

BMVA News

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

Volume 27 Number 1
September 2016

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 December 2016.

Contents

Editorial: <i>Is our Research Subject to Relaxation Oscillations?</i> . 1
The Sullivan Thesis Prize 2016 2
The BMVA Elections 2
<i>BMVA Distinguished Fellow 2016</i> : Edwin Hancock 3
Distinguished Fellow Acceptance Speech 3
BMVC 2016 Statistics 4
BMVC 2016 Prize-winners 4
Around and About at BMVC 2016 5
Deep Learning in Computer Vision 12
Plants in Computer Vision 13
Transfer Learning in Computer Vision 13
Report on BMVA Computer Vision Summer School 14

Editorial: *Is our Research Subject to Relaxation Oscillations?*

Those who attended the latest BMVC at York will have experienced an extremely well-organised event, with an excellent programme of papers, posters and exemplary keynote addresses. Not only were we brought bang up-to-date with good talks and the opportunity to meet people in our specific areas, but also we heard exciting stories about preserving our (well, Japan's at least) cultural heritage, and

finally getting to understand the true picture with regard to the future of autonomous vehicles: here I am hinting at the *likely* real future, but without the hype of certain well-known companies, which over-emphasise what has been achieved and how soon it will all happen. To cap it all, the Banquet took place at the National Railway Museum in York, which during the day is a great spectacle of engines and trains from the past, but by night acquires a wonderful lighting of its own, with the exhibits appearing in colourful subdued shades, ready to further entertain us at a highly enjoyable dinner. No expense seemed to be spared, and the brass band was 'out of this world', both visually and acoustically, with well-known tunes being played faultlessly and with flare. I guess we actually paid for it ourselves, one hint being that – after applying too late for university accommodation – I actually managed to obtain bed and breakfast even more cheaply only yards from the town centre, the downside being that the one taxi ride I had to get to the conference cost me £15, though I later found that the bus was three times faster and infinitely (literally, in my case) cheaper. As for the town centre, it is a treat in itself, though many delegates might well have missed it altogether by travelling direct from station to university and vice versa.

Having been to York quite recently, and having missed the shortest street with the longest name – Whip-Ma-Whop-Ma-Gate – during my only taxi-ride I was unexpectedly treated to a sight of the name-plate on the street itself. What does such a name signify? Well, centuries ago, it was the place where petty criminals were flogged, so the reasons for the name can easily be visualised. However, it seems that the name was actually adapted from an earlier Anglo-Saxon name of not dissimilar rhythm 'Whit-Nour-What-Nour-Gate', meaning 'Neither-One-Thing-Nor-The-Other-Gate'. Of course, one also has to ask why such a street is called a gate, even when obvious town gates were also called gates, as in other parts of York. That is because the word gate is a very old Germanic one, being related to German 'Gasse' – which to this day means an alley or lane – and Swedish 'gata', which still means street. Interestingly, the area overlooking the Wye valley in South Wales is known as Symonds Yat, and here 'yat' is an old English name for a gate or a pass. (In Germanic, 'g' can be either hard or soft, and in Old English a soft 'g' sounded like a 'y'.) While I am on the subject, around the corner from Whip-Ma-Whop-Ma-

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.

Gate is a street called the Shambles. How could anyone give such a name to a street? In fact, this is where the word ‘shambles’ originated:² the point being that the Shambles had many butcher stalls, and the blood of slaughtered animals used to trickle down the roadway in full view of passers-by. This led people to describe other such repugnant environments as quite a ‘shambles’. My apologies if I’m either causing you to feel queasy or am leaving you with the feeling that you need to go back to York to see the sights that you missed while being ensconced at BMVC.

