

# BMVA News

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

Volume 26 Number 3  
March 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** News<sup>1</sup> is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the Editor by 10 June 2016.

## Contents

Editorial: <i>Strategies for Thinking – from Deduction to Intuition</i> .....	1
Professor Ioannis Pitas has Moved! .....	2
Maria Petrou Prize .....	2
Call for Expressions of Interest to Host BMVC in 2018 .....	3
SketchX: Human Sketch Analysis and its Applications – Meeting Report .....	3
Deep Learning for Computer Vision .....	4
Upcoming BMVA Technical Meetings .....	5
BMVA Computer Vision Summer School 2016 .....	5
27 <sup>th</sup> British Machine Vision Conference .....	5
MIUA 2016 .....	6
Report on the first UK Manipulation Workshop .....	6
Report on ICCV 2015 .....	7
‘Epipolar Disorder’ .....	9
ICDP 2016 .....	9
Nominations for the BMVA Executive Committee .....	12
Nomination Form for the BMVA Executive Committee ...	12

## Editorial: *Strategies for Thinking – from Deduction to Intuition*

Suddenly the world seems to be buzzing with the exploits and endeavours of scientists. Indeed, the capabilities of the Google autonomous car have been widely known and accepted for quite some time, and the disasters that I predicted in that quarter haven’t happened; interestingly, the first accident caused partly by the car failed to make much of a shock wave, even though the apportionment of blame doesn’t seem to have been resolved with any generality. “What caused the accident?” you may validly ask. Apparently, it was the lack of any realistic algorithm for the bluffs that drivers make with each other when road lanes converge, and the car didn’t appreciate that larger vehicles such as busses have to be allowed to win priority.

If vision has been centre stage in the public mind for a few years, what of AI? Arguably, it has been there for far longer. After all, it is almost 20 years since the world chess champion Garry Kasparov was beaten by the Deep Blue computer, though it is now considered to have cheated, to the extent that it worked through all possible moves (up to ~20 moves ahead), and thus used brute force to achieve superiority. It clearly wasn’t ‘thinking’ in any humanly accepted sense, and couldn’t therefore adapt its methodology to activities other than rather artificial games in the Noughts and Crosses, Draughts and Chess series. Nevertheless, 20 years on, people wondered whether computers could tackle the ancient Chinese game of Go with any success, now that computers have become so much faster and more powerful. Hence, a Go tournament was set up and to everybody’s surprise (including that of the program designers) the AlphaGo computer managed to beat the human challenger in 4 out of 5 games – an achievement that was not expected for another 10 years or so. But how was this possible? In fact, it was not just power but deep thinking that did it, in the form of deep learning which is now a widely recognised computer vision tool that has advanced well above the noise level, seemingly faster than exponentially, during the past 2–3 years.

This achievement has hit the public press so hard in the past 2 weeks that I haven’t been able to avoid it: articles in *The Week*, in *Vision Systems Design* and two in the *Sunday*

<sup>1</sup> 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.


*Times* came to my notice: all had their own slant on it, and none spouted bad science or hysteria. The leading *Sunday Times* article by Dominic Lawson – who hails from a family more renowned for politics and economics than for science – seemed especially knowledgeable, with statements like “[Go’s] exponential branching factor on a scale that defies brute force calculation”. This illustrated that a  $19 \times 19$  Go board embodies unimaginable numbers of possible moves, so that it would be quite impossible to use the brute force methods employed by Deep Blue for supremacy at chess. *The Week* article twice enlightened me that to succeed at Go required ‘intuition’ instead of brute force. Apparently, that is what deep learning has to offer. But what is intuition? Here I am reminded of the Hough transform, whose claim to fame is that of recognising highly occluded objects, so inference is required rather than deduction. I can imagine that the true meaning of intuition is learning for itself (*inwards-tuition*). Well there is some truth in this. For the AlphaGo program starts with a knowledge of the moves, gets some tuition in what to do, and then practices interminably against itself, learning many (though not all) possible strategies and undergoing huge amounts of training until it is in a position to beat a human brain. Of course, it does not need to sleep but plays games continuously for months at a time until it can beat its other self. In that case how could a human possibly win? Well, the one game out of five that Lee Sedol won was achieved by making a totally unexpected move! Clearly, AlphaGo’s training included only playing against geniuses like itself, so it hadn’t been properly prepared, as we pattern recognition experts know it should have been, by making its training fully representative of what it might meet during testing. But I digress. We now see that intuition must be the process used in training a deep thinker, and inference must be the process of using it in testing, i.e., intuition leads to the capability for inference (which is generalisation over gaps in knowledge).<sup>2</sup>

