BMVA News

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

Volume 28 Number 1
September 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 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 December 2017.

Contents

Editorial: *Have Most of Us Missed the Bus?* 1
BMVA Distinguished Fellow 2017: Andrew Fitzgibbon 2
Around and About at BMVC 2017 3
BMVC 2017 Prize-winners 8
BMVA Executive Committee Elections 8
BMVA Symposium: Human Activity Recognition and Monitoring 9
Upcoming BMVA Meetings 9
BMVA Computer Vision Summer School 9
Conference on Automatic Face and Gesture Recognition 10
IJCV Special Issue on Deep Learning for Face Analysis 11
AI for Better Patient Outcomes 12
BMVC 2018 to be held at Northumbria University 13

Editorial: *Have Most of Us Missed the Bus?*

In the not too distant past, city buses were designed for convenience, and safety was a lesser consideration than it is today. In fact, it was possible to run for an accelerating bus and jump onto its rear platform with little risk of falling off, as long as one grabbed the convenient vertical pole seemingly provided for the purpose. I was reminded of this by the accelerating pace of developments in computer vision – particularly in relation to deep learning. However, in that area the rate of progress has now become so high that it will be difficult for many vision workers to catch up and cope with the flow. Indeed, it may be that those trying to join in are already too late to act as more than interested observers if they haven't taken part from The Beginning. 2012 was the year that the crucial developments in deep learning were taking place, marked in particular by Krizhevsky et al.'s paper on AlexNet. And in 2015 Simonyan and Zisserman's VGGNet architecture set the scene for much that was to follow.

Interestingly, I have to report that the rate of progress has not subsided at all. One has only to look at the papers just published in BMVC 2017 to confirm this. First, the number of papers in the conference proper, with 'deep' in the title, was some 22, and if one excludes poster papers it was 3, a number that is bolstered up to 5 if 'CNN' and 'convolutional' are included. Actually, these numbers may not sound hugely impressive, but if one looks at the lists of references in such papers, the picture becomes radically different. For example, Mordan et al.'s Best Science Paper contains 44 references, mostly about deep learning: what it more, the dates of these papers range from 1989 to 2017, though the actual distribution is 1989: 1, 2010: 2, 2011:1, 2012: 2, 2013: 1, 2014: 4, **2015: 14, 2016: 14**, 2017: 5. The 1989 paper is by LeCun et al., and can be regarded as the one that set the whole applecart going (but without the 'deep' epithet being immediately applied to it). Clearly, the distribution of these papers over time has a standard deviation of only a year or so, and there is no embarrassment from not including a representative set of key papers from previous decades. Long gone are the references to Azriel Rosenfeld's papers (few may recall that he was the undisputed 'Father of the Field' for at least 50 years): and now it is nearly respectable to include papers from just the past two or three years. However, Mordan et al.'s paper does deservedly cite Felzenszwalb et al (2010) and Ott and Everingham (2011). Nevertheless, as an editor for Pattern Recognition Letters, I shall now have to change my tune and not insist that authors pay lip-service to older key papers. Why bother? – they belong to an older era, and are interesting but no longer relevant. Another signal is that, with only two citations in the Mordan paper, PAMI may be approaching its 'sell-by' date. A further signal is the increasing use of ArXiv references, five in the BMVC deep

[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.

learning paper by Dong et al., but now widespread in this area, and again (in this paper) at the expense of PAMI.

Another thing that has bemused me recently is the general difficulty of getting people to review papers. The problem is easily stated: if people are any good they will be busy researching and writing papers and will not be wasting time reviewing them (or else they will be on a long time-fuse, which in practice amounts to the same thing); whereas young people who have not had much experience of writing papers will not be any good at reviewing them. [I emphasise that these are cynical, worst case paradigms and I am not trying to insult potential reviewers.] This puts editors in a difficult position. Furthermore, it's difficult to see how inexperienced people can ever *learn* to be good at reviewing: it's also difficult to see how penetrating judgements can be extracted from experienced people. Fortunately, it's usually possible to give people in the middle ground the chance to make useful comments, suggestions and criticisms about any paper: then it's up to the editor to put all the bits together and make global sense of them. The remaining difficulty is that the stamp collecting process can take time, though that is a sign of the times we are in. And so it will remain, while promotions and success are judged largely (in academia at least) by numbers of publications. Interestingly, I've seen a glimmer of light recently, in that papers on deep learning seem to be quite attractive for reviewers. Fewer potential reviewers need to be asked and those that are produce their reports more quickly. Evidently, potential reviewers are presently keen to get up-to-date about the latest deep learning work and to apply it themselves. Indeed, with the date relevance window that I hinted at earlier being something like 2 years, they simply have to take every opportunity to find out what's going on *before* it appears in journals and conferences; and papers for review are clearly prepublished material. Of course, the danger of learning in this way is that they have to make up their own minds how reliable the information they are obtaining is: fortunately, that will make them more conscientious and reliable reviewers, which is all to the good!