At conferences, I often find that one’s mind soon gets completely clogged up with new information – all of which needs to be memorised to gain full benefit from the lectures. No matter how hard I have tried, I have not overcome this problem. This time I resolved to try a different ploy – to look out for information on a rather restricted subset, specifically deep neural networks. I had recently read up many of the key events in this area over the period 2012–2015, and the architectural side seemed to be getting more and more exciting. Yet looking at the advance programme, I saw very little on this topic: indeed, only about 10–15% of the titles seemed to mention it, which was most unexciting considering my plan. However, on receiving the Abstract booklet, I realised that this was not the case. While only a small proportion of the papers had titles indicating deep neural network research, vastly more papers were actually using deep networks – but without including this in the titles. Were they pretending that deep networks are not a worthy subject, or merely using them as tools? Or perhaps, realising that deep networks are now *de rigueur*, there was no news value in this, and their own applications were clearly more important than mere everyday techniques.

But could such a view hold water? Well, the period 2012–2015 was a particularly fruitful one for deep networks: first, AlexNet capitalised on and supplanted LeNet; then ZFNet made advances in deconv nets leading to a better understanding of what convnets were able to achieve; then VGGNet took over supremacy with a brilliant masterstroke – both producing ‘Very Deep’ nets and showing that large depths need to be combined with the narrowest possible convolutional receptive fields (3×3) in order to optimise the system. This sounds very much like ‘Game, Set and Match’ to VGGNet, and apart from conv-deconv nets having recently led to automatic³ semantic segmentation, we end up with no further substantial advances being likely for yet another decade. Whether you believe this or not, it is clearly time for the status quo to be fully tested, and this means applying Deep and/or Very Deep nets to anything and everything, so that the real advantages and limitations can be worked out and fully documented. Interestingly, I have detected very little underlying theory in all this progress: such progress as has taken place has been achieved through inspiration, cunning, judicious use of available hardware (GPUs), and huge efforts to generate enormous datasets – or to simulate them by breaking down images into patches and by applying multiple transformations to them. I guess also, the fact that neural networks more or less died out in the 1990s meant that people were reluctant to employ them again – even when their time was due: this pent-up suppression eventually exploded on the scene in 2012, and

maybe we will soon return to it. Ideally, we would be working in a continuum in which progress is steady and fashions do not occur. However, the reality is that we operate in a situation where relaxation oscillations are inevitable (such as monkeys eating all the bananas, then the monkeys dying out leaving the bananas to grow again, and then a few monkeys returning, and so on). It only needs a time delay and a momentum term to keep such an oscillator going. Sounds as if we are part of a big neural network, in which our research is limited by time and momentum. One could even imagine we are being controlled not only by EPSRC but also by the Thought Police. Maybe Brexit will help ...

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

The Sullivan Thesis Prize 2016

This year’s Sullivan prize for the best thesis in a vision-related area has been won by Dr Xiatian Zhu of Queen Mary, University of London, with the thesis entitled “Semantic structure discovery in surveillance videos”. The reviewers were complimentary about the breadth and depth of the work, but also make a note about the clarity of the text and how it led the reader through the work. Recommended reading for students starting to write their own theses!

If you are interested in putting your own thesis forward for the 2017 prize, or the thesis of a student you have supervised or examined, please look at the relevant part of the BMVA website. The closing date is the end of March 2017.

Dr Adrian F Clark, M0PDF
BMVA Chairman
email: alien@essex.ac.uk

The BMVA Elections

Every year, half of the twelve elected members of the Executive Committee are available for election. Following the call for nominations in an earlier BMVA News, the following people were put forward for possible election:

- Marina Bloj
- Toby Breckon
- Adrian Clark
- Andrew Gilbert
- Xianghua Xie

As this is fewer than the six places available, all nominees are regarded as being duly elected.

Dr Adrian F Clark, M0PDF
BMVA Chairman
email: alien@essex.ac.uk

²In fact, going one stage further back, the word ‘shamble’ meant a bench such as that used by a butcher.

³Well, only if the right training is laboriously applied by human operators.

BMVA Distinguished Fellow 2016: Edwin Hancock



Every year, the BMVA awards a single Fellowship of the Association. The single criterion for the award of a Fellowship is that the recipient has made an outstanding contribution to computer vision.

