BMVA News

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

Volume 26 Number 2
December 2015

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

http://www.bmva.org/

BMVA News[1] is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the Editor by 10 March 2016.

Contents

Editorial: *Language Wars, Then and Now*

In the past 40–50 years computer vision has been hugely transmogrified from its original form which resembled image processing with a soupçon of statistical pattern recognition; first, 3D vision was added, and this introduced a large amount of complex mathematics which was used to wrestle with perspective plus a number of notable ambiguity and robustness problems; then general 3D shapes had to be dealt with and significantly more training/learning was required to cope with non-ideal situations including not only noise but also the textures of natural objects. People often had doubts about neural networks but the huge potential and actual power of deep learning networks eventually became apparent.

Meanwhile, workers had to struggle with a variety of computer languages in order to cope with the large amounts of statistics, probability calculations (not least for tracking objects using particle filters), 3D processing, and deep learning. The languages used in vision have included Fortran, Pascal, Prolog, C (with a variety of derivatives – C++, C#, …), Java, Matlab, Python, Julia, not to mention OpenCV. The less said about Fortran the better, though about 30 years ago I kept hearing that in the 2000s there would be one computer language, which would incorporate the good points of all the others: it would be called Fortran. Fortunately, this didn't come to pass and by 2003, much of the Particle Physics community had left it behind. Matlab originated around 1980, but (a) it was commercial and hence costly, and (b) it was targeted at mathematics users and was rather less useful for image analysis: hence it was ignored by many, in favour of C++ with extensive libraries. However, this has gradually changed and Matlab is now quite well adapted to computer vision needs. On the other hand, there is a love/hate relationship between it and C+Python. Arguably it is more polished than the latter, but it is accompanied by substantial costs that many cannot afford (though there is a trend to attractive package deals being made with universities for undergraduate and postgraduate use). But the high system reliability yet lack of flexibility that accompanies a commercial package is less attractive for some parties, such as (some) particle physicists which can't and mustn't be tied down by commercial enterprises. Currently, I find I'm caught 'between a rock and a hard place' in knowing which future to go for – Matlab or C+Python. This reminds me of the old arguments about whether to go with Algol or Fortran …

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

[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.

New Arrangements for BMVA Travel Bursaries

In 2016 there will be new arrangements for BMVA Travel Bursaries. The BMVA Executive Committee has decided to introduce four application deadlines during the year, the first of which will be 10 March 2016.

Please see the website link for news of later deadlines: http://www.bmva.org/bursaries

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

Sullivan Thesis Prize

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

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

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

Andrew Kay
http://invisibules.org
email: andrew.kay@sharp.co.uk

Upcoming BMVA Technical Meetings and Symposia

The articles below summarise the next four meetings in the BMVA Technical Meetings Programme: they are to take place in February, March, April and July 2016.

SketchX: Human Sketch Analysis and its Applications

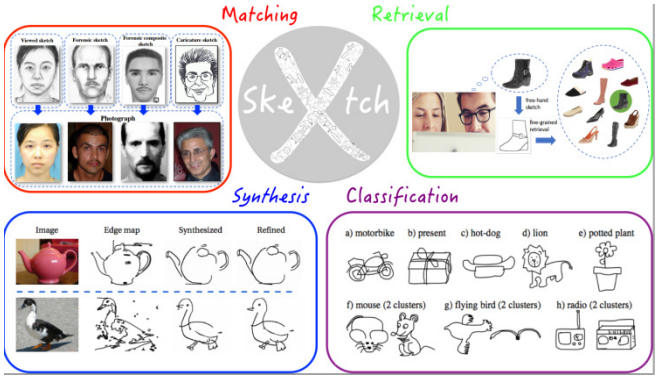

One-day BMVA symposium in London on Wednesday 24 February 2016

Chair: Yi-Zhe Song
Keynote speakers: James Wang (Penn State University), Metin Sezgin (Koç University), John Collomosse (University of Surrey), Charlie Frowd (University of Winchester and EvoFit), Pat Healey (Queen Mary University of London), Paul Rosin (Cardiff University)
www.bmva.org/meetings

The aim of this workshop is to showcase recent advances on human sketch analysis, and to promote discussions towards future research and commercialisation of sketches. The programme will feature a range of high-quality international keynote speakers.

Topic

Sketching comes naturally to humans. With the proliferation of touchscreens, we can now sketch effortlessly and ubiquitously by sweeping fingers on phones, tablets and smart watches. Studying free-hand sketches has thus become increasingly popular in recent years, with a wide spectrum of work addressing sketch recognition, sketch-based image retrieval, and sketching style and abstraction.

In this meeting we seek the opinions of speakers from leading research labs with an interest in human sketch analysis on the future of human sketch analysis and its applications. What research challenges remain? What are the key factors blocking progress? What new application opportunities can be envisaged? What elements are ready for adoption by industry?

Areas of interest covered by this meeting include:

- Sketch Recognition/Classification
- Sketch-Based Image/Video Retrieval
- Sketch-Based 3D Model Retrieval

- Sketch-Based Modelling
- Sketch Synthesis/Non-photorealistic Rendering
- Forensic Facial Sketch Synthesis
- Forensic Facial Sketch to Photo Matching
- Sketch for Human Computer Interaction
- Sketch for Learning and Education.