All that I have said so far points to deep learning being 'a good thing', and, as I have indicated in previous editorials, people are anxious to get on this new bandwagon: indeed, they have to be, as missing it may mean they never catch up. The unfortunate aspect of this is that we have 'chips with everything', or worse, we catch the British disease of having gravy or custard with everything (I often feel that gravy and custard were invented by the British to take away the poor quality of their meals). On the other hand, with the all-too-frequent appearance of Masterchef and similar programmes on TV over the past decade or two, the English have raised their palates and cooking capabilities to quite a high level, far supplanting the meals of 50 years ago. Perhaps we need to educate our deep learning and vision workers in the same sort of way: then the chips, gravy and custard syndrome will die a natural death. The trouble is that, if it takes a decade or two to achieve this, we will first have to know much more accurately what behaviour we are actually aiming to eliminate – and who can reliably predict that?

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

BMVA Distinguished Fellow 2017: Andrew Fitzgibbon

You often hear people being introduced on TV with the phrase that they "need no introduction" – and this is one of the few cases in the vision research community where this is genuinely the case.

Our Distinguished Fellow for 2017 studied first at University College Cork, then moved to Edinburgh where he contributed to, among other things, HIPR, the first web-based vision teaching aid. He then received a Royal Society Research Fellowship which took him to Oxford, where he most notably contributed to the *Boujou* system, which won an Emmy award. If you have seen a recent movie that has real actors and CGI, you have probably seen the fruits of this work.

In 2005, our Distinguished Fellow moved to Microsoft Research, where he was a major contributor to the development of the Kinect, which quickly became an important vision capture device. He has recently been working on the Hololens, which shows every sign as being as ground-breaking and important.

However, our Distinguished Fellow has not worked purely on development: he has amassed at least ten best paper awards at major conferences, and also the silver medal of the Royal Academy of Engineering and the BCS Roger Needham award. He is a past chair of the BMVA Executive Committee, and it was at his instigation that BMVC grew from a good national conference into the important international one that it is today. He is a Fellow of the Royal Academy of Engineering, of the British Computer Society and of the International Association for Pattern Recognition.

Our BMVA Distinguished Fellow for 2017 is Andrew Fitzgibbon.

Dr Adrian F Clark, M0PDF
BMVA Chairman
email: chair@bmva.org

Around and About at BMVC 2017

Andreas Geiger presenting his tutorial.

Lourdes Agapito giving her tutorial.

Chair T-K Kim taking questions for Lourdes Agapito.

Richard Szeliski giving his keynote lecture.

Taylor Mordan presenting his Best Science Award Paper.

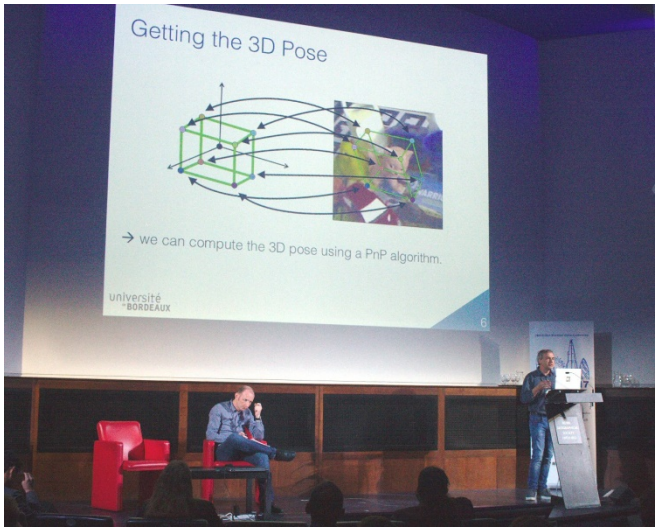

Andrew Davison chairing Vincent Lepetit's tutorial.

