BMVA News

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 June 2015.

Contents

Editorial: *Intuition for Robots?*

Every now and again one gets a thrill from reading an article, and it was a pleasure for me to read the report by Nadejda Roubtsova on 3DV 2014 (3D vision), now included on pp. 7–9 of this issue. Writing a good article involves two things – having exciting content and also writing in a

---

[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.

gripping way, so that the reader can't stop reading until he's finished. Arguably, the days of skilled letter and diary writers are nearly over, or at least far less honed and practiced with the advent of the computer and people's far busier lives: remember that there wasn't much to do in the long dark evenings, before cinema, TV and cheap telephone calls took over our lives, so inter-personal writing was at the core of what people did. Today's zero depth 140-character Twitter messages are a far cry away.

One thing that drew my attention in Nadejda Roubtsova's article was the report on Jean Ponce's talk about "a more generic approach to projective geometry *beyond the use of the fundamental matrix*" and involving "geometric constructions *giving clues to an object's shape*" (my italics). This was music to my eyes: at last the fundamental matrix, with all its mathematical insight and intelligence, is dead. Long live the new regime! For I have long believed that the human brain does not work mathematically, and certainly not by using differential geometry, matrix inversion and the like. Of course, our own computer vision voyages will have to take the path they will, not least with whatever mathematical paraphernalia we happen to have at our disposal; and in any case even if we manage to unravel the wiring diagram of the human brain, how would we know the program it is using? (Think about it: if an alien had one of our computers, but the stored program memory unit had not also been taken to his spaceship, how could he possibly work out its capabilities?) As a result, we have to cut our cloth according to our means and make use of the impoverished programming tools that we have developed to date.

But if advanced maths can't help us enough, maybe we can take some tricks on board. For example, how can one possibly catch a cricket ball without modelling its parabolic motion using differential equations? Maybe not so difficult: all that's necessary is to keep it on a constant bearing in the sky until it falls into our hands – and there's absolutely no mathematics in that. Even a toddler can do it.

Actually, I've often wondered what representation the brain uses when we're driving a car. Does it use the eye's 2D view? Or does it calculate a plan (ground plane) view by a suitable transformation? Or does it work on a more intellectually based abstract view based on lists of objects with accompanying numeric parameters? Or does its

working resemble a graphics Z-buffer approach? Actually, I favour the first of these, using as basis an 'out of sight, out of mind' model, i.e. don't worry about things that are not in the line of sight: they are irrelevant as there is at least one other object between them and us to make them redundant. It seems to me that all driving actions can be undertaken without recourse to any more complex model. But I was forgetting about navigation relative to a map. Well, that is an exception, but here again, we can form weak links between the 2D eye view and a similar view of a map: and don't forget that the 2D eye view includes all road signs and other information needed to make comparisons with a map.

My other bête noire regarding vision is the ubiquitous use of high accuracy probabilities to represent anything, as for example with particle filters. Here probabilities are all we've got, but to my mind the penny hasn't yet dropped on how to avoid all the maths that the brain almost certainly isn't using. Of course, if we think about Schwarzenegger's Terminator character, with his hugely accurate robot eyes, he must work by that sort of dead reckoning. But we limited humans – never. What we do is use 'intuition' instead. There, problem solved! However, my critics in the 3D vision and motion world will never let me get away with saying that, though ultimately, my contention is that we poor humans have to rely far more on context than robots seem to. We have brilliant, vivid, and surprisingly accurate mental models of what goes on around us, so we simply don't need to do all that mathematical and probabilistic calculation. And what is intuition anyway? Awareness, knowledge of the world, of what might happen, of what is happening, of what is likely to happen next, and where to look to find out and refine and update our world model. Of course, there is one major problem of working in this way: we can only see what we expect to see, and we miss other things in an unbelievable way. For instance, I sometimes press a key on my computer and wait for an action to occur, such as another screen of information appearing: then I find that it had appeared so quickly that I had missed it (the likelihood of this is all the greater if the previous and updated screens have some similarity, which happens quite often). There's also the example of the interviewer telling viewers to count how many times the dancers spin round, and they miss the gorilla that passes across the stage: one's concentration is high because one is locked into the current vital task and excludes information that is not relevant to it. Maybe not surprising when one remembers that the eye interrogates scenes using saccades. And even less surprising when one remembers that the eye has no other means of examining a scene: it's *all* done using saccades. Which proves that all human scene interpretation is done by using intuition and making relevant hypotheses and testing them.

Why is it done this way? One answer is that the eye channel provides information at far too low a rate to do anything else. Another is that asking questions and answering them is an overridingly powerful way of getting information. Another is that the available input data is so unreliable that there is no other way of working anyway. Maybe BMVA News will give a prize for the first reader to tell us which of these is the correct answer …

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

BMVA Computer Vision Summer School 2015

6–10 July, Swansea University
http://cvss.swansea.ac.uk
twitter: @BmvaCvss

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

BMVA runs an annual Computer Vision Summer School aimed at PhD students in their first year, though it will be beneficial to other 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. Places are limited to ensure good interaction in lab classes.

