

BMVA News

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

Volume 27 Number 4
June 2017



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 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 10 September 2017.

Contents

Editorial: <i>Short-cuts v. Scientific Rigour</i>	1
Travel Bursaries for International Conference Attendance.....	2
BMVA's Thesis Archive	2
BMVC 2017 – <i>Call for Participation</i>	3
BMVA Distinguished Fellow 2018 – <i>Call for Nominations</i>	3
Message from the BMVA Meetings Organiser.....	4
BMVA Symposium: Computer Vision in Cancer.....	4
BMVA Symposium: Human Activity Recognition & Monitoring.....	4
Review of BMVA Meeting on Dynamic Scene Reconstruction.....	5
Copy for the Next Issue of BMVA News.....	6

Editorial: *Short-cuts v. Scientific Rigour*

We often imagine Greece as being the ultimate source of thinking and logical argument: in fact, its civilisation originated at least as far back as 1250 BC; evidence for this is clear from the clay tablets found at Knossos and Mycenae, which show that the language of communication was recognisably Greek even then – as in the case of the (syllabic) Linear B words *ti-ri-po-de* (tripod) and *Ko-no-so* (Knossos). We are now in the 21st century of the modern era, but only in the last 300 years or so has true experimentally verifiable science evolved: indeed, it can reasonably be said

that Newton (1643–1727) was the father of the new scientific age.

Newton's laws were accurate to a high degree of precision (parts in a million) and represented absolute truth – until Einstein eventually showed how to make them even more accurate and absolute. Indeed, many fundamental atomic constants are now known to vastly better than 6-figure accuracy, and in physics at least, there is the pressure to explain everything – even in unified theories covering gravity, quantum mechanics and radiation – with this degree of success. Never mind for now that string theory hasn't provided the expected links needed for a theory of everything. Rather, my point is that we are punch-drunk by the idea that everything can be linked in this way. However, it is a pity that non-scientific disciplines don't aspire to this level of achievement: in particular, those that are numeracy-based such as economics are grossly limited by the fuzziness of the data they have to grapple with.

But there are other disciplines that lie somewhere in between, such as meteorology and vulcanology, that have also had limited success. Weather forecasts, though vastly improved even in my own lifetime, are still unreliable, while prediction of volcanic eruptions is so inaccurate as to be of little help in saving lives: interestingly, even the wild animals in the vicinity seem to know that something is afoot well before we do. Perhaps this provides a clue as to an underlying problem – namely lack of the right sort of sensory data, though one imagines that even if that were available, we wouldn't know how to use it to make the right predictions. However, in weather forecasting, the past 50 years has taught us the relevance of the butterfly effect – i.e., that even the flap of a butterfly's wings can upset the input data so much that ambiguity is introduced into the calculations, thereby preventing accurate predictions from being made. The introduction of chaos has the crucial effect of making deterministic Newtonian calculations impossible and cancelling out Newton's hopes that if only all the right data were available, pure calculation would permit the entire development of the universe to be mapped out. In fact, the chaotic component implies that only by gathering larger and larger amounts of input data can predictions be stretched out further into the future.

Meteorology was probably the first science to make the move to 'big data'. Other subjects will follow, perhaps more

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.

easily, as our computerised age involves collection of huge amounts of data on many fronts – e.g., financial data, accident data, personal data, genome data, and of course everything published under the auspices of Twitter, Facebook, and the like. What is different about social networks is that people are keen to add to the datasets for free (or, more accurately, they *act* as if they are), so it is almost effortless to pick up any relevant big data.

Meanwhile, big data is vital for deep learning networks, which are the modern *cri de coeur* in the vision community. By which I mean that it has become a necessity to train vision machines to carry out certain tasks because we are not clever enough to design them to do those tasks in conventional ways. Alternatively, it has become fashionable to use deep learning just to be seen to be ‘keeping up with the Joneses’, even when we are clever enough to manage without it.