Kristian Mikolajczyk, Andrew Davison, Roberto Cipolla, Pietro Perona, Christopher Zach and others in Jamie Shotton's workshop keynote.

Andrew Zisserman and Pietro Perona in discussion after a lecture.

Abhijeet Ghosh demonstrating Light Stage to Andrew Zisserman, Roberto Cipolla and Stefanos Zafeiriou.

On the Snapchat exhibition stand: Maria Sopo, Donna Denman, Rachel Greenfield, Qi Pan, Jennie Mandal, David McCabe, and Ed Rosten.

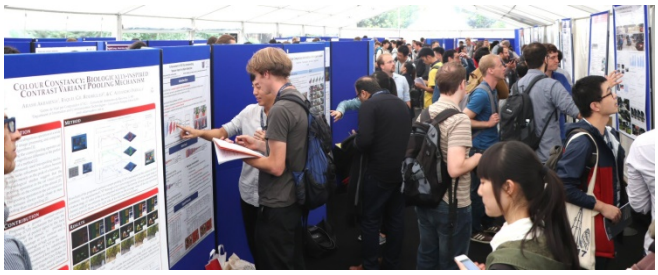

Activity during a poster session.

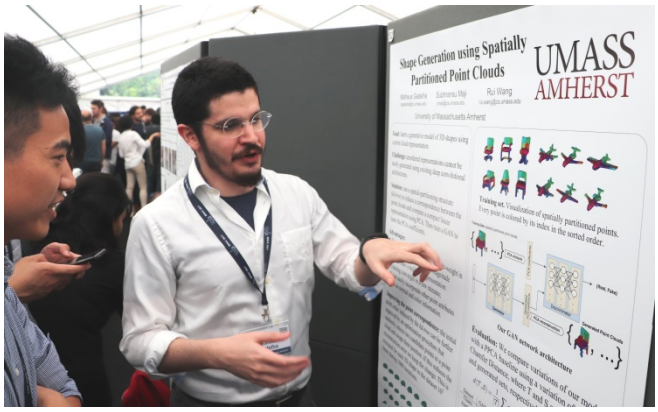

Matheus Gadelha presenting his poster.

Stephan Liwicki (Toshiba, front left) presenting his poster.

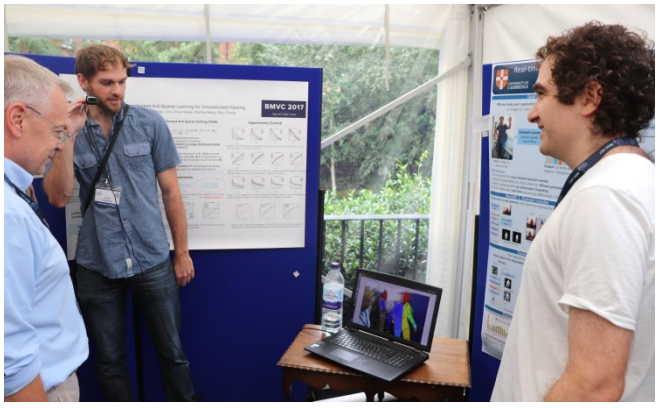

James Charles showing a live demo to Richard Szeliski and Edward Rosten as part of his poster presentation.

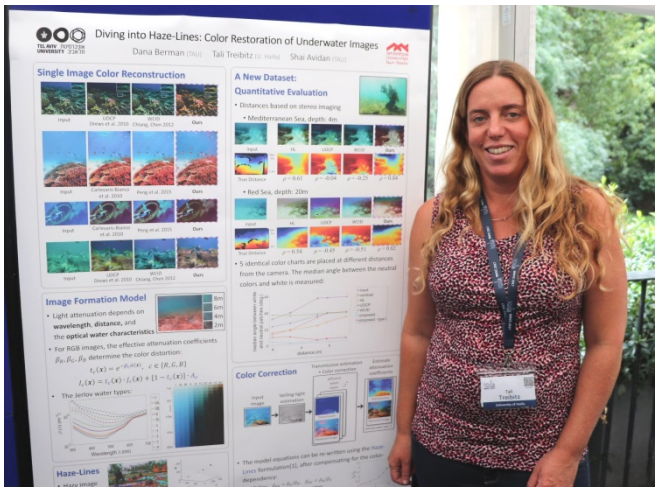

Tali Treibitz presenting her poster.