The 2015 Summer School will take place at Swansea University on 6–10 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 summer school programme are available on summer school website.

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.

Some quotes from delegates who attended the previous summer schools at Manchester University and Kingston University: "I think the variety of the topics is very good", "Diverse set of requirements were handled very well. I found it extremely helpful and fun!", "Nice overall atmosphere to get in touch with people working in a similar/related field". "Lecturers did a very good job in bringing the topics across". "Thanks for organizing! I've met cool people and learnt a lot".

Important dates

Early registration: 8 May
Late registration: 26 June
Summer school: Monday–Friday, 6–10 July

Registration

There are two tiers of registration rates: UK and Non-UK. Anyone who is currently studying or working in a UK institution that is physically based in the UK is payable at the UK rate. For the rest, the non-UK rate applies. If you have any questions, please contact the organiser.

Residential rates: includes all sessions, accommodation (four nights, Monday–Thursday), refreshments, breakfast and lunch (Tuesday–Thursday), reception, and banquet.

Non-residential rates: includes all sessions, lunches and refreshments, reception, and banquet.

Registration costs

	Early by 8 May	Late by 26 June
UK:	£410.00	£460.00
UK non-residential:	£250.00	£300.00
Non-UK:	£650.00	£700.00
Non-UK non-residential:	£490.00	£540.00

N.B. Onsite accommodation is only guaranteed for early registrations.

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

BMVC 2015 – Call for Papers

7–10 September 2015, Swansea, UK
http://bmvc2015.swan.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 26th BMVC will be held at Swansea University Singleton Campus on 7–10 September. The Singleton Campus is set in a rolling parkland overlooking the majestic sweep of Swansea Bay, the start of the famously dramatic Gower coastline comprised of twenty-one bays and coves.

BMVC 2015 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 conference tutorial, and has five associated workshops on the last day of the conference, including a PhD student workshop.

Topics of interest

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 visualisation, interaction, and graphics
- Visualisation for computer vision
- Visual analytics for computer vision
- Object detection and recognition
- Shape-from-X
- Video analysis and event recognition
- Illumination and reflectance.

Conference tutorial

BMVC 2015 features a half-day conference tutorial on 7 September. The tutorial is particularly beneficial to research students and early career researchers who are working in this field. We are honoured, with Chris Bishop, to have such a prominent researcher in the field of pattern recognition and machine learning to deliver the tutorial.

Professor Christopher Bishop

Chris Bishop has a BA in Physics with First Class Honours from Oxford, and a PhD in Theoretical Physics from the University of Edinburgh with a thesis on quantum field theory supervised by David Wallace and Peter Higgs. In 1998 he joined the Microsoft Research Laboratory in Cambridge where he became Deputy Managing Director, and later the Chief Research Scientist. He is a Partner in Microsoft, and is head of the Machine Learning and Perception group. In 2010 he was awarded the accolade of Distinguished Scientist, representing the highest level of research distinction within Microsoft, and was the first person in Europe to hold this title. At the same time as he joined Microsoft Research, he was elected to a Chair of Computer Science at the University of Edinburgh where he is a member of the Institute for Adaptive and Neural Computation in the School of Informatics. He is also a Fellow of Darwin College, Cambridge. He has been elected Fellow of the Royal Academy of Engineering, and Fellow of the Royal Society of Edinburgh, and has been awarded two Honorary DSc degrees. His research interests include probabilistic approaches to machine learning, as well as their applications in industry, commerce, and healthcare.

Keynote speakers

BMVC invites two leading researchers in the field to present their work at the conference. We are grateful to the following speakers who have agreed to give keynote lectures at the conference.

Professor Ron Kimmel

Ron Kimmel is a Professor of Computer Science at the Technion where he holds the Montreal Chair in Sciences. He held a post-doctoral position at UC Berkeley and a visiting professorship at Stanford University. He has worked in various areas of image and shape analysis in computer vision, image processing, and computer graphics. Kimmel's interest in recent years has been non-rigid shape processing and analysis, medical imaging and computational biometry, numerical optimisation of problems with a geometric flavour, and applications of metric geometry and differential geometry. Kimmel is an IEEE Fellow for his contributions to image processing and non-rigid shape analysis. He is an author of two books, an editor of one, and an author of numerous articles. He is the founder of the Geometric Image Processing Laboratory and a founder and advisor of several successful image processing and analysis companies.