But in truth, we could actually have reached the limit of what accurate equations (such as Newton’s) will do for us, so that conventional ways are no longer able to provide a way forward. One possible scenario is that we have taken all the easy routes up the mountain, and now all that is left is the far more difficult fuzzy path, in which there is no paving or even obvious direction forward. We were supremely clever in getting this far, with all sorts of elegant deterministic methods – edge detection, Hough transforms, RANSAC, invariants, the fundamental matrix, Kalman filters, particle filters, SIFT, SURF, HOG – you name it! But in the end we were found out because of the variability of real data in real applications: real data is fuzzy. And seeing the wood for the trees becomes far more difficult and requires vast data inputs to train the algorithms; *and* the algorithms are either unknown or need training in unusual ways; *and* we are not clever enough to design conventional algorithms to implement these unusual ways. Of course, it may only take a little more time to design the right sort of algorithms, but meanwhile, our deep learning colleagues are beating us to it. (Again, ‘if you can’t beat ’em, join ’em ...’) The only problem is that (as a good few stalwart scientists have been warning), we don’t really know what these deep networks are thinking, i.e., how they work, or *how reliably* they work. In any case, science advances only in small steps, each of which is examined minutely and tested rigorously for accuracy and reliability: what deep networks do does not fall into this category of sound science – and thus computer vision is falling into the hands of magicians, quacks and entrepreneurs.

Another view is that this is only a phase that CV is going through. Deep networks presently have the advantage that they offer the promised land by a known route, which will move the subject forward and at the same time offer an existence theorem of what is possible; and also show exactly how much big data is necessary to obtain the solutions. Once all this is known, the science itself will have been moved forward, though it will be up to real scientists to implement it by conventional means. But there is a caveat: we are running into murky waters, and the data will be fuzzy, so the science will have to show how to manage with that; i.e., we will be moving forward not from where we were yesteryear but from a new starting point. The new methodology will have to be called neural-replacement machine learning, and will be completely probability rather than statistically based. However, I’m not really claiming to have successfully combined the equivalent of gravitation

and quantum mechanics. I’m just offering a caveat, deep learning aficionados – back to the *real* future! And remember that when a player catches a cricket ball falling from high in the sky, he ignores his neural circuitry and merely runs to keep the ball at a constant elevation ...

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

Travel Bursaries for International Conference Attendance

The BMVA provides a number of travel bursaries for student members of the BMVA who are research students at UK institutions to present their work at significant international conferences within the BMVA’s remit. The maximum amount of a bursary is £750. In return for the bursary, students are asked to write a conference report for BMVA News – most of the conference reports you read in these august pages are from students who have received bursaries – or do some work for the BMVA. Details on eligibility and the application procedure are outlined on the BMVA website: note that the procedure has recently changed, and that strict deadlines now apply – as clearly indicated on the BMVA website.

Professor Lourdes Agapito
UCL
email: l.agapito@cs.ucl.ac.uk

BMVA’s Thesis Archive

You probably already know that the BMVA archives PhD theses of students studying at UK institutions: in fact, it has theses dating back as far as 1996. One or two other national vision associations have started archiving theses, so the UK is, as usual, leading the way in this.

However, our thesis archive continues to evolve in line with trends in open research. From this year, we are prepared to archive modest amounts of ancillary material such as software and datasets along with theses so that researchers will be able to reproduce the findings more easily. We are also happy to “retro-fit” such ancillary material with theses already in the archive.

The thesis archive is a useful resource. For example, it contains all winners of the Sullivan Thesis Prize. I would like to encourage everyone, including supervisors and students, to submit their theses to the archive. It is quick and easy to do, coming down to not much more than an email and providing a PDF version of the thesis text – see <http://www.bmva.org/theses:top> for details. Note that you do not have to have just completed your PhD to put your thesis in the archive: a recent submission was a decade after the award of the degree.

Adrian Clark
BMVA thesis archivist
email: alien@essex.ac.uk



28th British Machine Vision Conference, 4–7 September 2017, Imperial College London, London, UK

<http://bmvc2017.london/>

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 28th BMVC will be held at Imperial College London on 4–7 September. BMVC 2017 is a high-quality single-track conference, comprising oral presentations and poster sessions. It features two keynote presentations and four conference tutorials, and has associated workshops on the last day of the conference, including a workshop keynote.