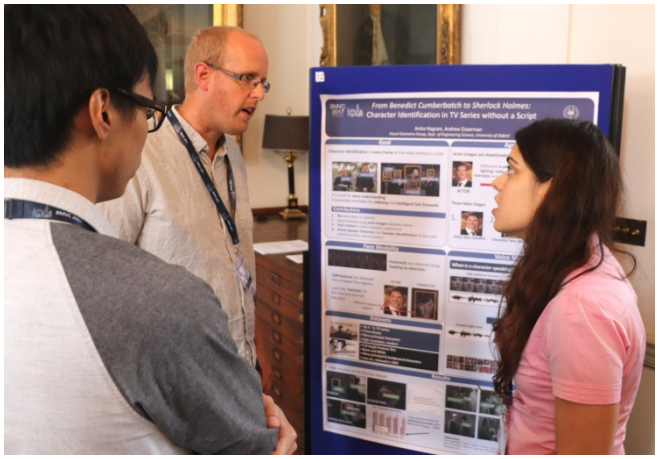

Arsha Nagrani in discussion with Andrew Gilbert about her poster.

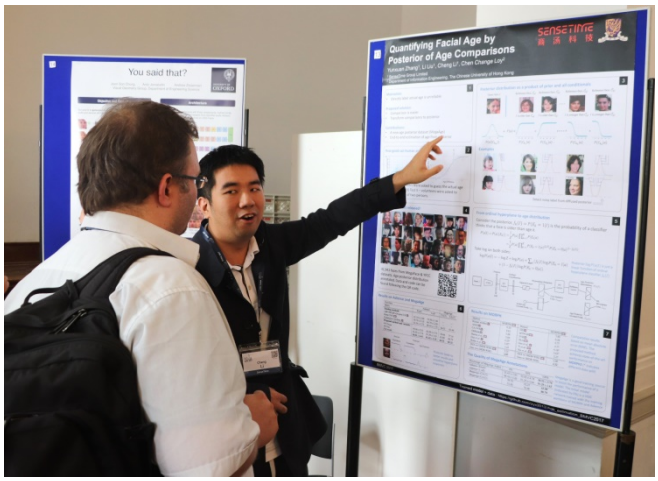

Cheng Li emphasises some of the work described in his poster.

Stefan Leutenegger (front right) listening to a poster talk.

Roey Mechrez (right) in front of his poster.

Joon Son Chung discussing his poster.

Mengjiao Wang, Arun Mukundan, Giorgos Tolias, Ondrej Chum, Stephan Liwicki and others at the conference banquet.

Hui Yu (front left) at a banquet table.

Akihiro Sugimoto and others awaiting the first course!

Tom Drummond (waving to the camera), with Ignas Budvytis, Olivia Wiles and others.

(right to left) Rigas Kouskouridas, Edward Miller, Vassileios Balntas, Chrysa, Guillermo Garcia-Hernando, Axel Barroso and Adrian Rodriguez.

Best Poster Award (Honourable Mention): Yuming Shen, flanked by organisers T-K Kim, Stefanos Zafeiriou; Gabriel Brostow and Krystian Mikolajczyk.[2]

Best Poster Award: Olivia Wiles, with Robert Deaves (Dyson) as the award sponsor.

Best Industry Paper Award (Honourable Mention): Bailey Kong.

[2]Authors not shown in these photographs are listed in the section BMVC 2017 Prize-winners on p. 8.

Best Industry Paper Award: Daniel Hernandez-Juarez.

Best Science Paper Award (Honourable Mention): Arun Mukundan and Giorgos Tolias.

Best Science Paper Award – Taylor Mordan, with Qi Pan and Edward Rosten (Snap, London) as the award sponsors.

Most Dedicated Organisers Award: Guillermo Garcia-Hernando and Athanasios Papaioannou.

BMVA Sullivan Best Thesis Prize: BMVA Chair Adrian Clark presenting the award to Karel Lebeda (University of Surrey).

BMVA Distinguished Fellow 2017: Andrew Fitzgibbon receiving his award from Adrian Clark.

Andrew Fitzgibbon, BMVA Distinguished Fellow 2017, making his acceptance speech.

May I take this opportunity to thank Tae-Kyun Kim and his team for providing such a wealth of photographs making it easy for me to compile this gallery of events from BMVC.

