision in the context of media production (e.g., filming sporting events using multiple drones).

Across the three main conference days, the time was divided between oral presentation and poster sessions, with the poster sessions timed to overlap with breaks to give plenty of time for discussions over coffee. Topics ranged from object detection, recognition and tracking, to behaviour analysis and scene understanding. Deep learning was of course a popular theme, but it was also good to see many classical approaches still being put to effective use.

We were also treated to four keynote speeches from leading academics. Professor Nikola Kasabov (Auckland University of Technology) spoke about spiking neural networks which aim to more closely match how brains process information. Professor Ramesh Jain (University of California, Irvine) presented his current research passion on addressing health issues using cybernetic principles, building on the progress in sensors, mobile, processing, and storage technologies. Professor Edward Delp (Purdue University) described his work in image-based plant phenotyping (estimating traits such as locations, leaf area, canopy cover, number of leaves per plant, etc.) using UAVs and ground-based sensing. The final keynote was presented by Professor Mengjie Zhang (Victoria University of Wellington).

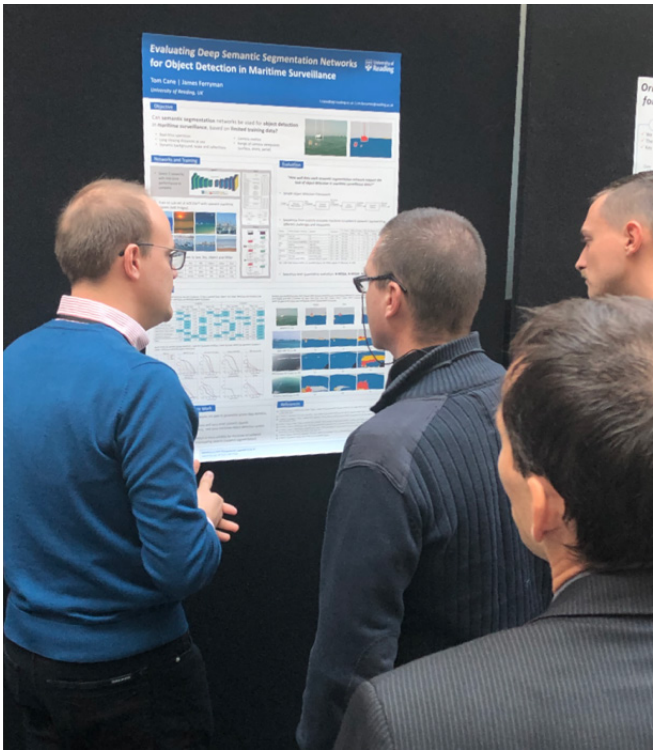

I presented a poster and 'spotlight' talk on my paper, "Evaluating deep semantic segmentation networks for object detection in maritime surveillance". The objective of this work was to establish how well deep semantic segmentation networks could support the task of object detection in maritime surveillance videos, based on very limited training data. In addition to gaining presentation experience, it was also a great benefit to hear other research relating to surveillance applications. There was only one other paper focussing on the maritime domain, but there were many transferable ideas from other applications which I look forward to exploring in the maritime context. It was also highly valuable to get input from other researchers on new directions to look into to extend my work.

The BMVA travel bursary made it possible for me to attend AVSS and discuss my work with other researchers in similar areas, as well as completely different ones. It was also an excellent opportunity to engage with other students and experienced academics from around the world. I would like to thank the BMVA for its generous support.

Tom Cane
University of Reading
email: t.cane@pgr.reading.ac.uk

Successes at the University of Surrey!

Professor Adrian Hilton

The Institute of Engineering and Technology (IET) has awarded Professor Adrian Hilton one of its highest honours – the IET Achievement Medal. He was presented with the award on Wednesday 14 November 2018 at the Brewery in London. The award recognises Professor Hilton "for his outstanding and sustained contribution in the field of computer vision engineering". Professor Hilton has been Director of the Centre for Vision, Speech and Signal Processing (CVSSP) at Surrey since 2012.

Professor John Collomosse

John Collomosse has recently been promoted to a personal chair in Computer Vision at Surrey's Centre for Vision Speech and Signal Processing (CVSSP). His main interests are in Deep Learning and applying Big Data problems in visual media.

Dr Jean-Yves Guillemaut

Jean-Yves Guillemaut was recently promoted to a Senior Lectureship at the CVSSP. His research centres on scene modelling from multi-view video input, with a focus on complex dynamic scenes.

Dr Andrew Gilbert

Andrew Gilbert has been appointed as Senior Lecturer in Film and Video Production within the Department of Music and Media at The University of Surrey. Andrew is well known to all members of the BMVA from his vitally important role as BMVA Meetings Officer – a duty he has undertaken with dedication and exceptional efficiency. Hopefully, his interesting move will serve to augment his links with the BMVA. We wish him the best of luck in his new position.

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