Professor Kristen Grauman

Kristen Grauman is an Associate Professor in the Department of Computer Science at the University of Texas at Austin. Her research in computer vision and machine learning focuses on visual search and object recognition. Before joining UT-Austin in 2007, she received her PhD in the EECS department at MIT, in the Computer Science and Artificial Intelligence Laboratory. She is an Alfred P Sloan Research Fellow and Microsoft Research New Faculty Fellow, a recipient of NSF CAREER and ONR Young Investigator awards, the Regents' Outstanding Teaching Award from the University of Texas System in 2012, the PAMI Young Researcher Award in 2013, the 2013 Computers and Thought Award from the International Joint Conference on Artificial Intelligence, and a Presidential Early Career Award for Scientists and Engineers (PECASE) in 2013. She and her collaborators were recognised with the CVPR Best Student Paper Award in 2008 for their work on hashing algorithms for large-scale image retrieval, and the Marr Best Paper Prize at ICCV in 2011 for their work on modelling relative visual attributes.

Paper submission

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. Please note that BMVC is a single-track meeting with oral and poster presentations. Paper submission instructions are available at the conference website: http://bmvc2015.swan.ac.uk/?services=call-for-papers

The paper submission site is now open: https://cmt.research.microsoft.com/BMVC2015/

Important dates

Submission deadline:	Monday 4 May
Author notification:	Friday 3 July
Early registration:	Friday 24 July
Late registration:	Friday 7 August
Conference tutorial:	Monday 7 September
Main conference:	Tuesday–Thursday, 8–10 Sept
Conference Workshops:	pm on Thursday 10 September.

BMVC 2015

Conference chairs

Xianghua Xie, Swansea University, UK
Mark Jones, Swansea University, UK
Gary Tam, Swansea University, UK

Dr Xianghua Xie
Swansea University
email: bmvc2015@swansea.ac.uk

Upcoming Technical Meetings

6 May	Vision for human–computer interaction and virtual reality systems. Manuela Chessa and Fabio Solari – programme below
17 June	Real-time 3D Scene Understanding in year 2020. Zeeshan Zia and Andrew Davison – programme below
17 July	Context Aware Cognitive Systems. Nicolas Pugeault, Sinan Kalkan, Frank Guerin and Angelo Cangelosi
14 Oct	BMVA technical meeting in face analysis and synthesis. Michel Valstar and Brias Martinez
9 Dec	Visual, tactile and force sensing for robot manipulation. Lorenzo Jamone and Emre Ugur.

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

Vision for human–computer interaction and virtual reality systems

One-day BMVA symposium to be held in London on Wednesday 6 May.
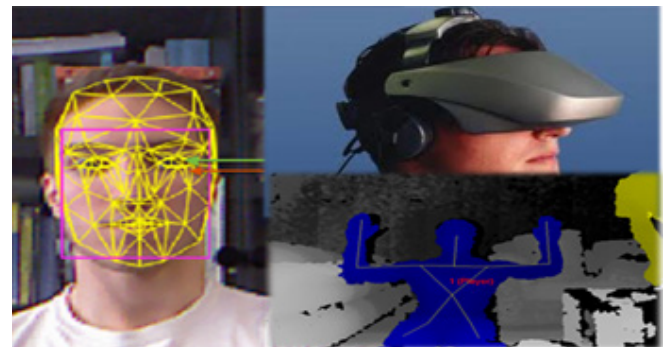
Chairs: Dr Manuela Chessa and Dr Fabio Solari

Keynote speakers: Dr Tae-Kyun Kim and Dr Guido Maiello

Programme

9.30	Registration + Coffee
10.00	Welcome and Introduction of the meeting. Manuela Chessa, University of Genoa
10.15	KEYNOTE SPEAKER I Depth perception and binocular vision in naturalistic virtual reality. Guido Maiello, University College London
11.00	Presenting natural focus cues in virtual/augmented reality. Simon Watt, Bangor University
11.30	Performance, experience and appearance: three dimensions of quality in 3D environments. PB Hibbard, RL Hornsey, P Scarfe, L O'Hare, T Zhang and HT Nefs, University of Essex
12.00	Lunch
13.00	KEYNOTE SPEAKER II: Human-action recognition. Dr Tae-Kyun Kim Imperial College London
13.45	A novel method for real-time multiple human action recognition. Victoria Bloom, Kingston

14.15 Hand pose and orientation estimation for ego-centric devices. Muhammad Asad, Greg Slabaugh, City University London

14.45 Understanding human motion and its qualities. Francesca Odone, University of Genoa

15.05 Coffee break

15.20 Self-organizing neural integration of pose-motion features for human action recognition. German I. Parisi, University of Hamburg

16.00 Good practices of hand gestures recognition for the design of customized NUI. Nicoletta Noceti, University of Genoa,

16.20 You-Do, I-Learn: unsupervised multi-user egocentric approach towards video-based guidance. Dima Damen, Teesid Leelasawassuk, Walterio Mayol, University of Bristol