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

Keynote speakers

- Richard Szeliski, Facebook, USA, <http://szeliski.org/RichardSzeliski.htm>
- Pietro Perona, California Institute of Technology, USA, <https://www.vision.caltech.edu/Perona.html>

Tutorial speakers

- Michael M. Bronstein, USI Lugano/Tel Aviv University/Intel, <http://www.inf.usi.ch/bronstein/>
- Lourdes Agapito, UCL, <http://www0.cs.ucl.ac.uk/staff/L.Agapito/>
- Shimon Whiteson, University of Oxford, <http://www.cs.ox.ac.uk/people/shimon.whiteson/>
- Andreas Geiger, MPI Tübingen / ETH Zürich, <https://avg.is.tuebingen.mpg.de/person/ageiger>

Workshop keynote

Jamie Shotton, HoloLens Science team at Microsoft, Cambridge, <https://www.microsoft.com/en-us/research/people/jamiesho/>

Important dates

Author Notification:	4 July
Workshop Paper Deadlines:	6–10 July
College Accommodation Deadline:	11 July
Camera Ready deadline:	18 July
Author Registration Deadline:	18 July
Early Bird Registration Deadline:	11 August
Conference Tutorials:	4 September
Main Conference:	5–7 September
Workshops:	7 September

Organizers

- General chairs: Tae-Kyun Kim, IC London; Stefanos Zafeiriou, IC London
- Program chairs: Krystian Mikolajczyk, IC London; Gabriel Brostow, University College London
- Advisory board: Andrew Davison, IC London; Maja Pantic, IC London; Yiannis Demiris, IC London
- Publicity chair: Giovanni Maria Farinella, University of Catania, Italy
- Tutorial/workshop chairs: Yannis Panagakis, IC London; Tania Stathaki, IC London
- Sponsorship/demo chairs: Ben Glocker, IC London; Stefan Leutenegger, IC London
- Local arrangements chairs: Rigas Kouskouridas, IC London; Anastasios Roussos, IC London; Guillermo Garcia-Hernando, IC London
- Website chairs: Athanasios Papaioannou, IC London; James Booth, IC London; Patrick Snape, IC London

Sponsorship

If you are interested in sponsoring BMVC 2017, 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 <http://bmvc2017.london/sponsorship/>

BMVA Distinguished Fellow 2018 – *Call for Nominations*

The BMVA Executive Committee seeks nominations for the *Distinguished Fellow 2018* award. This prestigious award is given to one person only each year in recognition to his/her services to the British Machine Vision community. The nominees must be distinguished researchers, based in the UK, who have contributed significantly to the field of research and the reputation of the British Machine Vision Community both nationally and internationally. Nominations, *with a few lines of rationale*, should be sent to Professor Roy Davies, Chair of the Distinguished Fellow Award Panel, by 1 September 2017.

Professor Roy Davies
Chair, Distinguished Fellow Award Panel
email: e.r.davies@rhul.ac.uk

Keynote Speakers: Professor David Hogg (University of Leeds), Dr Alessandro Vinciarelli (University of Glasgow), Professor Ian Craddock (University of Bristol), and Professor Yiannis Demiris (Imperial College London)

Call for contributions

This BMVA one-day meeting will present state-of-the-art developments in Human Sensing and is motivated by the attention of several computer science communities due to its connection to different fields of study. Human activity analysis and recognition has become a research area of great interest, as its strength in providing potential applications such as intelligent environments (smart home, smart vehicle, smart care home, smart factory, etc.), security and surveillance, human-robot collaborative tasks, human-machine interactions, assistive technologies, biomechanical study of athletes, physical activity and sedentary behaviour, virtual and augmented reality, physical therapy and rehabilitation, and many more.

Recent advances in visual, depth and inertia sensors, algorithms for data/signal acquisition and processing have led to advances in detection, tracking, analysis of human activities, as well as fundamental understanding of long-term modelling and recognition of human behaviour in complex scenarios. Human activity/behaviour includes varied modalities and numerous scales including single person, small group and larger group.

This one-day meeting will be dedicated to bringing together leading researchers, at various levels in their career, with expertise or strong interest in technical advances in activity analysis and recognition. The meeting also aims to bring together a collection of the latest approaches in this domain. We hope the meeting will stimulate future research, with both theoretical and practical perspectives, to stimulate further advances in the field.