Unfortunately, I am not wholly heartened by all this. Admittedly, brute force has rightly been replaced by intuition and inference, but the confidence that we now have in deep learning has a shaky foundation, as we have no real idea what goes on in the mind of a modern intelligent computer program such as AlphaGo. Instead we are left with no option other than to trust it as we know it has been infallibly trained on real data. As I used to say in the days of artificial neural networks – I have managed to replace my skill as an algorithm designer by my skill in finding the right data on which to train my computer, and of course I am ten times better at that. But either way, it remains a question of ‘garbage in, garbage out’.

Overall, the big test for computers is whether you trust them to drive you or fly you or whether you feel they are merely using you as data to train themselves in an interesting new game. But then, what is ‘interesting’ for a computer?

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

## Professor Ioannis Pitas has Moved!



Professor Ioannis Pitas (IEEE fellow, IEEE Distinguished Lecturer, EURASIP fellow) received the Diploma and PhD degree in Electrical Engineering, both from the Aristotle University of Thessaloniki (AUTH), Greece.

Since 1 January 2016, he has been Professor in the Department of Electrical and Electronic Engineering, University of Bristol, UK. He has served as a Visiting Professor at several Universities.

His current interests are in the areas of image/video processing, intelligent digital media, machine learning, human-centred interfaces, affective computing, computer vision, 3D imaging and biomedical imaging. He has published over 800 papers, contributed to 44 books in his areas of interest and edited or co-authored another 10 books. He has also been member of the program committee of many scientific conferences and workshops. In the past he served as Associate Editor or co-Editor of eight international journals and General or Technical Chair of four international conferences. He participated in 68 R&D projects, primarily funded by the European Union and was principal investigator/researcher in 40 such completed projects. He has 22200+ citations to his work and h-index 70+ (to March 2016).

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

## Maria Petrou Prize

**The deadline for submission of nomination and endorsement forms is 30 April 2016.**

The Maria Petrou Prize is to be awarded biennially at ICPRs to a living female scientist/engineer who has made substantial contributions to the field of Pattern Recognition (or a closely related field), and whose past contributions, current research activity and future potential may be regarded as a model to both aspiring and established researchers. This Prize honours the memory of Professor Maria Petrou as a scientist and engineer of the first rank, and particularly in her role as a pioneer for women researchers and highly successful role model. She is widely recognized for her extensive contributions to the field of image processing and pattern recognition. She also made significant contributions to the growth of IAPR, covering

<sup>2</sup>See the Web for many other attempts to identify intuition: naturally, some of these vary from my own interpretation.

significant leadership roles. The Prize consists of a suitably inscribed plaque and a cash amount.

The prize recipient will be selected by the Maria Petrou Prize Committee, subject to approval by the IAPR Governing Board, upon nomination by a member of a national member society of IAPR and by the endorsement of two more members, at least one of which must be a woman.

Members of the IAPR Executive Committee, as well as of the Maria Petrou Prize Committee, shall be ineligible for the prize and may not serve as nominators or endorsers.

The 2016 prize will be presented at the 23<sup>rd</sup> International Conference on Pattern Recognition (ICPR), Cancun, Mexico, 4–8 December 2016.

The prize recipient is expected to present an invited talk at the conference and to provide a contribution to a special issue of Pattern Recognition Letters, which will include extended versions of all papers that received an IAPR award at the 2016 ICPR.