16.50 Plenary discussion and demos (all).

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

Real-time 3D Scene Understanding in year 2020

One-day BMVA symposium to be held on Wednesday 17 June.
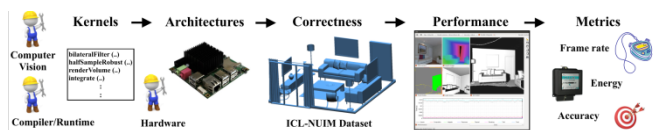
Chairs: Zeeshan Zia and Andrew Davison

Real-time 3D scene understanding is the capability which will enable the applications that have always been expected from sensing equipped AI robotic devices and systems which can interact fully and safely with normal human environments to perform widely useful tasks. A key reason that such applications have not yet emerged is simply that the robust and real-time perception of the complex everyday world that they require has simply been too difficult to achieve, algorithmically and computationally. This has been especially true in the domain of commodity-level sensing and computing hardware which is where the potential for real world-changing impact lies.

In the PAMELA research project (a collaboration between Manchester, Edinburgh and Imperial) we are asking questions about the way that computer vision algorithms will co-evolve with computer architecture and programming tools in the coming years. We predict that mobile-class hardware in the 2020s will be increasingly massively parallel, but also heterogeneous and specialised, and that power consumption measures will be critical. Computer vision is likely to be one of the key application domains driving the design of such hardware; but what are the essential algorithms in this fast-moving field which we should aim to optimise in architecture choices, or even to design specific processors for? And how can we give application programmers a usable and performance-portable interface to what might be complicated architecture under the hood?

In this meeting we hope to seek the opinions of speakers from the most interesting companies and academic labs with an interest in algorithms, processors and applications for real-time vision. It will be an extremely high quality meeting with a programme of 100% invited talks from industrial and academic leaders in the domain.

Speakers and Panellists for Discussion

- David Moloney, CTO at Movidius
- Gerhard Reitmayr, Principal Engineer at Qualcomm
- Jamie Shotton, Principal Researcher at Microsoft Research
- Mike Aldred, Electronics Lead at Dyson
- Simon Knowles, CTO at XMOS
- Simon Lynen, Researcher at ETH-Zurich and Google Project Tango
- Thomas Whelan, Dyson Research Fellow at Imperial College London
- Zeeshan Zia, Research Associate at Imperial College London
- TBC, Co-founder at Surreal Vision Ltd.

Registration

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

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

BMVA Workshop on Context Aware Cognitive Systems

One-day BMVA workshop to be held on Friday 17 July.

Chairs: Nicolas Pugeault, Sinan Kalkan, Frank Guerin and Angelo Cangelosi

Keynote Speakers

Professor Kenny Coventry (UEA), Professor David Hogg (Leeds), Dr Walterio Mayol-Cuevas (Bristol), Dr Ben Tatler (Dundee), Professor Florentin Woergoetter (Goettingen, Germany)

Context is the totality of the information characterising the situation of a cognitive system; e.g. it can include objects, persons, places, and temporally extended information related to ongoing tasks, but also information not directly related to these tasks. The way natural cognitive systems respond to a given stimulus also depends on other, apparently irrelevant stimuli which constitute what we call context. In fact natural systems not only consider these "apparently irrelevant" stimuli, they use them to their advantage: there is a wealth of evidence that context is used

routinely for reasoning and disambiguation, for tasks as varied as object recognition (eg., Palmer, 1975; Torralba, 2010), categorisation (Bar, 2004), language understanding (Coventry et al., 2010), problem solving (Cheng and Holyoak, 1985).

In contrast, artificial cognitive systems have long considered context as spurious, to be ignored, factored out or minimised in lab environments. This is now changing: in the recent years, there has been increasing efforts from the robotics and AI communities to research and design systems that can perform outside of closely controlled lab environments in complex, confusing, real-world situations. This will require a new generation of artificial systems, endowed with the cognitive tools to detect, assess, process and learn contextual relations (e.g., Choi et al, 2012; Li et al, 2012). This problem is intrinsically cross-disciplinary, touching aspects of computer vision, robotics, artificial intelligence, machine learning and psychology.

This workshop is an interdisciplinary meeting, bringing together psychologists, AI scientists and roboticists to discuss theories and challenges related to the representation, learning and processing of context in natural and artificial cognitive systems.

Call for papers

This one-day workshop will be the first interdisciplinary colloquium to discuss the multiple facets of context and the challenges it poses to cognitive systems. Contributions are welcome from all related fields, including new experimental results and models as well as position or review papers. Topics of interest include, but are not limited to:

- Theory of context
- Psychological perspectives
- Attention and context
- Experiments on context awareness
- Computational models of context
- Topic models, e.g., multimodal latent Dirichlet allocation to model context
- Context for machine learning algorithms
- Context in language interpretation, and language generation
- Context in other aspects of human-robot interaction
- Context in visual perception
- Context for navigation
- Context for robot manipulation – adjusting actions to the specifics of the current context
- Context for robot planning
- Using context to improve robot safety – e.g., moving carefully when humans are close.