It is fitting that, in a year when the British Machine Vision Conference is at York, the BMVA Distinguished Fellow should be Edwin Hancock. That is because Edwin is a Professor in the Computer Science department at York, where he has pursued a long and distinguished career.

Like many vision researchers in its early days, Edwin's first degree and PhD are in Physics — he will explain in his own words how he was bitten by the vision bug. Bitten he certainly was, because he has published around 700 papers, encompassing areas such as graph algorithms, shape from shading, texture, shape analysis, object recognition and machine learning. However, these are only part of Edwin's contribution. He has played a big part in journals and conferences, with a long-standing involvement in *Pattern Recognition* in particular. He has also been editor of the IET's *Computer Vision* journal.

For the BMVA, Edwin's most visible achievements have been chairing BMVC on its first visit to York in 1994, and then representing the BMVA on the Governing Board of the International Association for Pattern Recognition (the IAPR) for well over a decade. However, behind the scenes, Edwin has consistently provided sage and helpful advice to the BMVA's Executive Committee for about 20 years.

Edwin's contributions have already been recognised. He held a Royal Society Wolfson Research Merit Award during 2009–2014, he is a Fellow of the IAPR, the BCS, the IET, the Institute of Physics — and now he joins the select band of recipients of Fellowship of the BMVA.

Dr Adrian F Clark, M0PDF
BMVA Chairman
email: alien@essex.ac.uk

Distinguished Fellow Acceptance Speech

Thirty years ago to the day, I was on a train travelling to Bristol to present my very first vision paper at the Second Alvey Vision Conference.

About a year earlier I had switched fields from high energy physics to work with Josef Kittler. At the time it seemed a bit of a step in the dark, and perhaps even a step down from what seemed the more glamorous world of particle physics. I had worked on the first experiment to measure charmed particle lifetimes at SLAC (where the Physics Department had an infestation of Nobel Laureates), my PhD grandfather had discovered the pion (and got the Nobel Prize for it), the W and Z had been discovered 3 years earlier and the Higgs was expected soon (or so they thought). I was off to mix in different company for the first time. Expectations were not high. But I was nervous — this was to be my first conference presentation. In particle physics only the bosses got to give talks.

We, on the other hand, manned night shifts running the apparatus, spent long hours in darkened rooms visually scanning bubble chamber film, crawled inside powerful magnets (which were too delicate to power off) and wrestled with a pneumatic monster called the Hough-Powell (or huff-puff) device which temperamentally scanned for particle tracks on film.

My move from particle physics had been made possible by the Alvey Programme. This was a massive injection of cash into computing and AI, to revive UK research in these areas. The subject of AI had been all but killed off in the UK by the Lighthill Report, and Alvey was a desperate attempt to breathe life back into it. The field of computer vision had done well, with several well-funded long-term projects. The two largest of these were the so-called IKBS 3D Vision Project, which comprised GEC, Edinburgh, Sheffield, IBM and Sussex, and aimed to recover 3D shape from stereopsis and motion. The second was MMI007, which aimed to recognise 3D objects from 2D images, and comprised British Aerospace, Marconi, Bristol, Reading, RSRE and Surrey. The two projects were presenting their first results at the conference, and comprised a galaxy of talent including Andrew Blake, Hilary and Bernard Buxton, Bob Fisher, John Frisby, John Illingworth, Josef Kittler, John Mayhew, Dave Murray, Geoff Sullivan and Andrew Zisserman, to name just a few. Mike Brady had just returned from MIT, and was on the lookout for new talent to build a group at Oxford. The invited speakers were Shimon Ullman and Don Geman. Guy Scott won the best paper prize for his work on structure-from-motion, and later moved from Sussex to Oxford to join Mike, before embarking in a career in politics and becoming President of Zambia. This was primordial mass of talent from which the UK vision community was born, with its hallmarks of pragmatism and rigour.