The nomination must be made on a special nomination form.<sup>3</sup> Completed nomination and endorsement forms must be submitted in electronic form, and must be received no later than April 30 2016. The nominator and endorsers should email their completed forms directly to the Appointed Chair of the Maria Petrou Prize Committee at the email address below.

Edwin Hancock  
 Chair, Maria Petrou Prize Committee  
 University of York, UK  
 email: erh@cs.york.ac.uk

## Call for Expressions of Interest to Host BMVC in 2018

The venue and arrangements for BMVC are planned two years in advance. The venue for the 2017 event is now settled, so we are now starting to consider what will happen in 2018. To that end, the Executive Committee would like to hear from groups interested in hosting BMVC in 2018. At this stage only an expression of interest is required: following discussion by the Executive Committee, these expressions of interest may be followed up by requests to supply an official bid to hold the conference.

For the expression of interest please supply the following details:

- Main contact for the conference: name, email, phone.
- Prospective members of the conference organising committee: general chair, technical chair, local arrangements chair, and so on. There is no need for all these members to be from the same institution; indeed, we encourage you to spread the load.
- Provisional dates for the conference, with a confirmation of provisional booking, details of accommodation, rooms and arrangements for conference venue and meals.

BMVC is traditionally held in one of the first two weeks of September and runs as a single-track conference with

podium and poster sessions from Tuesday through to Thursday lunchtime. The pre-conference Monday afternoon is normally scheduled for a tutorial session, and separate workshops may be held on the Thursday afternoon or on Friday.

The main conference auditorium must be large enough to accommodate 300 delegates and should be a tiered auditorium with adequate A/V facilities. The poster room should be capable of holding 150 2m × 2m poster boards, ideally in the same contiguous space (e.g., a large hall). For further information on preparing a proposal to run BMVC please refer to:

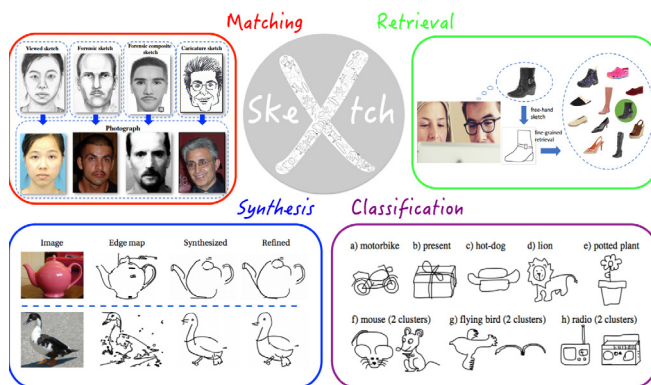
[http://www.bmva.org/w/bmvc\\_proposals](http://www.bmva.org/w/bmvc_proposals)

Please send expressions of interest to the BMVA Chair, Dr Adrian Clark, by 1 July 2016.

Adrian Clark  
 BMVA Chair  
 email: chair@bmva.org

## SketchX: Human Sketch Analysis and its Applications – Meeting Report

24 February at the British Computer Society, London



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. Aiming to bring together people working in this exciting and upcoming research area, BMVA organised a one-day workshop on human sketch analysis and its applications, chaired by Dr Yi-Zhe Song.

This workshop was structured so as to maximally stimulate idea exchange over a broad range of areas associated with sketches. It consisted of a total of 6 keynotes and 13 posters on a variety of topics, either directly related to sketches, e.g., facial sketch synthesis and sketch-based image retrieval, and those closely relevant topics that had not been specifically explored so far but were otherwise key in answering some of the fundamental issues on sketches, such as computational aesthetics, visual abstraction and non-photorealistic rendering. The workshop was attended by a

<sup>3</sup>Maria Petrou Prize Nomination and Endorsement Forms are available at: [http://iapr.org/fellowsandawards/awards\\_petrou.php](http://iapr.org/fellowsandawards/awards_petrou.php)

total of 29 people including academics as well as those from industry.