Andrew Gilbert
University of Surrey
email: ag0017@surrey.ac.uk

Computer Vision in Remote Sensing

One-day BMVA workshop in London on Wednesday 23 March 2016

Chairs: Hong Wei and James Ferryman, University of Reading
Keynote speaker: Uwe Stilla (Chair of ISPRS Inter-Commission Working Group III/VII: Pattern Analysis in Remote Sensing, Technische Universität München)
www.bmva.org/meetings

The workshop aims to bring together researchers in the combined field of computer vision and remote sensing to share ideas, outputs, and challenges for future research. It will feature a number of invited speakers, and a number of talks by interested parties.

Topic

Exploitation of spatial data to support earth observation is of significant interest to the computer vision research community. With increasing amounts of spatial data becoming available, and with rapid development of computational technology and remote sensing technology, there are demands to explore new algorithms which can automatically transform raw data to semantically meaningful information. Image analysis and computer vision technology will play an important role in this transformation.

In this workshop we hope to seek the opinions of speakers and participants from a wide range of interesting academics and companies with an interest in computer vision for remotely sensed data analysis and applications. What demands are there from the end-user point of view to make use of spatial data? What academic challenges remain? What elements are ready for adoption by industry? What are the key factors blocking progress? What new application opportunities can be explored?

Areas of interest covered by this workshop include (but are not limited to):

- Airborne and spaceborne remotely sensed data analysis
- Scene understanding and classification from remotely sensed data
- Fusion techniques in multi-spectral sensors
- Change detection from remote sensing
- Automated feature extraction for segmentation and object recognition in remote sensing
- Data-driven scene modelling for remote sensing
- 3D scene reconstruction from remotely sensed data
- Camera calibration, image registration in remote sensing.

Andrew Gilbert
University of Surrey
email: ag0017@surrey.ac.uk

Computer Vision for Automotive Applications – the Road Ahead

One-day BMVA workshop in London on Wednesday 13 April 2016

Chairs: Toby Breckon (Durham University), Anna Gaszczak (Jaguar Landrover)
www.bmva.org/meetings

Call for contributions

Automotive vision is an increasing area within computer vision research for both on- and off-vehicle applications. Whilst improved driver awareness is an obvious conduit, the greater potential for enhanced sensing within the wider road transport system offers numerous benefits for safety, security and possibly the environment alike.

Within this domain, this potential comes at the cost of hard constraints in terms of the obvious need for reliability and often real-time performance. However, the ever-reducing cost of computation coupled with increasing low-cost sensor capabilities now help drive both research and commercial systems for both on- and off-vehicle vision systems.

The aim of this meeting is to bring together researchers and practitioners, from both industry and academia,

interested in all aspects of automotive vision and in addition potential future applications of their work into this domain. Submissions are invited within in the following areas:

- autonomous driving
- driver assistance systems
- terrain/road classification
- highway monitoring and reporting
- driver monitoring
- improved driver awareness
- vision as a navigation aid
- automotive data fusion
- real-time obstacle detection
- pedestrian and vehicle detection
- vehicle situational awareness
- novel on-vehicle sensing
- vision for post-incident forensics
- roadside asset detection and cataloguing
- on-vehicle recognition and classification.

Other topics within the broadly applicable domain of automotive vision will also be considered for inclusion. Please submit an extended summary of your work by email attachment to Toby Breckon toby.breckon@durham.ac.uk by 3 February 2016.

Andrew Gilbert
University of Surrey
email: ag0017@surrey.ac.uk

BMVA Technical Meeting: Deep Learning for Computer Vision

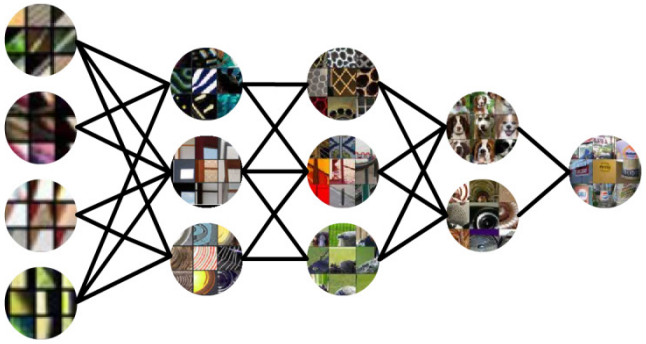

One-day BMVA symposium in London on Friday 8 July 2016

Chairs: Kai Arulkumaran, Anil Bharath
Keynote speakers: Andrea Vedaldi (University of Oxford), Ben Graham (University of Warwick), Samim Winiger (Samim.io), Sander Dieleman (Google DeepMind), Soumith Chintala (Facebook AI Research)
www.bmva.org/meetings

Call for contributions

The programme will feature a range of keynote speakers. In addition there will be poster sessions for interested parties. If you would like to bring a poster, please submit the proposed title and abstract via goo.gl/forms/Bn7o6TSaeM by Friday 15 April 2016.