After all, the conference was not so bad. They politely received my rushed and garbled talk, too full of equations on handwritten overheads. Old habits die hard. They were a convivial and irreverent lot who drank just as heavily as particle physicists and managed to look bright eyed and bushy tailed the next day. Many were also refugees from exciting but crowded areas of physics. I came across graphs and shape-from-shading for the first time. As they say, be careful what you allow into your head, it will find its way out one day.

I won't bore you with details of the intervening thirty years (it's extensively documented in the literature). Twenty-eight of them have been spent waiting for the Higgs to appear: twenty-five here in York. I cannot believe how productive those years have been. This is of course due to the patience, generosity of spirit and sheer braininess of my academic colleagues and research students. Not just that, but their willingness to embark on intellectual journeys into the unknown and the alacrity with which they engaged with whatever challenges confronted us.

This fellowship is an exceptional honour, and if it is not a contradiction, I do feel genuinely humbled, honoured and proud to receive it. It is of course a distinction not just for me, but for the colleagues, and in particular research students with whom I have worked over the years. Thank you very much indeed.

Professor Edwin Hancock
University of York
email: edwin.hancock@york.ac.uk

BMVC 2016 Statistics⁴

It is our great pleasure to welcome you to York for the 27th British Machine Vision Conference (BMVC). The University of York recently (2013) celebrated its 50th birthday and has grown rapidly since its founding. It is now home to more than 16,000 students. York is a campus University, sited in parkland and famous for its lakes and waterfowl. The conference is sited on the Heslington West campus on the outskirts of the historic city of York.

BMVC is one of the top events in the Computer Vision conference calendar and must now be considered as a truly international event with the majority of papers coming from outside the UK. This year, BMVC attracted a total of 365 valid submissions. Although this is lower than the record submissions level of recent years, it still represents a very active and healthy conference. The paper review process was unchanged from previous BMVCs, and we recruited more than 300 reviewers to process the papers. All papers received three reviews and each paper was then handled by two area chairs from our pool of 50 subject experts. Accepted papers required strong support from both reviewers and area chairs. We would like to thank all the reviewers and area chairs for their hard work and prompt responses.

Of the 365 submissions, just 144 were accepted for presentation in BMVC 2016, which is a 39% acceptance rate. Only the very highest quality papers were selected for oral presentation, with 38 papers gaining a podium spot, or 10% of the submissions. The accepted papers represent a truly international research community, with 18% of the papers from the UK, 36% from the EU excluding the UK, 22% from Asia, 20% from North America, and 4% from the rest of the world. As is now standard for BMVC, the proceedings are published entirely online, without the use of USB drives, for environmental reasons.

BMVC has always has strong links with industry, and again we are very grateful to our industrial sponsors for supporting the event. ARM, Disney Research, OSRAM, Ocado technology, HP, the IET, Edmund Optics and DigitalBridge kindly supported the main conference. Our thanks also go to NVIDIA, CRC press and Springer for sponsoring the best paper prizes.

We have put together an interesting programme of invited speakers and are delighted to welcome Dr Abhijeet Ghosh, Professor Katsushi Ikeuchi and Professor Raquel Urtasun to the conference. Dr Ghosh will deliver the tutorial on appearance modelling, and Professors Ikeuchi and Urtasun will give the two keynote presentations in the main conference.

BMVC 2016 has been organised by the Computer Vision and Pattern Recognition group in the Department of Computer Science at the University of York. The organisation of such a large conference would not be possible without the help of many people, and we are grateful to everyone who has contributed. Particular thanks must go to Bob French for his help in sorting out the logistical details and the student helpers for their support during the conference. We would also like to thank Xianghua Xie for his support and advice as the outgoing BMVC chair, and the BMVA committee for their extremely helpful suggestions and advice.

We hope you find BMVC 2016 in York both an enjoyable and a valuable experience.

