BMVA News

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

Volume 29 Number 3
March 2019

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 June 2019.

Contents

Editorial: *Teaching in an Age of Change*

A good many of the advances in physics over the past century or so involved the bombardment of atoms, nuclei and fundamental particles with electrons and other particles; in addition, the application of electromagnetic waves paid a considerable role, via the development of NMR and other important spectroscopic tools. Indeed, it became standard to probe materials in various controlled ways in order to get them to divulge their secrets. Interestingly, this systematic analytic approach was also rather independently adopted when we later tried to analyse images by using edge, corner, hole, circle, ellipse, line and other detectors to get them to divulge *their* secrets. But the underlying systematic analytic approach was still highly apparent. Whereas, of late, the emphasis in image analysis and computer vision has shifted to use of neural networks, maybe we should not read too much into this: *anything* you can throw at the data to find how it ticks could constitute a potentially viable procedure and lead to valuable results: the idea is that carefully probing the data gives you a significant chance of finding something useful.

Of course, throwing random entities or tools at materials or data is not unlikely to produce chaos. But while generating a totally new way of doing something might lead to a high chance of producing a complete shambles, it might also lead to a tiny chance of finding something highly unexpected and useful. So there is some premium on trying things out: and bear in mind the oft-repeated saying that if you don't attempt the absurd you won't achieve the impossible!

A lot has recently been said about the scientific value or otherwise of using neural-based approaches, but it seems distinctly possible that the next year or two could see a sudden definitive breakthrough in this area (reflect on the eventual successful discovery of the Higgs boson in 2012 and the failure to develop a provably applicable string theory over more than 30 years). With all the world-shattering impact this might have, or the chaos that might ensue from *thinking* that it has happened, we need to re-align our working lives accordingly. In particular, we need to decide which of the legacy methods should be retained and which of the new methods should be promoted. The BMVA CV Summer School is a case in point: should it retain all the standard approaches? And to what extent should it embrace the new, whether fully proven or not? It seems clear that a regular injection of new material is of some value. Indeed, this must be the case in all teaching situations. Just as a car learns to follow the road by updating the route and unlearning older information about the journey,

[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.

so we as teachers need to update our material and juggle our course modules to adapt to the developing situation.

From another point of view, it is often said that school courses such as GCSE and A-level are being dumbed down and that standards have dropped vastly since we ourselves took them. But we must compare like with like. And today's schoolchildren have to cope with huge advances in knowledge and technology that were unthought of when we were children – microelectronics, supercomputers, the Internet, social media, capability for searching for information in seconds using only a few keywords, and so on. Small surprise, then, that balances must be struck and that in some areas standards or content may indeed have fallen. All this must be managed with sensitivity and eased forward in acceptable directions.

But what is the best algorithm for achieving optimal adaptation? As hinted above, add a small amount of new and simultaneously do a small amount of forgetting of the old. However, I am here reminded of the problems I had in my O-level English – of learning to précis pieces of text. Theoretically, any piece of text could be reduced to any proportion between 0 and 1. But when you try to do this you find funny effects – one or two sentences refuse to be shortened: they must either be included in their entirety or eliminated altogether. The lesson that précising is nonlinear was one that it took me some time to master, though it has usefully stuck with me ever since. So the algorithm to use is not a simple moving average: possibly it but may be closer to a moving median. Alternatively, it is more likely to be something involving real thought – and not by an AI machine, which can't possibly at this stage know enough of the past and the future and of real life to form sound judgements of this type. As teachers, it seems that we have to take continuing responsibility for how we map what we know about the external world into our teaching programmes. An important remaining ambiguity is whether whoever pays the piper calls the tune.

Final thought: at some stage in the not too distant future, the advent of autonomous vehicles and other advances in use of vision may be pressing us to add ethics to our curriculum.

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

BMVA Symposium: High-Performance Computing for Computer Vision

One-day BMVA symposium, to be held at the BCS, London on 22 May 2019.
Chairs: Giuseppe Ciaccio and Nicoletta Noceti, Università degli Studi di Genova

Keynote Speakers

- Tae-Kyun Kim, ICL
- Marta M Betcke, UCL
- Juan Gómez Luna, ETH Zürich

Programme

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 Computer Vision.