Introduction

Deep learning is one of the hottest topics in computer science research right now, with numerous publications coming from both academia and industry. Since the success of Alex Krizhevsky's Convolutional Neural Network (CNN) on the ImageNet classification benchmark in 2012, deep neural networks have begun to dominate different areas in computer vision. Nowadays, some consider object classification to be an almost-solved problem, or at least in the regime of large amounts of data. Similarly, image segmentation has been massively improved through deep learning techniques. However, the lack of labelled data for supervised learning prohibits the application of large, vanilla CNNs to certain problem domains. Recent work has focused on how to tackle these more difficult challenges by combining traditional computer vision techniques with deep learning, as well as advancing more traditional deep learning methods themselves. The aim of the workshop is to explore not just the areas of computer vision that deep learning can advance, but also how the investigation of different areas in computer vision has led to advances in deep learning. On the one hand, more sophisticated architectures and training methodologies can be used to tackle more complex problems in computer vision. On the other hand, the domain of spatial and temporal dimensions in computer vision can be more intuitively explored by humans, leading to a better understanding of what, and possibly why, deep neural networks learn what they do. In particular, visualisation techniques have great potential in helping us understand both the models and the data that we work with.

Topics of Interest

With respect to computer vision, the areas of interest covered by this meeting include the following:

- Neural network architectures
- Deep learning optimisation techniques
- Visualisation of learned representations
- Practical applications of deep learning
- Generative neural network models
- Hybrid traditional CV/deep learning systems.

Registration

Book online at www.bmva.org/meetings:
- £10 for BMVA members
- £30 for non-members

including lunch

Andrew Gilbert
University of Surrey
email: ag0017@surrey.ac.uk

BMVC 2016

Announcement and Call for Papers

19–22 September 2016, York, UK
http://bmvc2016.cs.york.ac.uk/

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 27th BMVC will be held at the University of York, on 19–22 September 2016. The university is set on the outskirts of the historic walled city of York which features Roman and Viking history and is home to a number of world renowned tourist attractions ranging from the gothic cathedral York Minster to the National Railway Museum.

BMVC 2016 is a high quality single-track conference, comprising oral presentations and poster sessions (with oral acceptance <10% in the last 6 years). The conference features two keynote presentations and a tutorial, and has associated workshops on the last day of the conference, including a PhD student workshop.

Call for papers

Authors are invited to submit full-length high-quality papers in image processing and machine vision. Papers covering theory and/or application areas of computer vision are invited for submission. Submitted papers will be refereed on their originality, presentation, empirical results, and quality of evaluation. All papers will be reviewed *doubly blind*, normally by three members of our international programme committee.

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.

Accepted papers will be included in the conference proceedings, published and DOI indexed by BMVA.

Paper submission

Submission instructions, paper templates and other details are available on the conference website:

http://bmvc2016.cs.york.ac.uk/

Important Dates

Submission deadline:	9 May, 11:59 pm (Pacific time)
Author notification:	15 July
Camera ready deadline (including 1 p. abstract):	29 July
Conference tutorial:	Monday 19 Sept
Main conference:	Tu 20 – Th 22 Sept.

Invited speakers

Katsushi Ikeuchi, University of Tokyo
Raquel Urtasun, University of Toronto.

Tutorial

Measurement-based Appearance Modelling – Abhijeet Ghosh, Imperial College, London

Call for workshop proposals

Proposal deadline†	Friday 26 February
Notification:	Friday 11 March

†email to: bmvc16-group@york.ac.uk

Workshop proposals to be held in conjunction with BMVC 2016 are requested. Workshops will take place on the last afternoon of the conference on Thursday 22 September. The workshops are intended to be half-day events and to provide opportunities for in-depth discussions of focused topics, technical issues, or specific applications of computer vision and pattern recognition. Both established topics and new trends in related areas are welcomed.

The workshop organisers are responsible for soliciting papers, handling submissions, carrying out peer review, notifying authors, and collecting camera-ready papers within the specified deadlines. BMVC 2016 will provide meeting facilities, collect workshop registrations, and distribute electronic proceedings. Proposals will be relevant to BMVC and will be evaluated by the BMVC 2016 organising committee.

A workshop may contain both oral sessions and a short poster session. The overall length is restricted to a half day. Workshops will be made freely available to delegates that are registered to the main conference. Workshop presenters must register for the workshop event. Organisers may choose to invite speakers. The conference cannot provide financial support to invited speakers or the workshop organisers. One workshop registration fee can be waived for a workshop that has over 10 presenters.

Proposals should include the following information:

- Workshop title, and tentative call for papers for the workshop.
- Proposers' names, titles, affiliations, and primary contact email.
- Names and bios of any invited speakers and indication of whether they have agreed to speak.
- Proposed Workshop format.
- Description of relevance and viability.
- Topics that will be covered and rough program.
- Anticipated target audience as well as expected number of attendees, e.g. based on past events.

Proposals should be submitted by email to bmvc16-group@york.ac.uk by Friday 26 February.

Will Smith
University of York
email: william.smith@york.ac.uk

MIUA 2016

Announcement and Call for Papers

6–8 July 2016 at Loughborough University, UK

MIUA 2016 is the 20th in the series of annual meetings which is dedicated to research in the field of image analysis applied to the medical and biomedical sciences.