We invite contributions to the workshop in the form of an abstract, to be submitted via www.bmva.org/meetings by 24 April.

The workshop will consist of a mixture of oral presentations, poster sessions and a panel discussion. There will also be an organised after-meeting dinner – booking also through www.bmva.org/meetings.

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

Interested in Chairing your own Meeting?

We are currently *seeking volunteers* to host meetings for 2016 onwards. This is not a hard task and involves (i) organizing a program of speakers (by call for papers or invitation), (ii) chairing the meeting on the day – and that's it! All other arrangements (finances, food, registration, room-booking, etc.) will be taken care of by myself. Please contact myself, Andrew Gilbert (a.gilbert@surrey.ac.uk), if you are interested in hosting one of the future meetings.

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

ICDP 2015 – Call for Papers

The 6th International Conference on *Imaging for Crime Detection and Prevention* (ICDP 2015) will take place on 15–17 July at QMUL, and is organised under the auspices of IET's *Vision and Imaging* Network.

Co-sponsors include: BMVA, UKIVA, IEEE, QMUL, Kingston University, UCL, FONDECYT-Chile, LASIE FP7 Project.

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 aims 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, and to 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 – but only if at least one author registers and presents the work. Authors of exceptional papers will be invited to submit extended versions to be considered for publication in IET Computer Vision, IET Biometrics or IET Image Processing Journals. 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 April
Notification of acceptance: 25 May
Receipt of camera-ready papers: 15 June

> Professor Sergio A Velastin
> Universidad de Santiago de Chile
> email: sergio@velastin.org

MIUA 2015 – Call for Papers

This conference will take place on 15–17 July at Lincoln. http://www.miua.org.uk/

MIUA 2015 is the 19[th] in the series of annual meetings dedicated to communicating research progress in image analysis applied in the medical and biomedical sciences.

Technical papers (6 pages) and review papers (8 pages) are invited on topics from across the spectrum of medical image analysis – from theory and technical advances to novel practical applications.

MIUA operates a double-blind peer review system with three reviewers evaluating each paper. All accepted contributions will be published and the full proceedings (DBLP indexed) will be available at the conference and on-line. Selected papers will be invited for journal (ISI indexed) publication and prizes will be awarded for the best work.

Important dates

Paper submission open: 6 February
New paper submission deadline: 13 April
Author notification: 11 May
Upload camera-ready MIUA papers: 25 May
Conference: 15–17 July

We are also pleased to announce the distinguished keynote speakers for MIUA 2015:

- Professor Brian F Hutton (University College London)
- Professor Tim Cootes (University of Manchester)
- Professor Alejandro Frangi (University of Sheffield)
- Associate Professor Alfredo Ruggeri (University of Padua, Italy).

Further details of the conference, paper submission, scope and venue can be found at http://www.miua.org.uk/.

For questions about paper submission, academic content, sponsorships and all technical enquiries please contact: miua2015@lincoln.ac.uk.

In 2015, the historic city of Lincoln celebrates the 800[th] anniversary of the signing of Magna Carta, http://www.visitlincoln.com/magnacarta.

Chairs: Dr Tryphon Lambrou and Dr Xujiong Ye at the University of Lincoln.

> email: miua2015@lincoln.ac.uk

Report on 3DV 2014

Last year the International Conference on 3D Vision (3DV) was held on 8–11 December. The location of the conference in the vibrant city of Tokyo, the capital of the technological pioneer Japan, was highly suitable reflecting the dynamic nature of the research field of 3D vision. Unlike other broader computer vision conferences, 3DV specifically targets research in theory and applications of 3D vision, providing a forum for the members of this tight-knit community.

3DV came into existence in 2011 when two related conference 3DIM and 3DPVT joined forces to form what was initially called 3DIMPVT. Later re-branded as 3DV, the event continued to expand despite staying true to its focused single-track format. This year, in addition to the main conference, a full day of tutorials and workshops was organised to further explore the key interests of the community, such as RGB-D data and dynamic 3D reconstruction, in a more informal setting.

Ian Reid: building and using prior shape models, 3D tracking and SLAM

The main conference offered an abundance of keynote presentations as well as oral and poster paper deliveries organised in thematic sessions. A passion for research and vision for ever greater advances of the field in the future characterised all keynote talks.