After the initial registration and welcoming session in the morning, the workshop kicked off at 9.30am with a keynote from Professor James Wang (Penn State University). James offered an insightful overview of his work on image aesthetics and visual abstraction, in particular automatic aesthetic evaluation for photos and categorisation of oil paintings. Soon after that Dr John Collomosse (University of Surrey), delivered a keynote about his pioneering work on sketch-based image retrieval and sketch-based video retrieval, and concluded with a few key challenges of sketch research.

This second keynote was followed by a long coffee break, the first of three sessions where posters were presented. These included poster presentations from James Wang (Development and validation of the image stimuli for emotion elicitation), Zainab Harbi (Segmentation of clock drawings based on spatial and temporal features), Alexandra Bonnici (Vectorization of sketched drawings using co-occurring sampling circles), Tu Bui (Scalable sketch-based image retrieval using color gradient features), Shuxin Ouyang (ForgetMeNot: reversing human forgetting for forensic facial sketch matching), Yi Li (Free-hand sketch synthesis with deformable stroke models), Yonggang Qi (Making better use of edges via perceptual grouping), Ke Li and Kaiyue Pang (Fine-grained sketch-based image retrieval: the role of part-aware attributes), Haojing Diao (Sketch-based musical composition and performance), Qian Yu (Sketch me that shoe) and Stuart James (Towards sketched visual narratives for retrieval).

The next session began with a talk from Professor Pat Healey from Queen Mary University of London, who gave an interesting overview of sketch research in the context of multi-modal human communication. The talk was followed by an extended lunch session, where posters were intensively discussed. The afternoon session began with a keynote speech by Dr Metin Sezgin from Koc University, who presented his work on sketch-based intelligent user interfaces. His talk was followed by a keynote from Professor Paul Rosin from Cardiff University on visual abstraction and abstract non-photorealistic rendering of images. A long coffee break followed, during which posters were again deliberated. The workshop ended with a presentation by Dr Charlie Frowd from the University of Winchester and founder of EvoFit. He presented his forensic sketch synthesis system which can be used to synthesize forensic e-fits from human memory, and offered an overview of the field of forensic sketch synthesis in general.

The workshop successfully concluded at 5 pm, but discussions continued beyond the presentation room. In particular it was recognised by the organisers and keynote speakers that it would be beneficial to the larger vision community for a similar sketch-specific workshop to happen alongside mainstream vision conferences. A joint workshop proposal at ECCV'16 had resulted from this discussion.

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

## Deep Learning for Computer Vision

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

<http://www.bmva.org/meetings>

### Call for Posters

Chairs: Kai Arulkumaran, Anil Bharath  
Speakers: Alex Kendall (University of Cambridge), Andrea Vedaldi (University of Oxford), Ben Graham (University of Warwick), Nicolas Koumchatzky (Twitter Cortex), Samim Winiger (Samim.io), Sander Dieleman (Google DeepMind), Soumith Chintala (Facebook AI Research).

### 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](http://goo.gl/forms/Bn7o6TSaeM) by Friday 15 April 2016. An NVIDIA GPU will be awarded as a prize to the best poster presented at the meeting.

### Introduction

Right now, deep learning is one of the hottest topics in computer science research, 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](http://www.bmva.org/meetings):  
 £10 for BMVA members  
 £30 for non-members including lunch.

Andrew Gilbert  
 University of Surrey  
 email: [ag0017@surrey.ac.uk](mailto:ag0017@surrey.ac.uk)

### Upcoming BMVA Technical Meetings

- |             |  |
|-------------|--|
| 8 July 2016 | Deep Learning for Computer Vision:<br>Kai Arulkumaran and Anil Bharath                       |
| 19 Oct 2016 | Vision for Interaction: from humans to<br>robots: Nicoletta Noceti and Alessandra<br>Sciutti |
| 16 Nov 2016 | Plants in Computer Vision: Hannah Dee  |
| 25 Jan 2017 | Transfer Learning in Computer Vision:<br>Dima Damen  |

Andrew Gilbert  
 University of Surrey  
 email: [ag0017@surrey.ac.uk](mailto:ag0017@surrey.ac.uk)

### BMVA Computer Vision Summer School 2016

4–8 July, Swansea University, UK



**• 1 week of intensive training • 17 speakers from both academia and industry • 19 lectures including practical sessions • 21 years of BMVA tradition**

The BMVA annual Computer Vision Summer School, aimed at PhD students in their first year and researchers at an early stage in their careers. Despite the title, students from non-UK universities are welcome to attend, as well as students from UK universities.