Richard Wilson, Edwin Hancock,
Will Smith, Nick Pears, Adrian Bors
BMVC 2016 organising committee
email: william.smith@york.ac.uk

BMVC 2016 Prize-winners

Best Science Paper (sponsored by Springer and BMVA)

"Mapping auto-context to a deep, sparse convnet for semantic segmentation", *David Richmond, Dagmar Kainmueller, Michael Yang, Eugene Myers and Carsten Rother (University of Dresden, MPI Dresden)*.

Best Industry-related Paper (sponsored by NVIDIA and BMVA)

"EMVS: Event-based multi-view stereo", *Henri Rebecq, Guillermo Gallego and Davide Scaramuzza (University of Zurich)*.

Best Poster (sponsored by CRC press and BMVA)

"Better together: joint reasoning for non-rigid 3D reconstruction with specularities and shaping", *Chris Russell, Lourdes Agapito, Rui Yu, Andrew Fitzgibbon and Qi Liu-Yin (UCL, Microsoft)*.

BMVA Distinguished Fellow

Edwin Hancock (University of York)

Dr William A.P. Smith
University of York, UK
email: william.smith@york.ac.uk

⁴It seemed appropriate to include this article as is, even though it was written in advance to welcome delegates to the conference – Ed.

Around and About at BMVC 2016

Once again, BMVC was situated at York, in beautiful surroundings, both in the city and in the university itself – as you will see from the lake panorama below. This picture was taken from the door of the large area given over to conference poster presentations. As a result of man’s aquatic past, we are all calmed by water, even (or especially!) when we are subjected to as much new information as there is at a leading conference such as this. Thanks to Dr Guillermo Gallego (Zurich) for providing us with this lovely picture.



View from the poster presentation area.




Lectures in full swing in the Exhibition Centre’s comfortable, capacious theatre.



Invited speaker Raquel Urtaun (University of Toronto), who told us the thrilling story of how she and her team have moved us *scientifically* towards affordable self-driving cars.



Chairman Andrea Cavallaro (QMUL) and speaker Melih Kandemir (Heidelberg) listen attentively to a question from the floor.



Meetings between colleagues and old acquaintances were a valuable part of the poster sessions. Here, Huiyu Zhou (QUB) and Majid Mirmehdi (Bristol) meet Timothy Smith (Imagination Technologies), who is on the left.




Ruichi Yu (Maryland) responds to a question from the chair, Krystian Mikołajczyk (IC).



Poster sessions were interesting and well attended.




Adam Kortylewski (Basel) explains his work to another delegate.







Jiajun Wu (MIT) gives further insight into his poster.




Qi Liu-Yin and Rui Yu (both from UCL) stand by their prize-winning poster.




Hani Altwajry (Cornell) enthusing about his work.




Martin Šimonovsky (Paris Est) explains his work to another delegate.



Anurag Arnab (Oxford) gives further details about his poster.



Editor Roy Davies (RHUL) and Andrew Kay (invisibles.org) compare notes on summer school lectures, the newsletter, photography, cartoons and more.



Ryan Wolcott (Michigan) explains his poster to two visitors.



BMVA chair Adrian Clark (Essex) and Majid Mirmehdi (Bristol) relax for a moment.



General Chair Richard Wilson (York) double-checking that everything is proceeding according to plan.



Ancient and modern engines were among the fascinating exhibits in the National Railway Museum, York.





Some exhibits seemed positively antediluvian, yet it is salutary that rail travel arrived as recently as during Queen Victoria's reign.



Invited speaker Katsushi Ikeuchi (Microsoft Asia) and Edwin Hancock (York) look on.




The band was the centre of attention!



Farnoosh Heidarvinchek and Dima Damen (both from Bristol), Raquel Urtasun (Toronto) and Gabriel Brostow (UCL) circulating together.



The sight of all the brass was impressive – and the sound was superb!



Chris Russell (UCL), Oscar Mendez (Surrey) and Helen Bear (UEL).



Qi Liu-Yin (right) talks to his dinner companion.



