

BMVA News

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

Volume 27 Number 3
March 2017

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



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

BMVA News¹ is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the Editor by 10 June 2017.

Contents

Editorial: <i>Machine Vision in Historical Perspective</i>	1
Meeting Report on Plants in Computer Vision	2
Report on BMVA Symposium: Analysis and Processing of RGBD Data	3
Security and Surveillance	4
Dynamic Scene Reconstruction	5
BMVA Symposium: Computer Vision in Cancer.....	5
BMVA Computer Vision Summer School, 2017	6
BMVC 2017 – Call for Papers.....	6
Report on S+SSPR 2016, Mexico.....	7
Report on ACCV16 Assistive Vision Workshop.....	9
MIUA 2017	10
8 th International Conference on PR Systems.....	11
Lourdes Agapito promoted to a Chair!	12
The Southampton–York Natural Scenes Dataset	12
Nominations for the BMVA Executive Committee	14
Nomination Form for the BMVA Executive Committee	14

Editorial: *Machine Vision in Historical Perspective*

My late grandfather, Roy Davies, left me with trunks of old papers and books from which I have been trying to understand the state of Computer Vision in the 2010s and its meandering later development. Early on, the subject grew

out of image processing and there was little scientific methodology, algorithms being produced on a rather ad hoc basis: some worked reasonably, though many worked unreliably and non-robustly. With the extensive use of the Hough transform, RANSAC, robust statistics and other methods, things improved. However, from the onset, 3D vision developed – as it had to – on a stricter mathematical basis. For a while, progress with vision was dominated by the production of conventional algorithms reflecting what workers felt the subject ought to be like: thus there was a concentration on the structure of images, and algorithms had to proceed logically from image acquisition to feature detection to feature collation to high-level interpretation – albeit there was an understanding that downflowing information was also important: as with the human eye, queries were expressed by the brain and addressed by a saccadic interrogation system. Nevertheless, the whole architecture was one of logical progression and conventional algorithms. To be sure, there was an interlude in the 1980s when artificial neural networks seemed to be taking over, though this approach gradually fell into disrepute because their internal workings were not discernible, so people felt that they could not be trusted to work in real applications. The onus was still on conventional algorithms. But there was a limit to all this, as visual data is often too fuzzy and confused to make it possible always to rely on the logical feature detection to collation to interpretation route: if you assume it is nicely structured then you can't even tackle the more difficult problems such as medical imaging and novelty detection, and how to follow an infinitely variable 3D shape such as a shark as it flies through the air.

Ultimately, it all has to be done by training, but in a grossly more mature way than for the old ANNs. Still, CNNs came to the rescue (machine learning optimisation approaches were also used, but space precludes a discussion of this alternative here), and vastly more complex capabilities arose without the need for conventional algorithms. This time, people were less worried about trainable methods being unsafe, because huge image datasets were being employed; because deconvolution methods had been devised for getting an appreciation that these methods did in fact have a solid enough grounding to be reliable; and because it could be seen that new performance levels were being achieved in a quite wide set

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.


of applications, including object recognition, face recognition, and semantic scene segmentation. Indeed, the euphoria of workers around 2012–2017 was so great that papers on deep learning exceeded those describing conventional algorithms at most vision conferences. Subsequently, conventional work more or less died out, though not without huge objections from many, who felt that too few principles were being enunciated for the subject to be thought of as a scientific discipline (most workers had swung over to a mood of ‘if it works, don’t knock it’). After all, a fully trained system is not an engineered system, with proper specifications and rigorous proofs of success – though on the other hand, with really complex data that can’t be described in anything like complete detail, how do we even know that there is a solution and how to specify it?