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

BMVC 2017 Prize-winners

Best Poster

"SilNet: Single- and Multi-View Reconstruction by Learning from Silhouettes", *Olivia Wiles (Oxford University) and Andrew Zisserman.*

Best Poster (Honourable Mention)

"Unsupervised Deep Generative Hashing", *Yuming Shen (University of East Anglia), Li Liu and Ling Shao.*

Best Industry Paper

"Slanted Stixels: Representing San Francisco's Steepest Streets", *Daniel Hernandez-Juarez (UAB, Spain), Lukas Schneider, Antonio Espinosa, Juan Moure, David Vazquez, Antonio López, Uwe Franke and Marc Pollefeys.*

Best Industry Paper (Honourable Mention)

"Fine-Grained Forensic Matching", *Bailey Kong (UC Irvine), James Supancic, Deva Ramanan and Charless Fowlkes.*

Best Science Paper

"Deformable Part-based Fully Convolutional Network for Object Detection", *Taylor Mordan (LIP6, France), Nicolas Thome, Gilles Henaff and Matthieu Cord.*

Best Science Paper (Honourable Mention)

"Multiple-Kernel Local-Patch Descriptor", *Arun Mukundan (CMP, Czech Republic), Giorgos Tolias and Ondrej Chum.*

Most Dedicated Organisers

Guillermo Garcia-Hernando and Athanasios Papaioannou (both at Imperial College London).

BMVA Sullivan Best Thesis Prize

"2D and 3D Tracking and Modelling", *Karel Lebeda (University of Surrey).*

BMVA Distinguished Fellow

Andrew Fitzgibbon (Microsoft, UK).

Dr Tae-Kyun Kim
Imperial College London
email: tk.kim@imperial.ac.uk

BMVA Executive Committee Elections

The BMVA is run by an Executive Committee of twelve elected members and several co-optees. Elected members serve on the ExCo for two years, and six members of the ExCo are elected every year. The idea behind this is that the "new blood" members bring new ideas and stop the ExCo from becoming set in its ways.

The response to a request for nominations in an earlier issue of *BMVA News* was disappointing, with only one nomination received and that one was for re-election. I have asked the other ExCo members who served during 2015–17 to stay on for another term and they have kindly agreed; those people are:

- Neill Campbell
- Roy Davies
- Majid Mirmehdi

This allows the BMVA to continue its current activities unabated but makes it more difficult for us to work on new initiatives.

We are able to co-opt new members onto the ExCo, so if you are working in the vision domain and would like to be involved, please drop me a line. It doesn't matter whether you are working in academe or industry and whether you are academic or research staff or a student – the ExCo already includes people from all these backgrounds and we welcome more. If you have a burning new idea that you think would benefit vision research in the UK and want to bring it to reality, do contact us. It is worth pointing out that involvement in the BMVA is good for your CV!

Dr Adrian F Clark, M0PDF
BMVA Chairman
email: chair@bmva.org

BMVA Symposium: Human Activity Recognition and Monitoring

Chairs: Ardhendu Behera (Edge Hill University), Nicola Bellotto (University of Lincoln) and Charith Abhayaratne (University of Sheffield)

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).

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. Therefore, human activity analysis and recognition has become a research area of great interest, particularly in view of 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 a 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 latest approaches in this domain. We hope the meeting will stimulate future research, both theoretical and practical perspectives to stimulate further advances in the field.

Registration

Book online at www.bmva.org/meetings: £16 for BMVA Members, £36 for Non-Members, including lunch and subscription to the BMVA for non-members.

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

Upcoming BMVA Meetings

We have 5 exciting upcoming meetings in 2017 and 2018, as shown in table below: go to bmva.org/meetings to register or submit work to these meetings.

11 Oct	Computer Vision and Modelling in Cancer: Greg Slabaugh and Constantino Carlos Reyes-Aldasoro
8 Nov	Human Activity Recognition and Monitoring: Ardhendu Behera, Nicola Bellotto and Charith Abhayaratne
7 Feb	Autonomous Perception: Serge Thill and Maria Riveiro
6 June	Computer Vision for Smart Environments and assisted living: Francesca Odone
7 July	Robots SLAM: Oscar Mendez-Maldonado and Mihai Bujanca.

We are now looking for volunteers for two meetings in late 2018: please contact me if you are interested in finding out more.

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

BMVA Computer Vision Summer School

School of Computer Science, University of Lincoln, 3–7 July 2017

The 22nd BMVA Summer School was hosted by the University of Lincoln on 3–7 July 2017. The event ran smoothly, both in terms of lectures and general organization, including accommodation and social events. Further details are given below.