The 2016 Summer School will take place at Swansea University, Bay Campus, during 4–8 July. It will consist of an intensive week of lectures and lab sessions covering a wide range of topics in Computer Vision and Digital Image Computing. Lecturers are researchers from the most active Computer Vision research groups in the UK. A list of

invited lecturers and the summer school programme are available on the summer school website:

<http://cvss.swansea.ac.uk>

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

**NOTE: Places are limited to ensure good interaction in lab classes.**



### Important Dates

Early registration: 6 May  
 Late registration: 24 June  
 Summer school: Monday 4 – Friday 8 July

Dr Rita Borgo  
 Swansea University, UK  
 email: [r.borgo@swansea.ac.uk](mailto:r.borgo@swansea.ac.uk)  
 twitter: @BmvaCvss



### 27<sup>th</sup> British Machine Vision Conference

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 27<sup>th</sup> BMVC will be held at the University of York on 19–22 September. 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). It features two keynote presentations and a conference 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 on 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://bmv2016.cs.york.ac.uk/>

Paper submission is via CMT at:

<https://cmt.research.microsoft.com/BMVC2016>

### Important Dates

Submission deadline:	9 May, 11.59pm (Pacific time)
Author notification:	15 July
Camera ready deadline (including 1 page abstract):	29 July
Conference tutorial:	Monday 19 Sept
Main conference:	Tu 20 – Th 22 Sept

### Journal Special Issue

A selection of the best papers will appear in a special edition of the International Journal of Computer Vision (IJCV).

### Invited speakers

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

### Tutorial

Abhijeet Ghosh, Imperial College, London  
Topic: Measurement Based Appearance Modelling

Dr William AP Smith  
University of York, UK  
email: [william.smith@york.ac.uk](mailto:william.smith@york.ac.uk)

## MIUA 2016

6–8 July 2016 at Loughborough University, UK

MIUA 2016 is the 20<sup>th</sup> 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.

The conference proceedings covering theoretical and applied research in the area of medical imaging will be published online by *Procedia Computer Science* (Elsevier).

Conference Chairs: Professor Alastair Gale and Dr Yan Chen

For further details including registration information please visit the website, or use the conference email address:

<http://www.miaa.org.uk/>  
email: [miaa2016@lboro.ac.uk](mailto:miaa2016@lboro.ac.uk)

Sarah Lowe  
Loughborough University  
email: [s.e.lowe@lboro.ac.uk](mailto:s.e.lowe@lboro.ac.uk)

## Report on the first UK Manipulation Workshop

When I first joined you all in the world of computer vision, I did so with one key motivation: to design robots. However, whilst recent developments in vision have seen computers learn to recognise objects at an ability close to that of humans, and reconstruct the 3D world at even super-human capacities, hoping for similar levels of success in the physical, interactive world of robots, has often been frustrating. Due to its fundamentally challenging nature, state-of-the-art research in robot manipulation still addresses seemingly simple demands such as picking up objects from a table, and generalised methods for performing tasks with these objects remain very primitive. But it is for this very reason, encouraged by the abundance of remaining problems yet to be solved, that I recently began to move from a pure vision background into the world of visually-guided robot manipulation.

Bringing together the healthy breadth of talent across the UK's robotics community, this February saw a successful meeting in the form of the first UK Manipulation Workshop, held at the University of Birmingham. The two-day event involved a series of talks and poster sessions with 87 attendees, covering topics spanning mechanical design, tactile sensing, control, learning, and vision. Such a diverse mixture of skills and attitudes is somewhat rare at the typical

computer vision meetings I am accustomed to, but acutely demonstrates the scale and complexity of the problem we face in real-world robotics. In my opinion, coercing researchers into viewing the same problem but from rather unfamiliar, unnatural perspectives can only be a fruitful enterprise.