Funnily enough, it took another 16 years or so – until the mid-2030s – for this phase to pass and for the real science to return. For at that stage, it was increasingly found that the latest deep, deep networks, and in particular those including substantial feedback – deep, deep recurrent networks – were starting to automatically restructure themselves into more modular systems. Closer examination showed from their by then automatically generated architectures that they had in fact taken on something of the appearance and arrangement of parts of the human brain. There were fascinating parallels and fascinating differences. A whole panoply of visual areas and visual pathways that might be expected was present, though not identically or in identical numbers or even identically connected, but the overall layout nevertheless showed remarkable parallels. From the structure, one could imagine identifying most areas of the machine visual cortex, from V1 right up to MT (V5) – together with branches into the dorsal–ventral pathways. The lateral geniculate nucleus was generally conspicuous by its absence, though in the end it was found that this appeared when the eyes were made to operate in saccades in conjunction with the eye muscles. Next, various memory areas were connected to the purely visual areas: in the end these were related to short term memory in the hippocampus and prefrontal cortex. By the 2040s, structures concerned with long term memory and some apparently akin to the perirhinal cortex were identified.

Perhaps the disappointing aspect was the application of vision to vehicle guidance. In the 2010s, automatic vehicle guidance systems were in vogue, with competing designs by many prominent manufacturers, and the 2020s were poised for a big changeover – with humans naturally retaining the right to take over from the machine whenever they felt the need. However, the growth in the number of vehicles was steadily bringing road transport to a standstill, with the well-known frustrating relaxation oscillations in speed on motorways. By the late 2030s all the successes computer vision had achieved with vehicle control were suddenly eliminated. Vehicles now had to be managed into convoys in order to maintain anything like sensible flow rates. It was a scary period, as vehicles were typically separated by as little as 30 cm, and moved automatically from lane to lane and from one road to another under complete programme control. Things would never be the same again, in spite of the successes brought by computer vision. Nowadays, many of us prefer to travel by train, which doesn’t rely on vision at all ...

I wish my grandfather could have read this article.

Meeting Report on Plants in Computer Vision



This one-day BMVA symposium was held in London on 16 November 2016.

Chair: Hannah Dee

Introduction

In November I had the opportunity to attend a one-day BMVA workshop on plants in computer vision. As I have been working at a purpose-built plant phenotyping facility for the last few months, I found it to be extremely informative. In particular, it was useful to hear from many other researchers and find out how they have set about tackling and solving problems that plague so many of us.

Hanno Scharf (Forschungszentrum Jülich, Germany) opened the BMVA talks by highlighting the importance of purpose-built facilities, techniques and equipment for high-throughput plant science.

With so many industries depending on the 4 Fs (Food, Fuel, Feed and Fibre) there is huge pressure on the efficiency and accuracy of this branch of research. Many of the speakers reiterated just how vital computer vision, machine learning and automation are for continuous progression in research.

With the majority of research in plant phenomics requiring experiments to be done en masse, it seems that each project has come up with its own custom solutions for capture and analysis. A lot of facilities are making use of ‘LemnaTec’, ‘Elcom’ or similar systems for their automation and image capture and are augmenting these systems to fit their needs. A common adaptation presenters had made to these pre-built systems was an upgrade to higher resolution cameras, which produce larger but clearer data to work with.

Classification problems

Whilst I was aware of issues with classifications made based on visual analysis in the lab, I had not fully appreciated the problem until listening to Norman MacLeod’s talk. The ‘Climate leaf analysis multivariate program’ (CLAMP) system is used to identify a leaf by its physical characteristics. By sending samples to different labs and then comparing results, Professor MacLeod was able to comment on the variability of expert judgements. This talk, in particular, highlighted the possibly controversial issue that two experts may identify the same thing differently. As MacLeod pointed out, this is an issue that can be tackled by

the use of computer vision, in the form of an unbiased system for routinely and consistently evaluating leaves or other living matter.

Using machine learning to aid ID of plants

In order to extract useful information from images, it is vital to clearly identify regions of a plant and to segment them correctly. This is exactly what Drs Hannah Dee and Jonathan Bell have been using machine learning techniques to achieve. Using *Arabidopsis* they have used convolutional neural networks in order to separate and uniquely identify each individual leaf of a plant throughout its life. This allows for the examination of any particular leaf at any stage in growth, and to compare it with any other sample.

Building 3D models

As automation overtakes manual tasks in working physically with plants, precision becomes an increasingly important factor. Without this accuracy we would have machines destroying plants accidentally. This emphasises the importance of building accurate, usable and manipulable 3D models for use in directing machines towards their targets i.e., individual leaves, stalks or stems.