The conference provides an opportunity to present and discuss research in medical image understanding and analysis, which is a rapidly growing subject with ever increasing real world applicability.

Confirmed keynote speakers for this conference are:

- Dr Sue Astley, Manchester University, UK
- Professor Nico Karsemeijer, Radboud University, The Netherlands.

Technical papers (6 pages) and review papers (8 pages) are invited for peer review covering theoretical and applied research in the area of medical imaging.

The conference proceedings will be published online by Procedia Computer Science (Elsevier).

Conference Chairs: Professor Alastair Gale and Dr Yan Chen

Key Information

Paper submission opens:	18 January
Paper submission deadline:	1 April
Conference dates:	6–8 July
Further details:	http://www.miua.org.uk/
Contact email:	miua2016@lboro.ac.uk

Sarah Lowe
Loughborough University
email: s.e.lowe@lboro.ac.uk

A Robotics View of ICIP 2015

The 22nd edition of the International Conference on Image Processing was held during the last week of September in the picturesque Canadian city of Québec. ICIP is an internationally recognised conference thanks to the wide range of fields covered and the quality of the accepted papers. A total of 2306 papers were submitted from 64 countries, of which 1048 were accepted (45% acceptance rate). The conference lasted four days, and a variety of events and activities took place in it.

On the first day tutorials were provided on various areas of image processing. Unfortunately I could not attend them due to problems with flights. The schedule on the other three days was very similar: a plenary talk first thing in the morning, followed by three sessions of two hours each. While a plenary talk was devoted to the hot topic of deep learning, other relevant fields such as computational imaging were also covered. In this talk new advances and breakthroughs on image sensors were described, along with algorithms that make the most of captured images. Most of the papers were presented as posters and a small fraction (59) in oral sessions. I found this format very convenient, for I could attend the oral talks of my interest while I was still able to cover the poster presentations. I had many enriching discussions, learnt about various methodologies and got in contact with researchers from my area thanks to this format. During the sessions 34 works were shown in the format of Show & Tell, where the work is presented interactively. Relevant companies like Google, Microsoft Research, Nvidia and Cisco have sponsored, hosted workshops and performed exhibitions at the conference. There was a queue of job applicants at their booths at all times.

An overview of the audience during the plenary talk on advances on computational imaging.

Many social activities were scattered throughout the whole event. On Sunday evening a welcome reception was held at the Citadelle de Québec, a historical fortress of the city. On Monday the Women in Signal Processing Luncheon took place for the first time in the history of ICIP. The purpose of this activity is to foster support and offer networking opportunities by engaging women across the signal processing profession to exchange ideas and experiences from academia, research and industry in an informal setting. The Banquet & Awards Ceremony took place on Monday night. The awards were presented by the iconic Mrs. Lena Söderberg, whose famous picture has inspired so many image processing works. I must say that people got really excited to finally see her and that her presence was welcomed by several ovations from the researchers. Two new awards were introduced: the best industry paper and the best video of the 3-minute thesis video contest. This contest challenges the communication skills of PhD students on their theses. On Tuesday another new activity was held: the Get-Together Students-Employers Luncheon, where ICIP put together PhD students and industrial employers, an excellent occasion to find job opportunities and research candidates. Also on Tuesday a Young Professional Event was celebrated in the gorgeous, nineteenth century National Assembly Library. Even though I attended the conference as a student I could be in this event so I met other students and young professionals. It was for me a great way of doing networking, as I could discuss my work with other researchers.

Lena greets the audience of ICIP 2015.

I presented my poster with the title "Improved 3D Sparse Maps for High-Performance SFM with Low-Cost Omnidirectional Robots". I found a good reception of my work and it received great exposure, as well as valuable feedback from researchers. Overall, this conference has been an exceptional opportunity to get in contact with the research community and appreciate other works in my field. I am very thankful to BMVA for providing me with funds to cover the expenses for this trip.

Pedro Cavestany Olivares
Cranfield University
email: p.cavestany@cranfield.ac.uk

A 3D View of ICIP 2015

The 2015 International Conference on Image Processing (ICIP) was held in Québec City, Canada, on 27–30 September 2015. It is known to all that ICIP is one of the largest international conferences on image processing. This year it attracted 2306 paper submissions from both academia and industry, and the acceptance rate was around 45%. I was involved in two poster sessions and one 'show & tell' session. The topics of my posters were "Random Cascaded-Regression Copse for Robust Facial Landmark Detection" and "Fitting 3D Morphable Model using Local Features", and the topic of my show & tell session was "Landmark Detection and 3D Face Reconstruction using Modern C++". All of them were collaborated with Mr Patrik Huber and supervised by Professor Josef Kittler and Dr William Christmas, from the Centre for Vision, Speech and Signal Processing (CVSSP), University of Surrey. Our presentations were very successful and attracted many attentions.

Tutorials and Plenaries

Nine tutorials focusing on different aspects of image processing and analysis took place on 27 September. These tutorials were excellent and very well presented. Each tutorial cost 250 CAD (around £125) for an attendee. That was not very cheap. However, there were also three plenary talks at the conference. These plenaries were free to all the attendees.