Eraldo Ribeiro (Florida IoT) and Raquel Urtasun (Toronto) sharing a joke.



Nick Pears (York) and Adrian Clark (Essex).



Qi Liu-Yin, Chris Russell and Rui Yu accept the Best Poster Prize on behalf of themselves, Lourdes Agapito and Andrew Fitzgibbon.



Guillermo Gallego and Davide Scaramuzza (University of Zurich) receiving their Best Industry-related Paper on behalf of themselves and Henri Rebecq.



Carsten Rother (Dresden) receives the Best Science Paper prize from Edwin Hancock on behalf of himself, David Richmond, Dagmar Kainmueller, Michael Yang and Eugene Myers.



Krystian Mikolajczyk (IC) outlines his plans for next year's BMVC, which is to be held at Imperial College.



Edwin Hancock receives his *Distinguished Fellow* Award from Adrian Clark (BMVA Chair).




The marvellous atmosphere at the Banquet, with tables almost touching the exhibits.



After the Banquet, the number of guests gradually dwindles, though some find it hard to leave ...

Deep Learning in Computer Vision

Report on the BMVA Symposium held in London on 8 July 2016



Deep learning is an exciting driving force behind state-of-the-art image recognition systems, and is now fully penetrating the field of computer vision. The aim of this meet-up was to look further into the areas of computer vision that could be advanced by deep learning and how computer vision has advanced deep learning.

The event kicked off with a talk from Alison Lowndes – Deep Learning Solutions Architect at NVIDIA – about the hardware that makes deep learning possible – the GPU. Alison introduced three key concepts:

1. NVIDIA’s deep learning SDK, which allows integration of GPU’s with popular frameworks such as theano, caffe and others.
2. NVIDIA GPU, a chip designed to perform multiple tasks simultaneously, significantly speeding up computation, allowing us to train our networks within realistic time frames.
3. Providing a deployable product: GPU’s can and will be found in all devices: cars, phones and drones, so the step from bench to product is also important.

Deep learning not only requires processing power, but also large amounts of labelled data. Sander Dieleman – from Google Deepmind, explained how he won several Kaggle competitions classifying galaxies and plankton by using a series of innovative approaches to augment the datasets available. He also described the large ensemble of networks used to achieve the winning result, and further used test data as a regulariser – which was an interesting approach.

Unlabelled data is often easier and cheaper to acquire than labelled data. Soumith Chintala – Facebook researcher – introduced generative adversarial networks, which are able to learn concepts primarily from unlabelled data, and require very few labels to achieve good classification. This could allow deep learning to be applied to tasks that would otherwise not be possible due to lack of data.

It can take 20 years to design a new text font, but deep generative models can learn a design space that can be easily explored by anyone, making design accessible to a wider user base. Samim Winiger, founder of samim.io showed us how generative models, specifically generative adversarial networks, can be used to learn to interpolate between images in high dimensions to create a smooth, explorable design space and find new designs. The design spaces described were not limited to fonts, but extended to music and patterns. This still left the question of how to

visualise these high-dimensional spaces in only 2 or 3 dimensions.

One of the missions of Twitter Cortex is to make sense of content, be it a photo, a tweet or a video. Content on Twitter is ever evolving, new concepts and ideas are born, and as a consequence the Taxonomy required to make sense of content cannot be static. Nicolas Koumchatzky, described the challenges and approaches involved in developing a taxonomy for understanding photographic content on twitter and some of the software contributions twitter has made to mitigate these challenges.

Deep learning can seem like a “black box”: visualising what’s going on in this “black box” can be highly insightful. Andrea Vedaldi from Oxford University, demonstrated visualisations that maximally activate neurons in deep networks, identifying concepts that networks learn. Visualisations reveal key features captured by the networks as well as when the network learns to be invariant. These visualisation techniques can be a powerful tool for understanding what a network is learning, and potentially how to improve learning.