9:30 Keynote: Tae-Kiun Kim "Machine learning for 3D vision"
10:10 Oscar Rahnama "Developing real-time stereo algorithms for FPGA-based platforms"
10:30 Alessio Tonioni "Real-time self-adaptive deep stereo"
10:50 Coffee Break
11:20 Matteo Poggi "Real-time monocular depth estimation without GPU"
11:40 Armin Mustafa "4D Vision in the wild"
12:00 Keynote: Juan Gomez Luna "From GPGPUs to processing-in-memory: high performance and energy efficiency for computer vision workloads"
12:40 Lunch
13:50 Keynote: Marta Betcke "High-resolution tomographic imaging with learned priors"
14:30 Olivia Wilson "Mega maps: large scale deep learning of aerial imagery"
14:50 Bolin Pan "Photoacoustic imaging using curvelet sparsity with optical flow constraints"
15:10 Coffee Break
15:40 Nargiza Djurabekova "Reconstructing ankle and foot dynamics using optical flow"
16:00 Elena Nicora "Why high-performance computing is important in robotics"
16:20 Final discussion
16:30 End

BMVA Symposium: Deep Learning in 3D

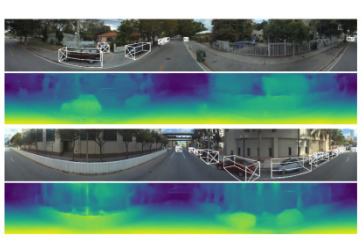

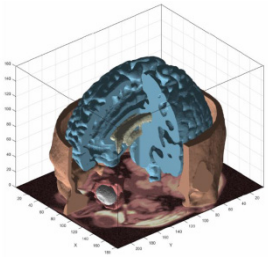

This technical meeting was held at the BCS headquarters in London on 20 February 2019. It aimed to explore the key challenges of combining deep learning with 3D vision.

Chairs: Chris Holder, Christopher Willcocks and Grégoire Payen de La Garanderie, Durham University
Organiser: Andrew Gilbert, University of Surrey

The morning session started with a keynote from Adrian Hilton on 3D shape capture across time – hence his title: "4D Vision". Through his work with the entertainment industry, he highlighted specific challenges for 3D vision moving from controlled performance capture environments to real-world unconstrained environments.

In the next talk on 3D Pick & Mix, Adrián Peñate-Sánchez reasoned not only about whole objects but also about their parts. This work learns a shared embedding between images and shapes using synthetic images and a realistic renderer. This embedding can be used to retrieve 3D models of furniture using a combination of query images representing the desired parts.

Subsequently, Olivia Wiles talked about recovering the 3D shape of sculptures from a single image. The dense depth estimator is trained by extracting sparse point correspondences in a multiview setup – an approach that she called Pointillism. During training, camera pose is estimated using SIFT and assuming an affine camera model.

The next talk, by Andrew Gilbert, delved into volumetric performance capture using a minimal number of viewpoints. The work is based on a convolutional auto-encoder. The network is trained on 8 input views: however, only 2 inputs are required at test time.

Finally, in the last talk of the morning session, Viswadeep Sarangi looked at the clinical evaluation of machine learning approaches for 3D gait classification. Interpretability is a key part of the approach due to the medical setting and can also be used to provide insights on which gait features are relevant to classify different gaits.

The afternoon session started with a keynote from Alex Kendall who showed that 3D geometry understanding is still important in an era of end-to-end deep learning. For instance, an end-to-end lane-following vehicle can be trained in the real-world – rather than simulated – using reinforcement learning in very few iterations if the solution is adequately constrained using geometry.

In the first talk of the second session, Michael Edwards introduced Graph Convolutional Network as a technique to replicate convolutions and pooling on a graph structure rather than the classical 2D or 3D grid representation of an image. This is used to reduce the size of the required network to segment structures in large 3D medical images.

Our last speaker, Georgi Tinchev, presented his work on relocalising in a map using a laser-based point cloud. The approach allows localisation in difficult environments such as a forest. Moreover, the model is kept small in order to run within the limited computational payload of a drone.

Last but not least, we had four spotlight presentations on a plug-and-train loss for single view 3D reconstruction by Eduard Ramon Maldonado; ground classification in 3D point clouds of large-scale heritage sites by Dimitrios Makris; fruit detection with 3D deep neural networks by Justin Le Louëdec; and the recovery of superquadric parameters by Aleš Jaklič.

The chairs would like to thank Adrian Hilton, Alex Kendall, and all the speakers and attendees for making this meeting a success.

Grégoire Payen de La Garanderie
Durham University
email: gregoire.p.payen-de-la-garander@durham.ac.uk

Upcoming Meetings

We have 5 exciting varied meetings planned so far in 2019 – check out the website, www.bmva.weebly.com, for all the details.