Ian Reid from the University of Adelaide (Australia) presented the evolution from 2D tracking to state-of-art 3D deformable object tracking with 6 degrees of freedom where the target is *a priori* unknown. Reid's group continues to formulate new challenges for 3D tracking algorithms. The newest addition is described in this year's 3DV submission "3D tracking of multiple objects with identical appearance using RGB-D input" by CY Ren et al. that, having shown impressive results, was given the 'Best Paper' award at the conference.
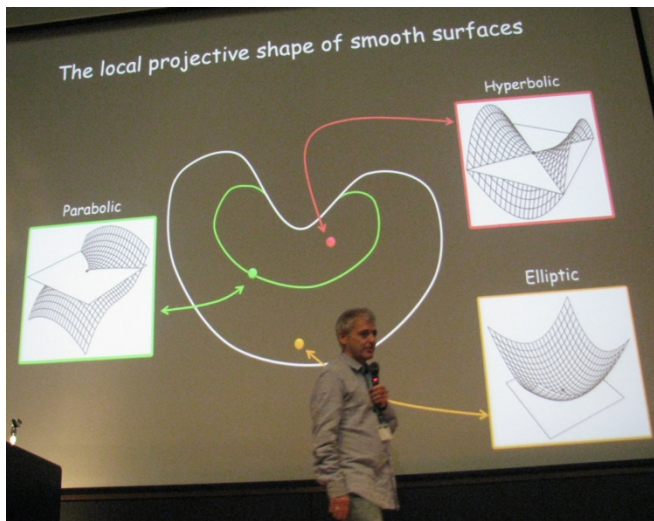

Jean Ponce: Towards geometric foundations for computer vision

Jean Ponce from ENS (France), a well-established figure in the 3D reconstruction community, presented his vision about a more generic mathematical approach to viewing projective geometry beyond the use of the fundamental matrix. The speaker introduced a series of purely geometric constructions giving clues to an object's shape. Light incidence and curvature could according to Ponce play the central role in the process of 3D reconstruction instead of the universally targeted depths and normals.
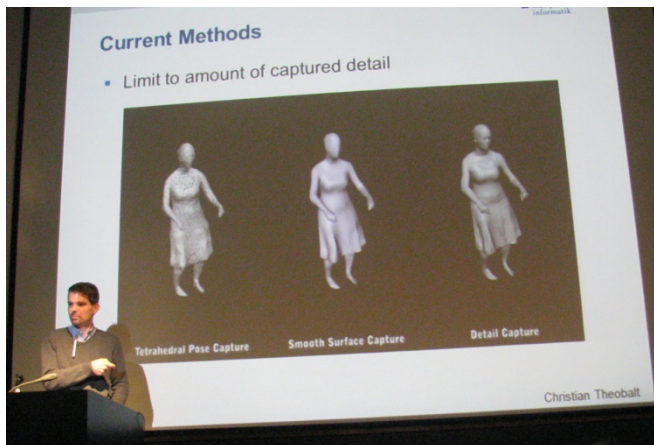

Christian Theobalt: Capturing and editing the real world in motion

Christian Theobalt from Max-Planck-Institute for Informatics (Germany), known as an individual linking 3D reconstruction and graphics communities, presented a subjective view of the 3D model capture technology currently lagging behind in development relative to the model rendering technologies. His vision is 4D model capture in uncontrolled environments in such a manner that permits alteration by for example the post-production artist. Theobalt presented impressive research results towards the goal of rectifying what he sees as a development imbalance between reconstruction and graphics.
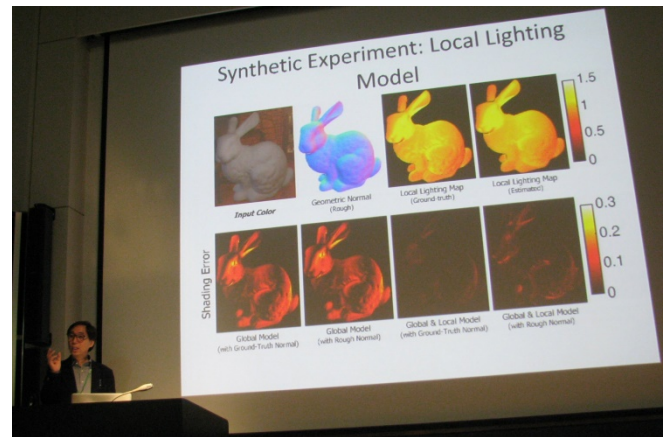
In So Kweon: Per-pixel representation for accurate geometry recovery
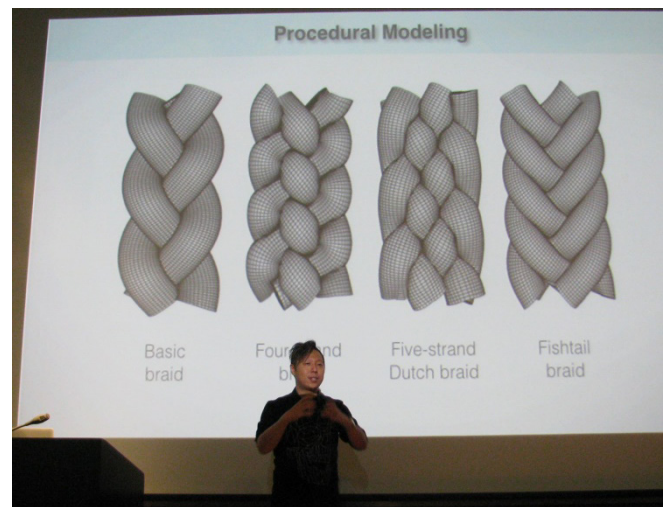
Hao Li: Democratizing 3D human capture: Getting Hairy!