Bird's eye view of Québec City

Tutorials

- Deep learning in image processing and vision (invited) – Yoshua Bengio and Roland Memisevic
- Hevc/h.265 video coding standard (v. 2) including range, scalable, and multiview extensions – Dan Grois, Benjamin Bross, Detlev Marpe and Karsten Suehring
- Image processing for cinema – Marcelo Bertalmío
- Visual saliency: fundamentals, applications, and recent progress – Ali Borji, Neil DB Bruce, Ming-Ming Cheng and Jian Li
- Computational photography (invited) – Mohit Gupta and Jean-François Lalonde
- Example-based super-resolution – Jordi Salvador and Mehmet Türkan
- Perceptual metrics for image and video quality in a broader context: from perceptual transparency to structural equivalence – Thrasyvoulos N Pappas and Sheila S. Hemami
- Spectral methods in 3D data analysis – Michael Bronstein
- Sparse stochastic processes: a unifying statistical framework for modern image processing – Michael Unser.

Plenaries

- Deep Learning – Yoshua Bengio
- Advances in Computational Imaging – Shree K Nayar
- Challenges and Opportunities in Biological Imaging – Michael Unser.

These talks provided valuable opportunities to broaden your horizons, not only for fresh PhD students but also for those who were interested in exploring new ideas and advanced techniques. The highlight is that Professor Yoshua Bengio was also invited to give a 3-hour tutorial and a plenary talk, both on deep learning. If you are interested in any of these tutorials or plenary talks, please go to the official webpage of ICIP 2015 for more details.

'Show & Tell' Sessions

Besides the traditional oral and poster sessions, the conference provided 'show & tell' sessions for the authors. It was fantastic. The show & tell session was a great idea for all attendees. The conference provided a platform for the speakers with a big TV, with that you could show your live demo with your innovations and new ideas on image and video analysis. We demonstrated our software for 3D face reconstruction from video sequences. The software was programmed in modern C++, which is also real-time. Our show & tell session was very successful and impressive. For the software, the reader is referred to my friend's webpage http://www.patrikhuber.ch/. You can download and play with it. With this software, you can track your face from a webcam and obtain your 3D face shape and texture information. Also, I will release my code for facial landmark detection and face tracking very soon on my personal webpage https://sites.google.com/site/zhenhuaswebpage/.

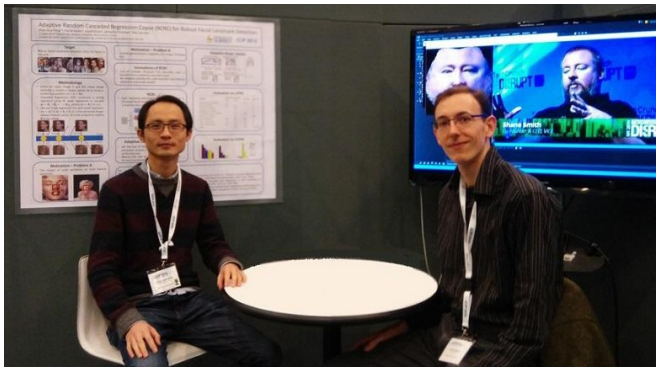

'Show & Tell' Session with Patrik Huber

Social Events

All the attendees were invited to the welcome reception. It was located at the Citadelle of Québec that is the largest British fortress in North America and Canada's oldest military building, with more than 300 years of history since the time of New France. The welcome reception was a great opportunity for the attendees to celebrate their arrival at the conference, to get to know each other and to exchange ideas.

The best paper awards were announced at the conference banquet, including three best student paper awards, two best paper awards and one best industry paper award. The highlight is that Lena Söderberg attended the banquet and presented awards to the winners. In addition, since ICIP 2015, the conference gives an industry best paper award for papers with the first author from industry.

The conference also provided many other social events for the attendees, such as the students-employers lunch and the women in signal processing lunch. Note that men are also welcomed to the women in signal processing lunch :)

Welcome Message from ICIP 2016

ICIP 2016 will be held in Phoenix, Arizona, USA, on 25–28 September 2016. Please visit http://icip2016.com/ for more details.

Zhenhua Feng
University of Surrey
email: z.feng@surrey.ac.uk

Report on CVPR 2015

The 28th IEEE Conference on Computer Vision and Pattern Recognition (CVPR) was held in June at the Hynes Convention Center in Boston, Massachusetts, USA. CVPR is an annual conference on computer vision and pattern recognition. According to Google Scholar's h-index metric in 2015, CVPR is the number one venue in Computer Vision and the seventh in Engineering and Computer Science. As the leading IEEE publication, it has higher h-index than any other conference in all fields, and is ranked to the top 70 of all publications. I was privileged to have the opportunity to attend and to present my recent research work at the event.