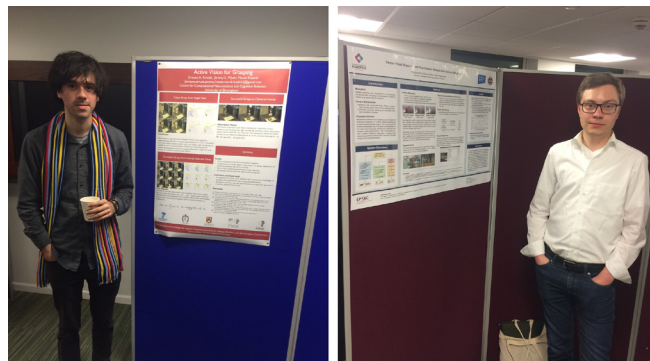


Professor Jeremy Wyatt, host and organiser of the workshop, presenting an overview of his group’s research on planning for object manipulation.

Over the course of the meeting, it became clear to me that vision has a firm place in manipulation as being an assistant to the wider problem, rather than a solution in its own right. Design of compliant hands that would conform passively to the shape of an enclosed object, or artificial tactile skin with sub-millimetre sensory resolution, will play as much a part in breakthrough robots as traditional vision-based tasks such as recognising, reconstructing, or tracking objects. Nathan Lepora (Bristol) spoke about tactile pressure sensors based on miniature cameras recording the displacement of a pattern of markers on the fingertip, which reminded me of structured light methods for depth cameras. Sethu Vijayakumar (Edinburgh) discussed visual servoing for guiding an arm towards a target under occlusion and camera movement. Marek Kopicki and Jeremy Wyatt (both at Birmingham) presented their approach to robotic grasping, whereby visual cues indicate suitability for both local fingertip targets and global hand configurations, which are later fused for an optimum grasp. Paul Siebert (Glasgow) spoke about clothing recognition based on learning local 3D textures extracted from depth images. Maurice Fallon (Edinburgh) discussed his experiences with the DARPA Robotics Challenge and in particular humanoid robots, such as using stereo vision to plan stable stepping points over uneven floors. Finally, an industrial talk by Graham Deacon (Ocado) revealed some exciting commercial applications of robots, in picking up and placing objects by locating pre-determined grasp points in images, soon to be automated in their delivery warehouses.

As for myself, I spoke of my recent work on Convolutional Neural Networks and synthetic training data for multi-view object recognition (to appear at this year’s CVPR) and robust object grasping. Indeed, it was a rare treat these days to be the only speaker discussing deep learning applications, and such methods are still seemingly fresh to the robotics community. If computer vision is the applied, real-world spin-off from machine learning, often taking a couple of years to apply its findings, then perhaps robotics is that further spin-off from computer vision.

However, of course, this leads to the following question: How can we generate an ImageNet-scale dataset of real robot manipulation sequences...?




Ermano Arruda (Birmingham) and Wolfgang Merkt (Edinburgh) presenting their work at the poster session.

Edward Johns, Dyson Fellow  
Imperial College London  
email: e.johns@imperial.ac.uk

## Report on ICCV 2015

A BMVA Travel Bursary allowed me to attend the International Conference on Computer Vision (ICCV) during 11–18 of December. The location for the conference was the city of Santiago, Chile, from where we enjoyed warm summer temperatures right next to the Andes.



Outside the conference venue

The eight-day event was centred around the main four-day conference, with satellite workshops and tutorials on the first and last two days. This report is only for the main conference that started on 13 December, when about 1400 delegates arrived in the underground Centroparque Convention Center for the largest ICCV conference to date. An opening speech by Ruzena Bajcsy, Professor at

University of California Berkeley, officially started the conference. During her speech, she stressed the need for development of the underlying theories within computer vision, and I was reminded of the high standards of research this conference prides itself in. Subsequently, a number of oral presentations were given on the topic of Vision and Language, followed by a poster session.