Attendance

There were 65 registered attendees: 45 were from the UK and 20 from outside the UK; 10 were from the University of Lincoln and 2 were from industry. Most were PhD students, and there was 1 Professor. More information about the student profiles will be available after feedback forms have been completed and analysed.

Lectures

16 lectures were given at the summer school.
Co-organiser lectures: 2 (1 from UoL and 1 from UEA).
Gender split: 2 female, 14 male.
Geography: 15 from UK institutions, 1 from Switzerland (University of Zurich).
Industry representation: 2 (Dyson and Microsoft), 14 from academia.
New speakers: 5 first time lecturers.
One lecture was delivered as a hands-on tutorial demonstrating the use of Matlab for computer vision.
The programme and the lecture notes are available on http://cvss.blogs.lincoln.ac.uk/programme/

A lecture in progress

A Matlab session in progress

Poster session

25 posters were displayed.
1 best poster and 2 finalists were awarded on the last day:

- Ghalia Hemrit (best poster), University of East Anglia, "A Curious Problem with the Colour Checker Dataset"
- Bruna Pearson, Durham University, "Real-time Computer Vision – Using Visual Saliency for Monocular Guidance in Mobile Robots"
- Roberto Dyke, Cardiff University, "Non-rigid Registration for Large-scale Deformations of 3D Meshes".

Dinner at the Wig and Mitre restaurant

The Summer School attendees

General organization

The total number of University accommodations used for the event was 55. There were 8 coffee breaks, 4 lunch breaks, 2 sessions (the opening and poster sessions) with refreshments, social events including ten-pin bowling, dinner in the Wig and Mitre restaurant, and tasting of local beers and selected delicacies.

Nicola Bellotto, Tryphon Lambrou, Michael Mangan (all at Univ. of Lincoln) and Michal Mackiewicz (UEA)
email: nbellotto@lincoln.ac.uk

Conference on Automatic Face and Gesture Recognition

A view of Washington DC

The 12th IEEE International Conference on Automatic Face and Gesture Recognition (FG 2017) was held in Washington DC, US on 30 May–3 June. It is one of the most important conferences on Face and Gesture Recognition. It covers a broad range of topics, such as face recognition, facial expression analysis, and gesture and action recognition. The conference presents research that advances the state-of-the-art in these and related areas, leading to new capabilities in various application domains. FG 2017 received 172 submissions, 92 papers were accepted: 17 submissions were selected for presentation in the Oral track and 75 for presentation in the Poster track. The conference lasted for 5 days with a main technical program, various workshops, tutorials, challenges and technical demos.

Talks during lunch time

On the first day of the conference, there were four tutorials and the Doctoral Consortium, which I mainly attended. Doctoral Consortium aims to provide a forum for PhD students to discuss their research and career objectives with the senior researchers. 15 PhD students including me were selected to attend and each student was allocated with

a mentor working in a similar area. Doctoral Consortium provided a networking environment that enabled the doctoral students to establish new contacts and collaborations with other researchers and to equip the new generation of bright researchers with career options in academia and industry.

Award session during the banquet

The main conference lasted for 3 days. Each day started with a keynote in a different research area. On the first day, Margaret Livingstone, Takeda Professor of Neurobiology at Harvard University, presented latest findings in the organization of visual processes in temporal cortex. After the keynote there were three oral sessions and 8 papers were presented on face recognition and facial expression analysis. After that a poster highlight session with 30 papers was held and then the posters are presented. 5 demos were also exhibited in the afternoon. In addition, there was a special session on *Remote Physiological Measurement from the Face and Body*.

Panel session of FG 2015

The second day of the main conference started with a keynote given by Christopher Boehnen, Senior Program Manager at IARPA. He presented on the vital role of face recognition in presentation attack detection and unconstrained face verification and recognition. After the keynote, 11 papers on gesture recognition, facial landmarks detection, etc., were presented. Next, a poster highlight session was held in the afternoon. The second day of the main conference ended with a banquet and award ceremony. The last day of the main conference was mainly about behavioural analysis and emotions, which is the area I work on and I also presented my paper on that day. The keynote was given by Stan Sclaroff, Associate Dean and Professor of Computer Science at Boston University. He presented on deep models for video-based gesture and activity recognition and understanding. In addition to the oral and poster sessions, there was a panel discussion session "*Where is the 'social' in face and gesture research? From individual behaviour to dyadic, multi-party and crowd analysis*". Development of the analysis of facial expression, emotions, etc., was discussed and interesting topics were suggested.