We welcome contributions to this workshop in the form of oral presentations, posters and demos. Suggested topics include:

- Multi-modal (visual, depth and inertia sensors) human activity modelling and recognition
- Human motion analysis
- Scene analysis and understanding
- Real-time activity recognition and monitoring
- Social signal processing
- Emotion recognition
- Single-user, multi-user and group activity recognition
- Human pose recognition
- People detection and tracking
- Activity modelling and recognition through logic and reasoning
- Long-term modelling and recognition
- Data/pattern mining based approaches to activity recognition
- Context-aware activity monitoring
- Activity recognition for personalised services.

The above list is not limited. Therefore, you are welcome to submit an abstract on human activity recognition and monitoring related research. The piece of research work could be recently published, in progress, or novel. You may also include links or pointers to web-based materials, demonstrations or papers giving more details. We encourage

submissions from students, academics and industry, including interdisciplinary work and work from those outside the mainstream computer vision community.

Submission deadline

All those interested in presenting at this meeting are invited to submit a summary of their talk, by 28 September 2017, to <https://goo.gl/nXXy3f>

Registration

Book online at www.bmva.org/meetings: the costs – £16 for BMVA Members, £36 for non-Members – include lunch and subscription to the BMVA (in the case of non-members).

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

Review of BMVA Meeting on Dynamic Scene Reconstruction

London, 21 June 2017

Chairs: Marco Volino, Armin Mustafa, Jean-Yves Guillemaut

Scene reconstruction is now seen as a somewhat core component of computer vision, something that has been enabled by advances in both science and hardware. Problems that could not have even been contemplated five years ago are quickly being tackled and solved. This technical meeting provided a chance to discuss the field and, despite high temperatures outside, some 60 people attended the event.

After a brief opening from the chairs, Michael Zollhöfer started the day with a summary of some of the work being done at the Max Planck Institute, ranging from more traditional lines of reconstruction all the way through to new and novel techniques, including combining the facial features of one person with the facial expression of another. Also of note was an introduction to Opt, Michael's own GPU-based optimiser and solver, used to allow for real-time processing of traditionally difficult tasks.

Following on from this, Christian Richardt presented some fascinating work reconstructing scenes from non-ideal cameras (such as hand-held mobile phone cameras), producing good results from two shaky cameras situated some distance apart. Armin Mustafa then gave an overview of her two latest papers, before Jonathan Starck presented the latest works from Foundry, demonstrating the differences between academic problems and industrial applications. A brief insight into node-based effect processing later, Nadejda Roubtsova demonstrated her findings, determining reflectance in scenes and allowing for some element of relighting, as well as allowing for albedo changes to objects within a scene.

Following lunch, Lourdes Agapito gave us a brief history lesson of the field followed by some very striking facial models from RGB video. Produced without any kind

of deep learning, the results were certainly as vivid as promised from the talk title, capable of conveying emotion as well as being accurate representations of the captured face. Benjamin Biggs then gave an overview of his intended PhD work in animal tracking, intending to monitor animals in captivity and keep watch for prolonged periods. Returning to summaries, Dan Casas treated the audience to various advances from his department, including producing a body model for use in a VR environment from a pair of head mounted cameras.

Rilwan Basaru provided an excellent summary of his PhD work, using a pair of poor quality RGB cameras to determine hand depth. By leveraging a series of random forests, he demonstrated that lack of any calibration information and comparatively low resolution can be overcome to track hand depth effectively. Finishing the talks of the day, Stamatia Giannarou brought medical applications to the table, using medical imaging results to provide brain reconstructions suitable for surgeons to consult during brain tumour removal operations.

Finally, a short but lively panel session was held, during which two key topics were discussed. The first covered the divide between academia and industry, and the challenges that need to be faced in order to bring the two components more closely together; while the second focussed more on

datasets, and the desire to have newly created datasets released at the same time as any paper that uses them, so that replication of results (and later comparisons) can be achieved without the need to recreate suitable datasets.

Overall, the day was a success, with a large amount of excellent work being presented, as well a lot of interesting discussions being held in the panel session. Certainly, there are still many challenges to overcome in the field, however, as was summarized by Lourdes, a huge amount of progress has happened over the last few years, and certainly the next few look very promising.

Matthew Shere
University of Surrey
email: m.shere@surrey.ac.uk

Copy for the Next Issue of BMVA News

Please put the following entry and deadline in your diary: copy for the September issue of BMVA News should reach The Editor by 10 September 2017.