Some visualisations show that trained CNNs learn a set of hierarchical features. Nick Kingsbury drew parallels between trained CNNs and his work on multi-layer filter banks – viz., using banks of pre-defined, lossless wavelet transforms at different scales throughout the network, to describe images.

Ben Graham from the University of Warwick turns the curse of dimensionality on its head, by taking advantage of datasets that contain spatially sparse examples, where the greater numbers of dimensions increase the likelihood of a sparse sample. Ben explains how a sparse image can be more efficiently processed by only considering non-zero parts of the data, with applications in chemistry, biochemistry and robotics.

In the poster competition sponsored by NVIDIA and Cortextica, Wenzhe Shi from Twitter’s recently acquired deep learning start up, Magic Pony Technology, won first prize for his poster “Real-time single image and video super-resolution using an efficient sub-pixel convolutional neural network”. Second prize went to Adeline Paiement from University of Bristol for her poster “Skeleton-free body pose estimation from depth images for movement analysis”.

Thanks to Kai Arulkumaran and Anil Bharath for chairing this event and inviting excellent speakers. Thanks also to Andrew Gilbert, the meeting organizer.

Antonia Creswell
Imperial College London
email: ac2211@ic.ac.uk

Plants in Computer Vision



One-day BMVA symposium in London, 16 November 2016
Chair: Hannah Dee

www.bmva.org/meetings

Keynote Speakers: Sotirios Tsafaris, Chancellor's Fellow, University of Edinburgh and Hanno Schar, Head of Quantitative Image Analysis, Forschungszentrum Jülich, GermanyImage

Programme

New methods in plant biology have led to an explosion of data types and methods of data acquisition, and much of this data is image-based in nature. Computer vision, image analysis and image processing techniques are being applied to more plant data than ever before.

- 9:00 Registration opens
- 9:30 Keynote 1: Hanno Schar – Forschungszentrum Jülich, Germany
- 10:30 Jonathon Gibbs – A fully automated active vision cell for 3D reconstruction of plant shoots
- 10:50 Coffee Break and Posters
- 11:50 Gytis Bernotas – 3D plant surface reconstruction using the Photometric Stereo technique
- 12:00 Dominic Williams – Segmentation and splitting of hyperspectral images of raspberries in field conditions
- 12:20 Norman MacLeod – Automating leaf physiognomic character identification for taxonomic, phylogenetic and climate change research
- 12:40 Lunch and Posters
- 13:20 keynote 2: Sotirios Tsafaris – Edinburgh University
- 14:20 Jonathan Bell – Detecting leaf-from-leaf occluding edges in Arabidopsis thaliana
- 14:40 Andrew French – New analysis and traits for automated phenotyping
- 15:00 Christian Fournier – Phenomenal: a software framework for model-assisted analysis of high throughput plant phenotyping data
- 15:20 Rick van de Zedde – Plant architecture measurement using a 3D reconstruction method suitable for high throughput plant phenotyping
- 15:40 Coffee Break
- 16:00 Marco Aita – Circumnutation in 3D
- 16:20 Milan Sulc – Fine-grained recognition of plants in the wild
- 16:40 Shafiekhani, A – Vinobot and Vinocular: two robotic platforms for field phenotyping.


Posters

- A Raspberry Pi based solution for high quality automated plant imaging using RFID technology – Katie Awty-Carroll, Colin Sauze, Kevin Williams
- Developing continuous Gravimetric systems for use in plant Phenotyping – Nathan Hughes
- Watering plants makes them grow a bit later – K.S. Williams, R.D. Boyle
- Producing and analysing a seed germination time series dataset – Danny Awty-Carroll, Paul Robson
- Flower identification on species level with uneven classes and few images – Vasiliki Simaiaki, Andrea Mirabile, Eduard Vazquez
- Visibility in foliage – Miles Hansard
- Root system recovery from X-ray CT image data of fully-developed wheat plants – Stefan Mairhofer, Brian Atkinson, Craig J. Sturrock, Sacha J. Mooney, Malcolm J. Bennett, Tony P. Pridmore.