View of Boston city near the conference site

Conference registration desk

Out of 2123 valid submissions to the conference, 1815 were fully reviewed while the others were withdrawn before review, or administratively rejected for either technical or ethical reasons. For paper selection, 66 researchers were invited as Area Chairs (ACs) and a team of 1238 experienced reviewers from the broader computer vision and pattern recognition community was recruited for high quality review. In the end, the conference accepted 602 papers for publishing (28.4% of valid submissions), 71 of which were allocated as oral presentations (3.3% of valid submissions). All papers were presented as posters. Full camera-ready paper PDFs are now available on the Computer Vision Foundation website as well as the conference program page. This year, following the successful tradition at BMVC, CVPR authors were required to provide a one-page extended abstract describing their papers, which can also be downloaded individually from the conference program page.

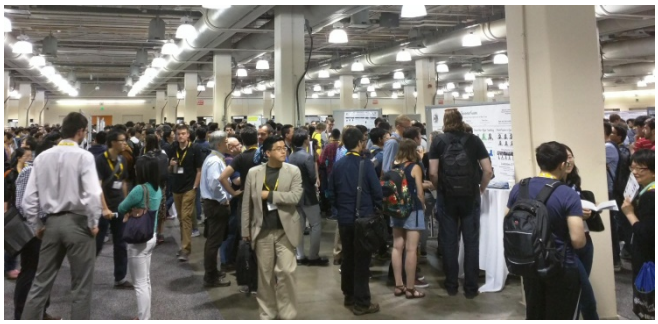

Poster session

In addition to the main three-day program (Mon–Wed, 8–10 June) of oral and poster presentations, plenary talks, demos, exhibitions, a doctoral consortium and social functions, CVPR 2015 held a number of co-located events, including 18 tutorials (Sunday 7 June) and 31 workshops (Thursday and Friday, 11–12 June). The tutorials covered a wide range of topics in computer vision and pattern recognition, for instance, 3D vision, OpenCV, OpenVX, optimization, deep learning with Torch and Caffe, etc. The workshops covered a wide range of theories and applications like biometrics, large-scale applications, embedded vision and Women in Computer Vision.

Plenary talk by Professor Yann LeCun

Not surprisingly, in this year's CVPR, Deep Learning attracted a lot of attention. A large proportion of papers were related to it. Several tutorials, workshops were specifically focused on deep learning applications, thus were especially popular among attenders. Professor Yann LeCun, the director of Facebook AI research and Silver Professor in New York University, gave a plenary talk on "What's wrong with deep learning". As a pioneer researcher in the field, he talked about several limitations of deep learning and introduced several solutions proposed by his team and other groups in the past decades. He also demonstrated the impressive diversity of Computer Vision applications such as OCR, body pose estimation, semantic segmentation, etc. Several professors and researchers predicted that learning large Deep Learning model for instance ConvNets from videos without annotations is going to be 'hot' at next year's CVPR.

The other plenary speaker, Professor Jack L. Gallant, from University of California, Berkeley, addressed the topic of "Reverse Engineering the Human Visual System", which was about how his lab uses neuroimaging methods to study how human brains represent and process sensory and cognitive information. He talked mainly about an approach in which researchers measure brain activity – by using functional MRI while the subject is watching naturalistic movies – and then decode brain activity and reconstruct the structural and semantic content in the movies, by applying the corresponding estimated models.

Plenary talk by Professor Jack L. Gallant

Several companies and research groups exhibited their technical products, attracted potential employees and promoted their on-going projects actively. Giant firms such as Microsoft Research, Google, Amazon, NVIDIA, Facebook, Intel, A9, as well as startups like SenseTime, Cogtu and MegVii served as platinum sponsors of CVPR this year. They demonstrated their solutions to a wide range of computer vision problems and showed many exciting products. Their representatives were sitting behind desks, introducing their vacancies and building connections.

Exhibitions of advanced computer vision products

This year's Best Paper was awarded to Richard A Newcombe, Dieter Fox, Steven M Seitz for "Dynamic Fusion: reconstruction and tracking of non-rigid scenes in real-time". The companies mentioned above sponsored all the paper awards.

In addition, there was a Doctoral Consortium designed for students who were close to completion, or had recently finished, their doctorate degrees. It was organized by Professor Adriana Kovashka from University of Pittsburgh and Professor Christoph Lampert from the Institute of Science and Technology Austria, and sponsored by Amazon and the US National Science Foundation. 31 students from different countries and areas were invited to present their PhD work. Each student was assigned a mentor from the community's senior members, based on the similarity of their research interests. I was very lucky to get involved and to have Professor Fernando De la Torre from the Robotics

Institute at CMU as my mentor. We had an inspiring discussion about my work on face analysis and its future trends in the research community as well as in commerce.

Overall, it was a highlight of my academic life. My poster, entitled "Mirror, mirror on the wall, tell me, is the error small?" interested the crowds and received appreciation. Last but not least, I must express my gratitude to BMVA for providing me funds for this trip.