Further, In So Kweon from KAIST (Korea) talked about the difficulties in obtaining high resolution un-textured models suitable for such unforgiving applications as 3D printing. Last but not least, Hao Li from the University of Southern California (USA) gave a talk with a provocative title "Democratizing 3D Human Capture: Getting Hairy!". Numerous interesting projects were discussed in the talk. The most memorable research problem shown was about modelling the structure of human hair, both loose and in plaits, and its motion patterns. The knowledge could be used in the entertainment industry to improve the realism of virtual character appearance, something that the industry continuously strives to do as the human eye remains hard to fool.

'Best Paper Runner-Up' award went to David Ferstl et al. for "aTGV-SF – dense variational scene flow through projective warping and higher order regularisation".

The topics covered in the oral and poster sessions of the main conference were highly representative of the key interests of the 3D vision community of today. Represented were the time-honoured multi-view stereo brought to ever higher levels of achievement, the ever so popular RGB-D methods and the emerging research cluster explicitly targeting reflectance complexities during reconstruction. Further, there was work on semantic analysis of 3D captures involving scene flow computation, tracking and recognition. The 'Best Paper Runner-Up' award this year went to the paper "aTGV-SF – dense variational scene flow through projective warping and higher order regularisation" by David Ferstl et al. (Graz University of Technology, Austria) from the *Flow and Motion* session. Lastly, work on registration, pose estimation, calibration and 3D representation could not be omitted from the scope of 3DV. The versatility in the scope of topics at 3DV is a testament of the expansion the 3D vision field has been undergoing in the last couple of decades and its usefulness in a large variety of applications. Let us hope 3D vision will continue to thrive and flourish for many years to come.

At the izakaya

When writing about a successful conference, one cannot neglect to mention the social event. This year's 3DV banquet was held at a Japanese izakaya En-Ya. An izakaya is a Japanese pub – an establishment where both food and alcoholic beverages are served. As is customary in Japanese households, the banquet guests were asked to remove their shoes and were seated on the tatami mats on the floor. It felt a bit awkward and uncomfortable at first but the feast that followed made it all worth it. The guests were treated to traditional Japanese dishes pre-ordered by the conference organisers, some of which were cooked directly on the table. In order to independently explore the exotic oriental cuisine some resorted to the help of volunteer translators since no English menu could be produced especially for the conference, the options being too numerous. Naturally, the participants indulged in sake and Japanese beer – fully earned after 4 conference days – and the sounds of toasting constantly rang through the air.

All in all, 3DV 2014 was a great success, a memorable experience both academically and culturally. After months of hard work with the satisfaction of a job well done, the organisers could pass the 3DV baton to the organising committee of 3DV 2015. The announcement of 3DV 2015 was made at the closing ceremony. In 2015 the conference will be held in Lyon, France. A lot has been promised to prospective participants: a beautiful setting, excellent French cuisine at a stylish restaurant, less of a jetlag for European participants, and of course the best and the latest research of the 3D vision world! I would strongly encourage researchers working in the field to consider submitting to 3DV 2015.

Let me conclude by thanking BMVA for sponsoring my conference trip to Japan. Should you have any questions about 3DV 2014, please don't hesitate to contact me.

Nadejda Roubtsova[2]
University of Surrey
email: n.roubtsova@surrey.ac.uk

Vision for Robotics – Meeting Report

One-day BMVA Symposium held in London on Wednesday 3 December 2014.

Chairs: Marco Paladini and John Illingworth

BMVA Meetings seem to have risen in popularity since I last attended one. With 80+ attendees from both academia and industry, the Vision for Robotics meeting was not only a success but also a very good platform for discussion and networking.

At 10 am, shortly after everyone had registered and had their caffeine fix, the meeting started. An introduction by chair John Illingworth quickly led into the first keynote speaker. Andrew Davison (Imperial) gave a very interesting talk on "Scene understanding for domestic robots", where he presented some of the work that has been done at Imperial's Dyson Robotics Laboratory. After a brief introduction to Simultaneous localisation and mapping (SLAM), Davison

---

[2]Presenter of the paper "Colour Helmholtz stereopsis for reconstruction of complex dynamic scenes" at 3DV 2014.

showed a very sleek presentation of the Dyson360. The talk then organically moved onto SLAM frontends and finished off with a very interesting demonstration of SLAM using event-based cameras. Jeremy Wyatt (University of Birmingham) then gave an interesting talk about compositional hierarchy and its use in visual recognition to guide dextrous gasps. Margarita Chli (University of Edinburgh) talked about the use of vision systems for UAV navigation. She discussed the fundamental aspects of a SLAM system and gave a short introduction to the sFly Project, which was able to have autonomous UAVs flying and mapping an environment. Paul Siebert (University of Glasgow) finished the morning session with a presentation on a robot that is able to perform dextrous manipulation of clothing. Dextrous Blue (the robot) is able to sort clothes by color, identify the type of clothing, unfold and fold garments.