Presenting my paper

Four workshops were held on the last day of the conference. I concentrated on the *Facial Expression Recognition and Analysis Challenge* (FERA 2017), which was more directly related to my PhD research. This challenges includes two sub-challenges (1) Occurrence Sub-Challenge to detect AU occurrence in 9 different facial views and (2) Fully Automatic Intensity Sub-Challenge to detect both the occurrence and intensity of AUs in 9 different facial views.

This year's Conference on Automatic Face and Gesture Recognition was very well organized in terms of both technical program and social events. More information about FG 2017 can be found on the conference website: http://www.fg2017.org/. Finally, I would like to thank the BMVA for providing the generous support that enabled me to attend this conference.

Wenxuan Mou
QMUL
email: w.mou@qmul.ac.uk

IJCV Special Issue on Deep Learning for Face Analysis

Scope

Deep learning is one of the most important breakthroughs in the field of artificial intelligence over the last decade. It has achieved great success in speech recognition, natural language processing, computer vision, and multimedia. Many face analysis tasks, including face detection, alignment, reconstruction and recognition, benefit from the powerful representation learning capability of deep learning techniques. Not only has there been a constantly growing flow of related research papers, but also substantial progress has been achieved in real-world applications such as

security, video surveillance and human–computer interaction.

While substantial progress has been achieved in face analysis with deep learning, many issues still remain and new problems emerge. For instance, the scalability of deep networks to large-scale unconstrained recognition needs to be improved. 'In-the-wild' facial attribute recognition with imbalance class distribution is still challenging. The accuracy and efficiency of detecting faces with a wide range of scales in crowded scenes still see considerable room for improvement.

This special issue presents a great platform to make a definitive statement about the state of the art by providing a significant collective contribution to this emerging field of study. Specifically, we aim to solicit original contributions that: (1) present state-of-the-art theories related to deep learning for face analysis; (2) develop novel methods and applications; (3) survey the recent progress in this area; and (4) establish benchmark datasets.

Topics

The list of possible topics includes, but is not limited to:

Theory

- Deep learning
- Cross-domain feature learning and fusion
- Transfer learning
- Multitask learning
- Generative adversarial learning
- Multi-instance learning
- Weakly supervised learning
- Reinforcement learning
- Zero-shot / One-shot learning.

Applications

- Face detection
- Face alignment and tracking
- Face recognition
- Face verification
- Face clustering
- Face attribute recognition (including age and gender)
- Facial expression recognition
- Face hallucination and completion
- 3D face reconstruction
- Face parsing
- Face sketch synthesis and recognition.

Paper Submission

Authors are encouraged to submit original work that has not appeared in, nor is in consideration by, other journals. Previously published conference papers can be submitted in extended form (with additional supporting experiments and a more detailed technical description of the method). Manuscripts will be subject to a peer reviewing process and must conform to the author guidelines available on the IJCV website at *Instructions for Authors*.

Manuscripts can be submitted to: http://visi.edmgr.com by selecting *S.I.: Deep Learning for Face Analysis* in the section *Choose Article Type*.

Important Dates

Submission deadline:	15 Jan 2018
First review decision:	15 April 2018
Revision deadline:	30 June 2018
Final review decision:	1 August 2018
Final manuscript submission:	1 Sept 2018
Online publication:	October 2018.

Guest Editors

Rama Chellappa, University of Maryland, USA
Xiaoming Liu, Michigan State University, USA
Tae-Kyun Kim, Imperial College London
Fernando De la Torre, Facebook, USA
Chen Change Loy, The Chinese University of Hong Kong, HK SAR, China.

Dr Tae-Kyun Kim
Imperial College London
email: tk.kim@imperial.ac.uk

AI for Better Patient Outcomes

This IET Symposium will take place on Wednesday 18 October 2017 at Microsoft Research Lab, Cambridge, UK.

Register at: www.theiet.org/aihealthcare

Programme

08:30 Registration and refreshments.

09:30 Chair's welcome and introduction
Dr Kenji Takeda, Director, Azure for Research, Microsoft.