The conference hall

Later in the morning, a plenary talk was given by Stephen Boyd of Stanford University on the topic of Convex Optimisation. He has developed software to solve any convex optimisation problem with minimal user interaction – the user only needing to express the problem mathematically, typically in less than 10 lines of code. This reduces the need for users to implement their own convex optimisation approaches, and, he claimed, to even understand how convex optimisation algorithms operate. In the afternoon, we returned for an oral and a poster session on the topic of 3D Computer Vision.



A poster session

For the most part, day two of the main conference proceeded in a traditional fashion, with two oral sessions in the morning. Lunch was often obtained from a nearby shopping centre, Parque Arauco, where a very large range of snacks and meals were available – sufficiently many for the ICCV delegates. A subsequent oral session on Statistical Methods and Learning, followed by a poster session, were held in the afternoon.


At the end of the day, the awards for the conference were given. The honourable mention for the best paper went to Saining Xie and Zhuowen Tu for “Holistically-Nested Edge Detection”, with the best paper to Peter Kontschieder, Madalina Fiterau, Antonio Criminisi, and Samuel Rota Buló for their paper “Deep Neural Decision Forests”. Other prizes included the Mark Everingham prize, for a selfless

contribution of significant benefit to the Computer Vision community going to Daniel Scharstein, Richard Szeliski and team for the Middlebury Dataset; and to Andrea Vedaldi for his VLFeat Software. The PAMI Distinguished Researcher Award went to both David Lowe, for his work on SIFT, and Yann LeCun, one of the founding fathers of the Deep Learning architecture. It was clear throughout the awards ceremony that Deep Learning is continuing to develop as a hugely important topic in Computer Vision, as demonstrated by both awards and the content of the acceptance speeches. A subsequent drinks reception was held in the park above the conference centre to provide further networking opportunities for the delegates.



A recording of Andrea Vedaldi's acceptance speech

The next day I presented my own paper entitled “Globally optimal 2D-3D registration from points or lines without correspondences”, as a poster. The two and a half hours for each poster session flew by as a presenter since many delegates were eager to ask questions and learn about as much work as they could in the allotted time. The poster session gave me a fantastic opportunity to describe and discuss my work with experts in the field, and provided me with valuable feedback and ideas for future research. During the rest of the day, a doctoral consortium was held for invited PhD students to discuss their own work with a nominated expert in the field, after which a poster session was given and subsequently an oral session on Vision and People.



Drinks reception

The final day of the conference went smoothly with orals and poster sessions across a range of topics. In the morning the sessions were centred around Computational Photography, Face and Gesture, and Motion and Correspondence. The final session of the day was on Video – Actions, Surveillance and Tracking. At the end of the




conference, a closing ceremony was held, including a prize for the best demo. During the conference, delegates had voted for the location of ICCV in 2019, and we subsequently learnt during the closing ceremony that Seoul won by a large proportion of the votes. First, however, ICCV 2017 will be held in Venice. A final drinks and nibbles reception was provided in the park above the conference centre that was enjoyed by all.

I would sincerely like to thank the BMVA for their generous support of this travel grant. Without it I would not have been able to attend, and as such could not have had such interesting discussions with some of the very best in the field. Furthermore, attending ICCV has allowed me to observe the best research by the top researchers throughout computer vision. I would also like to thank my supervisors Jean-Yves Guillemaut and David Windridge for their help and support throughout my PhD.

I would like to acknowledge Charles Malleson for allowing me to use some of his photographs in this report.

Mark Brown  
University of Surrey  
email: mb00336@surrey.ac.uk

## ‘Epipolar Disorder’



In my opinion you don't need glasses; but I will write you a prescription for a new fundamental matrix.

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

## ICDP 2016

The 7<sup>th</sup> International Conference on Imaging for Crime Detection and Prevention (ICDP 2016), organised by IET's Vision and Imaging Network.

ICDP 2016 will be held at Universidad Carlos III de Madrid, Spain on 23–25 November 2016.