A busy discussion-led coffee break

Lunch allowed some much needed relaxation and food, as well as important time to talk to other researchers and network. The afternoon session started with a very interesting talk by Bruno Siciliano (University of Naples Federico II), who discussed Visual Servoing for industrial applications, monocular ball catching, the ARCAS Project, and some very cool tracking of deformable materials with a demo using pizza dough. Diego Boesel (CSEM) followed with a similar talk, except his Visual Servoing involved much larger scales and moving robotics platforms. Zeeshan Zia (Imperial) moved the conversation back to SLAM. He proposed SLAMBench, a multi-platform, multi-language benchmark platform and language (PAMELA). This session ended with Johannes Hiltner (MVTec), who gave an industrial talk about HALCON, an off-the-shelf image processing library and prototyping tool.

A coffee break followed, along with the last session of the day. The keynote, Paul Newman (University of Oxford), gave a presentation about Vast-Scale Localisation (Vast-Scale meaning space AND time). His main goal is to become extremely good at localising in a well mapped area. He has hours of footage of the same trajectory in different weathers, lighting and even traffic levels. This allows his localisation algorithm to be illumination invariant, suppress distractions (such as busses) and use 'place dependant' features to make weak localisations. During all this, he made the point that with how cheap memory is, storing large maps is no longer prohibitive. Pedro Cavestany (Cranfield University) followed, he presented his work on low-cost ground-based robots (Rovio) that perform collaborative Structure from Motion (SfM). The event ended with Frederic Fol Leymarie, Prashant Aparajeya and Daniel Berio (Goldsmiths University) who showed us a very well structured and thought out way of doing human-robot gesture transfer. They demonstrated it by tracking the motions of a graffiti artist and reproducing them using a robot hand.

As the day ended, people gathered outside the seminar room and continued to discuss their work. We were soon told the discussions could carry on at the pub, where we headed for drinks.

Oscar Mendez
University of Surrey
email: o.mendezmaldonado@surrey.ac.uk

STOP PRESS – Dundee MIA Workshop

The Dundee 1-day Medical Image Analysis Workshop ran on 27 March and was a success. It was organised under the auspices of sicsa (the Scottish Informatics and Computer Science Alliance).

A PDF of the book of abstracts was produced for the MIA Workshop (~13 MB) and can be downloaded from the following link:

http://mia.computing.dundee.ac.uk/wp-content/uploads/2015/03/Dundee_MIA_Workshop_2015_Book_of_Abstracts.pdf

Andrew McNeil
University of Dundee
email: aymcneil@dundee.ac.uk

BMVA Election

Nominations are requested for the forthcoming election of BMVA Executive Committee members.

We are looking for new members with the ideas and time needed to help enhance our research community. Nominees must be paid-up members of the Association and agree to serve for a period of two years. Members of the Committee are expected to take an active role in BMVA activities. They must also participate in the bimonthly committee meetings taking place at locations deemed mutually most convenient to committee members.

Completed nomination forms should be sent to the Secretary at the address below and must be received by 20 June. Each form must be signed by the individual standing and by one other BMVA member.

The BMVA Committee has ten elected members, five of whom are elected each year. Details can be found at:

http://www.bmva.org/w/executive_committee

Should we receive 6 or more nominations it will be necessary to hold elections. Candidates will then be asked to send a 200-word statement to *helen.cooper@surrey.ac.uk*, summarising relevant experience and interest in BMVA activities, for distribution to BMVA members, A postal ballot will be held, with voting papers returned by 20 August. The results will be announced at this year's BMVC, and in BMVA News.

BMVA Committee Nominations 2015

If you wish to be nominated for the BMVA Executive Committee, please fill in the following application form and return it either by electronic signed copy (*helen.cooper@surrey.ac.uk*) or by post to:

Dr Helen Cooper (BMVA Secretary)
Centre for Vision Speech and Signal Processing
Faculty of Engineering and Physical Sciences
University of Surrey
Guildford
Surrey, GU2 7XH, UK

————————————————fold———————————— here————————————————

Nomination

I wish to be nominated for the BMVA Executive Committee

name: _____ date of birth: _____

signature: _____

email address: _____

BMVA registration no.: _____

I hereby nominate the above BMVA member for a place on the BMVA Executive Committee

name: _____

signature: _____

BMVA registration no.: _____

(Lost/forgotten registration numbers can be obtained upon request from *helen.cooper@surrey.ac.uk*.)