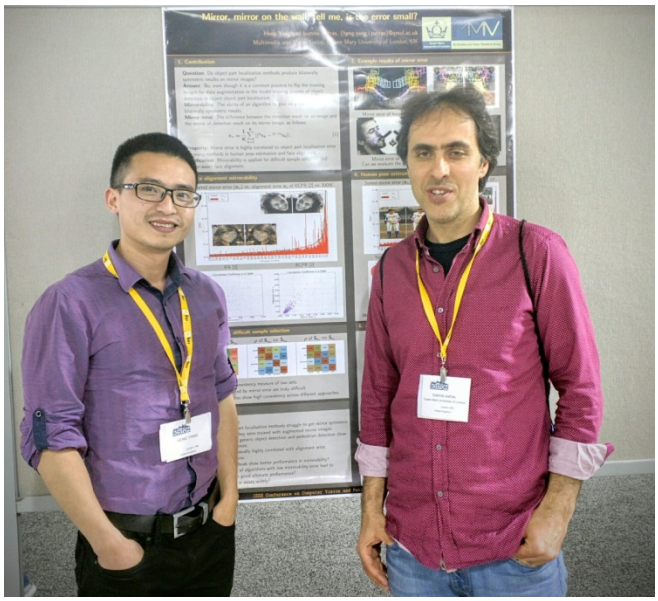

Heng Yang presenting his poster with his PhD supervisor Dr. Ioannis Patras (QMUL).

Heng Yang
University of Cambridge
email: heng.yang@cl.cam.ac.uk

Report on MICCAI 2015

The International Conference on Medical Image Computing and Computer Assisted Intervention (MICCAI) is one of the most prestigious conferences in the field of medical imaging. This year, the 18th MICCA was held in the scenic city of Munich, on 5–9 October. MICCAI annually attracts the leading scientists in the field, and it is a great privilege to be accepted there. This year, MICCAI received 810 submissions, selecting 263 papers for publication, and as is traditional in MICCAI these were mainly poster presentations. Three keynote speeches were presented, Professor Vasilis Ntziachristos and Professor Franz Pfeifer (both from TUM, the host university), and by the physics Nobel Laureate Professor Gerd Binning.

The conference was held in Munich Philharmonic Hall (the Gasteig). This was a very impressive building that really added to the vastness of the conference. Being situated along the river, a 30-minute walk from my hotel allowed for a nice relaxed stroll to and from the conference, a welcome break to the hectic schedule, running most days from 08:30 to 18:00. Most of the time was devoted to poster presentations: the sessions were all named after German cities and started with an overview of that city. I felt this was a very nice touch, as it allowed the hosts to share a lot of culture with the delegates who, like myself, did not have time to fully explore the region.

Wednesday morning saw the two local keynotes give their talks. These made a great start to the day, starting with Professor Ntziachristos giving an overview on new optical imaging techniques, particularly multi-spectral opto-acoustic tomography (MSOT). This was followed by Professor Pfeifer's talk on X-Ray phase-contrast imaging. The third keynote was scheduled for that evening's gala dinner.

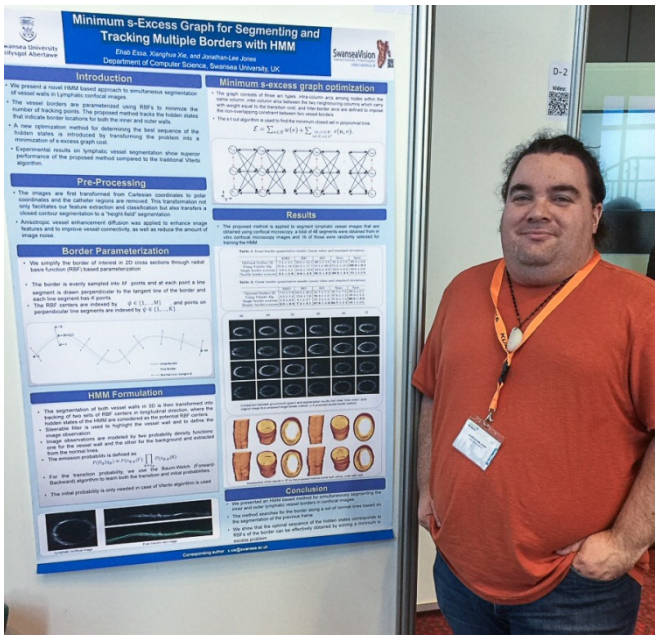

Jonathan-Lee Jones presenting his poster

My poster presentation was in the first session on the Tuesday. I have had experience doing poster presentations before, but they were nothing like this. I talked about the work and answered many questions (some rather in depth) almost constantly through the entire session (and half of the following lunch break). Luckily, my poster was not too far away from the refreshments, so I could grab a much-needed coffee in the one short respite I had. It was a great experience though, as some of the people interested in the poster were people that I had cited myself, and whose work I was familiar with. It was both challenging and rewarding talking through my work to people whose name had come up very often in my own research. In the other poster sessions, I was able to look at the other posters. Although some were in fields other than my own, it was interesting to see the use of methods such as those I have employed, applied to different applications and in different fields. There were also several interesting posters that had a similar end-goal to some of my research, but had gone about it in a very different manner. Talking to the authors of that work proved very useful in understanding the problem from a different point of view, and will perhaps lead to new ideas in my own work.