Upcoming meetings include:

10 April	Visual image interpretation in humans and machines: the role of learning and experience – Andrew Schofield
22 May	Computer vision + high-performance computing – Nicoletta Noceti & Giuseppe Ciaccio
17 July	Geometry and deep learning – Vassileios Balntas & Krystian Mikolajczyk
25 September	Video understanding – Hilde Kuehne, Jeurgen Gall, Dima Damen & Ivan Laptev
27 November	Generating data in computer vision and machine learning – Nathan Olliverre.

Registration for all BMVA Symposium Meetings
Note the different website to register for BMVA Technical Meetings: bmva.weebly.com
Book online at bmva.weebly.com: BMVA Members, £16; non-Members, £36 (both prices include lunch)

Call for Abstracts: IET Human Motion Analysis for Healthcare Applications

This event will take place at Savoy Place, London on
26 June 2019
N.B. BMVA members can register at IET rates!

Advances in technology in human motion capture have been quite remarkable during the last decade. Specialised depth sensors are now embedded even on mobile devices. Recent progress in Deep Learning allows standard video cameras to capture 3D human motion. Their main advantages are their

non-intrusive nature, low cost, and widely available support for developers offered by large corporations or Open Communities. They inspired numerous healthcare-related ideas and projects were developed in areas such as Medical Disorder Diagnosis, Assisted Living, Rehabilitation and Surgery.

In Assisted Living, human motion analysis allows continuous monitoring of elderly and vulnerable people and their activities to potentially detect life-threatening events such as falls. Human motion analysis in rehabilitation provides the opportunity for motivating patients through gamification, evaluating prescribed programmes of exercises and assessing patients' progress. In operating theatres, surgeons may use a gesture-based interface to access medical information or control a telesurgery system. Human motion analysis may also be used to diagnose a range of mental and physical diseases and conditions.

This event will discuss recent advances in human motion sensing and provide an application to healthcare for networking and exploring potential synergies and collaborations.

Call for participation
Abstracts are invited from academia, industry and healthcare professionals related to healthcare applications of human motion analysis. The event programme will include keynote presentations and oral and poster presentations of accepted abstracts.

Topics of interest include but are not limited to:
computational techniques for human motion analysis
biomechanics
challenges and opportunities in healthcare for human motion analysis
gesture recognition
action recognition
action quality
facial expression recognition
gait analysis
fall detection
people tracking and identification
re-identification
identity concealment
gamification and serious games.

These should be related to healthcare applications such as
rehabilitation
assisted living
activity monitoring
medical surgery
gesture-based interfaces
diagnosis of medical disorders
behavioural change for healthy living.

Submit a one-page abstract
Submit a one-page abstract using the template at https://events.theiet.org/media/1160/human-motion-analysis-abstract-template.docx by 17 April 2019 to dimitrios.makris@ietvolunteer.org for a chance to deliver an oral or poster presentation at this June event.

Keynote speakers
Professor Antonis Argyros, University of Crete, Greece
Sarah Taylor, St George's University Hospital, UK
Simon Yarwood, InnovateUK, UK
Aliah Shaheen, Brunel University London

Chairs
Professor Dimitrios Makris, Kingston University, London
Dr Daniel Abasolo, University of Surrey, UK

The event is co-organised by the IET Vision & Imaging TPN and the IET Healthcare TPN.
It is also sponsored by the BMVA.

Key dates
Call for participation deadline: 17 April
Notifications of acceptance: 13 May
Final programme confirmed: 17 May
Event: 26 June.

Register here: https://events.theiet.org/human-motion-analysis-for-healthcare-applications/

MIUA 2019

University of Liverpool, UK, 24–26 July 2019

MIUA 2019 is the 23rd conference of the Medical Image Understanding and Analysis series organised in the UK 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.

Early-bird registration deadline: 28 April 2019.

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 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.

Conference topics

- big data processing
- biomarker discovery
- clinical and scientific evaluation of imaging studies
- computer-aided pathology
- computer-assisted surgery
- data compression
- data fusion
- decision support
- human computer interaction
- image enhancement
- image interpretation
- image-guided intervention
- image reconstruction
- image perception
- image registration
- image segmentation
- intelligent imaging systems
- machine learning in imaging
- modelling and simulation
- motion analysis
- multi-modality image analysis
- pattern and feature recognition
- protocol development and standardization
- quantitative image analysis
- shape analysis
- software development
- super-resolution algorithms
- statistical methods in imaging
- texture analysis
- virtual reality.

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

Metropolitan Police Commendation!