09:35 Keynote address: InnerEye – Assistive AI for Cancer Treatment

- AI is being used to augment human ingenuity and provide much needed productivity improvements in medical imaging
- InnerEye helps radiation oncologists deliver better care, more efficiently and consistently to their cancer patients
- It builds upon many years of deep research at Microsoft in computer vision and machine learning, as used in Kinect and Hololens.

Dr Raj Jena, Scientific Officer, Honorary Consultant Clinical Oncologist, Cambridge University Hospitals NHS Foundation Trust.

10:05 Predicting health outcomes for smarter healthcare systems

- Learning health systems are integrated healthcare system which harness the power of data and analytics to learn from routinely collected health data
- Predictive analytics play a key role in producing the intelligence for learning health systems
- This talk with focus on opportunities and challenges in outcomes prediction from routine data.

Niels Peek, Director, Greater Manchester Connected Health City and Professor of Health Informatics, University of Manchester.

10:35 Machine learning for the next generation of healthcare interventions
- Dynamically forecasting patient health in critical care
- Phenotyping complex diseases
- Understanding antibiotic resistance from the tuberculosis genome.

Dr Tingting Zhu, Junior Research Fellow, Computational Health Informatics (CHI) Laboratory, University of Oxford.

11:05 Refreshments and networking.

11:35 Deep learnt and healthy: What AI and machine learning can do for lung cancer care
- How machine learning is being applied to healthcare
- Understanding some of the challenges in lung cancer
- The role of machine learning in patient stratification and diagnosis.

Dr Timor Kadir, Chief Scientist and Technology Officer, Optellum.

12:05 Understanding Patient Data: Opportunities, ethics and the public view
- How should we talk to patients and the public about new data-driven technologies in healthcare and research?
- Exploring the ethical and social issues that are likely to arise
- Finding out what the public think.

Dr Natalie Banner, Policy Adviser, Wellcome Trust.

12:35 Lunch and networking.

13:35 IBM Watson
- What is Watson, its underpinning technology and why it is useful to healthcare problems
- How Watson is being leveraged in cancer and oncology and what it is delivering (at scale now)
- What next with Watson (AI) and healthcare.

Stephen Boyle, Clinical and Healthcare Subject Matter Expert, IBM Watson Health.

14:05 Cancer image analytics
- Computational pathology is poised to change the status quo in pathology practice for the better
- Analytics of information-rich, high-resolution whole-slide images (WSIs, the so-called Big Cancer Image Data) could yield direct benefits for cancer diagnosis and prognosis
- Morphological motifs extracted from histology image data are likely to lead to patient stratification for precision medicine.

Professor Nasir Rajpoot, Department of Computer Science, University of Warwick.

14:35 Refreshments and networking.

15:05 From pathway to patient: a model-based approach to clinical decision support
- Codification of clinical pathways is allowing for doctors to make more informed decisions
- This will become increasingly powerful as we are able to seamlessly link population and individual behavioural data to them
- The data products that will emerge from this will utilise both machine learning and system-based modelling techniques
- Industrialisation of the workflows around these techniques will provide an opportunity to scale these approaches.

Karl Surmacz, Principal Data Scientist, McLaren Applied Technologies.

15:35 Panel discussion: Take the opportunity to put questions to our expert panel, and share solutions with your peers to overcome the challenges you are facing:
- Dr Neelam Dugar, Consultant Radiologist, Informatics Advisor – Royal College of Radiologists
- Niels Peek, Director, Greater Manchester Connected Health City.

16:30 Chair's final remarks.

16:40 Close.

Maggie Copping
Event Producer, IET Events
email: mcopping@theiet.org

northumbria UNIVERSITY NEWCASTLE

British Machine Vision Conference 2018

3rd - 6th Sept 2018

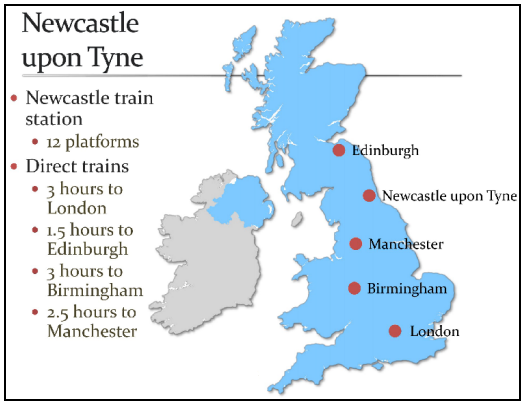

Dr Hubert PH Shum
Northumbria University
email: hubert.shum@northumbria.ac.uk