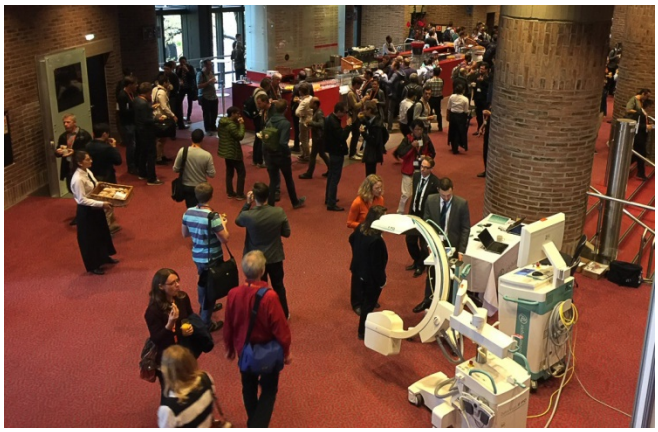

Munich itself, a very beautiful and friendly city, had its Oktoberfest running in the background: indeed this traditional event took on the role of the conference banquet. This was held at the Löwenbräukeller, a traditional Bavarian beer hall in the old Löwenbräu brewery. This was a nice

touch, as it allowed delegates to experience the traditional Bavarian festival despite the tight schedule of the conference. Indeed, it was one of the most interesting gala dinners I have ever attended, and was a little more involved and lively than most.

Jonathan-Lee Jones
Swansea University
email: j.jones.560541@swansea.ac.uk

Report on BMVA meeting: The Computational Face – Automatic Face Analysis and Synthesis

British Computer Society, London
14 October 2015

Automatic face analysis from visual data and its synthesis has attracted a lot of attention in the recent years due to its potential applications in the area of biometrics, forensics, psychology, mental health, human computer interactions, etc. To bring together people working in this exciting research area and to share their knowledge, BMVA organised a one-day symposium on this topic chaired by Dr Michel Valstar, Dr Brais Martinez and Dr Yorgos Tzimiropoulos.

The symposium was attended by approximately 70 people including people from academia as well as industry. After the initial registration and coffee in the morning, the first session of the symposium kicked off at 10 am with the keynote speaker Professor Maja Pantic. She gave a detailed overview of the field of facial expression recognition. In her talk, Professor Pantic described the importance, the current developments and the upcoming challenges in this field. Soon after that Dr Evangelos described about his work on subpixel temporal registration of face images in a video. In this work, he and his team posed face registration as a statistical learning problem to align faces in consecutive frames of a video. In the next talk, Dr Charlie Frowd described about his computer based system called 'EvoFit' which can be used to synthesize face images from human memory. The system works by combining a list of images selected by the user from an image database, according to his memory of the face to be synthesized. This talk was followed by a short coffee break.

The next session began with a talk from Dr Darren Cosker who gave an interesting overview of the area of dynamic facial processing and capture. In his talk, Dr Cosker described the main challenges in this field both in academia and industry. His talk was followed by a presentation from Dr William Smith in which he described his work on an example-based approach to face synthesis. In this approach, patches from example face images are combined together for modelling facial texture. Just before going for lunch, Patrik Huber, from the University of Surrey, described a C++ library he had developed for facial landmark detection and face reconstruction in 3D. This presentation was followed by lunch. The lunch break also saw some interesting poster presentations from Alessio Dore (Face detection, tracking and recognition for Interactive videos), Ali Bukar (Non-linear AAMs for age estimation and synthesis), Georgia Rajamanoharan (Multi-view Constrained Local models for large head angle facial tracking) and Joanna Olszewska (Automated Face detection: Challenges and Solutions).

The post-lunch session began with a keynote talk from Professor Tim Cootes, titled 'Facial feature detection using statistical shape models'. In his talk Professor Cootes gave an overview of statistical shape and appearance models for facial feature tracking. He also discussed regression-based methods for tracking facial features. In the next talk, Dr Ioannis Patras gave a very interesting presentation about how face alignment in mirror images can be used to estimate the alignment errors in the original images. This talk was followed by a presentation from Thomas Heseltine from Aurora Computer Services, on deep learning for biometrics. In his presentation, Thomas talked about the large improvement in performance he and his colleagues achieved by using deep learning methods on the task of face recognition and facial feature tracking. This talk was followed by another short coffee break.

The last session of the day began with a keynote talk from Professor Richard Bowden, titled 'Faces in Communication'. In his presentation, Professor Bowden gave an interesting overview of the field of sign language and lip reading. He described his current work and the main challenges facing this field. This talk was followed by another interesting talk from Dr Timothy Hospedales from Queen Mary University, London. In his talk, Dr Hospedales described about his work on cross-modality matching of forensic sketches and caricatures to photos. Dr Hospedales and his colleagues used a mid-level attribute representation which can be learnt independently within each modality and are invariant to cross model gap. The symposium ended with a presentation from Dr Anastasios Roussos from Imperial College, London. In his presentation, Dr Roussos described about his work on a variational approach to the problem of dense 3D face reconstruction from a video sequence.

The symposium formally ended at 5 pm but the discussions continued beyond the presentation room and over to the dinner table. Overall, the symposium was successful in bringing together people in this field and provided a platform for exchange of ideas and networking opportunities in academia and industry.

Shashank Jaiswal
University of Nottingham
email: psxsj3@nottingham.ac.uk