### Call for Papers

General Chair: Sergio A Velastin (Universidad Carlos III de Madrid)  
Co-Chair: Dimitrios Makris (Kingston University, London)

### Aims and scope

Crime and anti-social behaviour have a significant cost for society and business alike. Just in the UK, anti-social behaviour alone accounts annually for around £3.4 billion of taxpayers' money with incidents of graffiti and vandalism estimated to cost around £600 million/p.a. Surveillance systems of all kinds are thus extensively deployed in public and private locations to deter, prevent and control. The last years have also seen an increased awareness on the vulnerability of public spaces to attacks. However, there are serious limitations to the use of conventional monitoring systems where human operators are asked to survey a large number of cameras or go through enormous amounts of recorded material for forensic investigations. Computer-based technologies are increasingly becoming researched in what is now known as video analytics, propelled by advances in processing power, embedded computing, IP-networking technologies, volume storage, cheap cameras, etc. The realisation of such advances into working systems can have a major impact on Society but also on individual liberty. This conference follows the successful IDSS (Intelligent Distributed Surveillance Systems) events held in 2003 and 2004 and ICDP 2005, 2006, 2009, 2011, 2013 and 2015 to bring together researchers, industry, end-users, law-enforcing agencies and citizens groups to share experiences and explore areas where additional research and development are needed, identify possible collaboration and consider the societal impact of such technologies.

Full papers are invited on all aspects of Imaging Surveillance technologies, from academia, industry, NGOs and others, to be selected for oral presentations or posters through a peer-review system (see also: <http://www.icdp-conf.org>). An indicative, non-exclusive, list of relevant topics is:

- Surveillance systems and solutions (system architecture aspects, operational procedures, usability, scalability)
- Multi-camera systems
- Information fusion (e.g., from visible and infrared cameras, microphone arrays, etc.)
- Learning systems, cognitive systems engineering and video mining
- Robust computer vision algorithms (24/7 operation under variable conditions, object tracking, multi-camera algorithms, behaviour analysis and learning, scene segmentation)
- Human machine interfaces, human systems engineering and human factors

- Wireless communications and networks for video surveillance, video coding, compression, authentication, watermarking, location-dependent services
- Metadata generation, video database indexing, searching and browsing
- Embedded systems, surveillance middleware
- Gesture and posture analysis and recognition
- Biometrics (including face recognition)
- Forensics and crime scene reconstruction
- X-ray and terahertz scanning
- Case studies, practical systems and testbeds
- Data protection, civil liberties and social exclusion issues.

Accepted papers will be published on the IET's Digital Library, indexed by Inspec (and normally by IEEE Xplore) only if at least one author registers and presents the work.

Authors of exceptional papers may be invited to submit extended versions to be considered for publication in one of the following peer-reviewed Journals: IET Computer Vision, IET Image Processing or IET Biometrics. There are delegate fee discounts for authors, students and members of the IET and sponsoring organisations.

#### Key dates

Receipt of full papers (maximum of 6 pp. using the prescribed PDF format)	15 Aug
Notification of acceptance	3 Oct
Receipt of camera-ready papers	24 October

Professor Sergio A Velastin  
University Carlos III Madrid Spain  
email: [sergio.velastin@ieee.org](mailto:sergio.velastin@ieee.org)

This page is deliberately left blank.

# Nominations for the BMVA Executive Committee

Nominations are requested for the forthcoming election of Executive Committee members of the BMVA. Nominees must be paid-up members of the Association and agree to serve for a period of two years. A member of the Committee is expected to participate in roughly five committee meetings per year, taking place typically in London and Birmingham. Reasonable travelling expenses to attend meetings are paid by the BMVA.

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

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

<http://www.bmva.ac.uk/admin/exco.html>

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

Adrian Clark  
BMVA Chair  
email: [chair@bmva.org](mailto:chair@bmva.org)

## Nomination Form for the BMVA Executive Committee

### To be completed by the Nominator

As a fully paid up member of the BMVA, I,

Name: .....

Address: .....

.....

wish to nominate:

Name: .....

Institution: .....

for the Executive Committee of the BMVA.

Signed: .....

### To be completed by the Nominee

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

Name: .....

Address: .....

.....

email: .....

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

Signed: .....

### Brief Biography of Nominee

*Please send completed nomination form by post or email to:*

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