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

Transfer Learning in Computer Vision

BMVA Symposium
Wednesday 25 January 2017, London
Chairs: Dr Dima Damen (Bristol), Dr Gabriela Csurka (Xerox) and Dr Timothy Hospedales (QMUL)



Keynote Speakers

- Professor Massimiliano Pontil (Istituto Italiano di Tecnologia and UCL – Transfer Learning (TL) Methodologies
- Professor Peter Hall (Bath) – TL across visual domains: from images to sketches.

Invited Speakers

- Dr Michel Valstar (University of Nottingham) – TL for facial action unit detection
- Dr Teo de Campos (University of Surrey) – Transductive TL for activity recognition
- Dr Adrien Gaidon (XRCE) – TL from gaming environments to real images.

Call for Papers

The Computer Vision community is in need of moving beyond dataset or task-specific methods towards those that can efficiently adapt to new tasks or domains in a supervised, semi-supervised or unsupervised manner. We aim in this technical meeting to bring together leading researchers, at various levels in their career, with expertise or strong interest in TL for Computer Vision problems, in order to discuss current challenges and propose future directions including potentially establishing a continuous forum or a workshop series.

We are inviting researchers to present short talks and posters that address the motivation, methodologies, challenges and applications of using TL in Computer Vision. These could include ongoing or recently published works, relevant but not limited to:

- TL for a specific computer vision task (e.g., object/event/activity detection, recognition, retrieval, tracking)
- Cross-Dataset or Cross-Annotations TL
- Cross-Domain TL
- Cross-Task TL
- Cross-Modality or Cross-Sensor TL
- Unsupervised and Incremental TL
- Datasets, challenges and protocols for evaluating TL methods.

Call for Participation

All those interested in presenting at this meeting are invited to submit a summary of their talk by 1 November 2016 (firm deadline): <http://goo.gl/forms/AkABYVIRIH0qxen42>

Registration

Book online at www.bmva.org/meetings: £10 for BMVA Members, £30 for non-Members, including lunch.

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

Report on BMVA Computer Vision Summer School 2016

The BMVA Computer Vision Summer School is probably the longest running annual event in the field. Its 21st edition was held on 4–8 July in Swansea University's new £450M Bay Campus. The summer school has successfully become a self-sustained event and still offers substantial discount to students and early career researchers from UK institutions, e.g., UK non-residential early rate of £260 versus £490 for non-UK. A total of 62 delegates from 13 different countries, with 15 attendees from outside UK, attended this year's summer school. The group photo below was taken in front of the iconic Great Hall main building towering over the new Swansea University Bay Campus designed by the world renowned architect Dr Demetri Porphyrios.



The attendees had a busy week of lectures on a range of current topics in computer vision. A total of 15 speakers from academia and industry delivered 17 ninety-minute lectures and 2 lab sessions. The poster session also proved to be popular. The Best Poster prize went to B. Gatto of Federal University of Amazonas for "A deep network model based on subspaces". The recipient was awarded £100 and Professor Roy Davies's book *Computer and Machine Vision: Theory, Algorithms, Practicalities*. Two runner-up prizes (£50 each) were awarded to M. Wray and D. Moltisanti of Bristol University for "SEMBED: semantic embedding of egocentric videos" and R. Busatto of University of East Anglia for "Visual navigation system for people with sight loss". The attendees also enjoyed the summer school banquet at the Meridian Tower (the tallest building in Wales), overlooking the majestic Swansea Bay.



B. Gatto



R. Busatto



M. Wray and D. Moltisanti

On behalf of the organisers, I would like to thank the speakers, attendees, and volunteers. Swansea has been glad to host this very successful event for the past three years; it is now time to pass the baton: please keep an eye open for announcements about the 22nd edition in 2017!

Dr Rita Borgo
Swansea University
email: r.borgo@swansea.ac.uk