Emeritus Professor Roy Davies, Department of Physics, Royal Holloway University of London, has recently been commended for his dedication and service to the Metropolitan Police. Professor Davies is internationally renowned for his expertise in machine vision and has assisted with research into the use of video analysis in collision reconstruction.

Detective Superintendent Andy Cox
Metropolitan Police Service

BMVC CARDIFF 2019

30th British Machine Vision Conference
9–12 September 2019
Cardiff University

https://bmvc2019.org

The British Machine Vision Conference (BMVC) is one of the major 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

Cardiff University on 9–12 September 2019. BMVC 2019 is a high-quality single-track conference, comprising oral presentations and poster sessions. It features three keynote presentations and two conference tutorials, and has associated workshops on the last day of the conference. Full-length high-quality papers cover theory and/or application areas of computer vision. 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 visualization, interaction, and graphics
- Object detection and recognition
- Shape-from-X
- Video analysis and event recognition
- Illumination and reflectance.

Important dates

Submission deadline:†	29 April
Author Feedback:	10 June
Camera-Ready deadline:	14 July
Early Registration Deadline:	31 July
Conference Tutorials:	9 September
Main Conference:	10–12 September
Workshops:	12 September.

†*Note: there will be no extensions this year.*
See also https://bmvc2019.org/dates/

Keynote speakers

- Cordelia Schmid, French Institute for Research in Computer Science and Automation
- Antonio Torralba, MIT
- Adrian Hilton, University of Surrey.

Tutorials

- Professor Michel Valstar, University of Nottingham, “Computational Face Analysis”.
- Professor John Collomosse, University of Surrey
- Professor Ondrej Chum, Czech Technical University in Prague, “Robust visual search and matching”.

Sponsorship

If you are interested in sponsoring BMVC 2019, would like to discuss a custom package, or require more information, please contact the Sponsorship Chair. More information on the sponsorship packages is available at:

https://bmvc2019.org/sponsorship/

Dr Kirill Sidorov
Cardiff University
email: sidorovk@cardiff.ac.uk

BMVA CV Summer School, 2019

UNIVERSITY OF LINCOLN

BMVA runs an annual Computer Vision Summer School aimed at PhD students in their first year, though it will also be beneficial to other researchers at an early stage in their careers. Despite the title, students from non-UK universities are welcome to attend, as well as students from UK universities. Places are limited to ensure good interaction in lab classes.

The 2019 Summer School will take place at the University of Lincoln during 8–12 July. It will consist of an intensive week of lectures and lab sessions covering a wide range of topics in Computer Vision. Lecturers are researchers from some of the most active research groups in the UK and abroad.

In addition to the academic content, the Summer School provides a networking opportunity for students to interact with their peers, and to make contacts among those who will be the active researchers of their own generation.

Nicola Bellotto
University of Lincoln
email: cvss@lincoln.ac.uk
website: http://cvss.blogs.lincoln.ac.uk

This page is deliberately left blank.

Nominations for the BMVA Executive Committee

Nominations are requested for the forthcoming election of Executive Committee members of the BMVA. 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 participate in roughly five committee meetings per year, taking place typically in London and Birmingham. Reasonable travelling expenses to attend meetings are paid by the BMVA.

Completed nomination forms should be sent to the BMVA Chair and must be received by 1 July. The nomination form must be signed by the individual standing and by one other member, and should also include a brief biographical statement for distribution to BMVA members.

The Executive Committee normally consists of ten elected members, five of whom are elected each year. Details of the current members may be found at:

http://www.bmva.org/executive_committee

If more than five nominations are received for the elected places then a postal ballot will be held. Voting papers will be sent out in early July and will need to be returned by the end of August. Each member will be able to vote for up to five candidates. The results will be announced at BMVC and in *BMVA News*.

Adrian Clark
BMVA Chair
email: chair@bmva.org

Nomination Form for the BMVA Executive Committee

To be completed by the Nominator

As a fully paid up member of the BMVA, I,
Name: .....
Address: .....
.....
.....
wish to nominate:
Name: .....
Institution: .....
for the Executive Committee of the BMVA.
Signed: .....

To be completed by the Nominee

I am a fully paid up member of the BMVA and am willing to serve for a minimum period of one year on the BMVA committee.
Name: .....
Address: .....
.....
.....
email: .....

I attach below a brief biography for distribution to BMVA members.
Signed: .....

Brief Biography of Nominee

Please send completed nomination form by post or email to:

Dr Adrian F Clark
Computer Science and Electronic Engineering
University of Essex
Wivenhoe Park
Colchester
Essex, CO4 3SQ
email: chair@bmva.org