Dr Rick van de Zedde's team have done this very successfully with a multi-camera approach in order to build 3D models on the fly, allowing seedlings to be rapidly measured, sorted, and processed within milliseconds. Using multiple cameras and a volumetric subtraction technique to carve away anything from the model which is part of the background, they develop a 3D model enabling the seedlings to be assessed. Van de Zedde also presented work currently in use automatically propagating roses by cutting away unnecessary parts of the plant and allowing the stem to grow successfully unhindered. In addition to performing clever maintenance for plants, this system is able to capture visual data along with physical tissue information by tactfully sampling leaves without harming the plant.

A related 3D approach came from Jonathan Gibbs of Nottingham. His "Active Vision Cell" uses a robotic arm with 6 degrees of freedom and intelligently captures enough images of a plant, from as many angles as it requires in order to fully create a 3D model without any gaps.

It's always interesting and informative to hear of other fields' take on computer vision and biology. Dr Marco Aita – a physicist by trade – has been working on mapping the circumnutation in plants over time. What's most interesting about this is the breadth of past research that has been done on the subject, yet how little is still known about the seemingly random twists and turns of a plant during growth. Dr Aita has been using a dual camera set-up to reconstruct this phenomenon in 3D.

Hyper-spectral imaging

Another interesting technique is the use of hyper- and multi-spectral imaging to filter out or investigate spectra of light beyond the visible. This appears to have huge potential in field work on larger scales, as Hanno Scharr said in his keynote. Being able to process entire fields of crops at once could allow identification of diseases, pests or malnutrition, effectively reducing the need for wide-scale treatments and allowing for smaller more precise treatments. As Jonathan

Gibbs pointed out, hyper-spectral imaging on a plant-per-plant basis has largely been ignored and could allow for easier depth perception in single camera image analysis. However, a lot of work would be needed on identification of spectra of interest per species. Gytis Bernotas used spectral imaging to investigate photometric stereo, and Dominic Williams has found success in identifying raspberry plants in the field using 400–1000 nanometre (visible/near-infrared) cameras.

These approaches show that looking outside the visible and combining information from multiple wavelengths can help with plant imaging.

Take away notes

It is clear that this field of research is still very much at the blossoming stage: there are so many avenues yet unexplored and a lot of excellent work being produced. As the divide between Biology and Computer Science becomes blurred, it is evident that both fields can greatly help to improve the other.

Hannah Dee
Aberystwyth University
email: hmd1@aber.ac.uk

Report on BMVA Symposium: Analysis and Processing of RGB-D Data



This one-day BMVA meeting was held at the BCS, London on Wednesday 1 March 2017.

Chairs: Paul Rosin and Yukun Lai, Cardiff University

This BMVA Technical meeting covered a wide range of topics within the field of RGB-D research, some presenters focusing on the fundamental challenges of RGB-D systems and others detailing complex systems enhanced by RGB-D technology.

Amir Atapour-Abarghouei proposed an effective method to fill in holes found in RGB-D images; these holes are primarily caused by limitations in current RGB-D capture systems. Although not infallible, the algorithm presented was capable of producing results that were far more plausible than those of other modern algorithms in a variety of different challenging scenarios. Through a set of real world and synthetic benchmarks, Amir demonstrated that

his method was capable of accurately correcting blemishes at the edges of objects, as well as convincingly recovering larger holes within an image. Current RGB-D camera limitations appear to be here to stay, so it is foreseeable that hole filling methods like this will play a crucial role in future RGB-D processing pipelines.

I found Victor Adrian Prisacariu's morning talk of particular interest, presenting a system capable of large-scale 3D scene reconstruction with loop closure. Prisacariu conducted a live demonstration of his 3D reconstruction algorithm running on a tablet device, which – for me – helped exemplify how much commodity hardware has advanced in recent years and how accessible such technologies are now.

Just after lunch, Dr Dima Damen gave a fascinating talk on the challenges and opportunities of RGB-D action/activity recognition in the context of collaborating with several research projects, including SPHERE – a large healthcare project at Bristol University. Covering opportunities in RGB-D technology such as the automation of healthcare monitoring for elderly patients and the integration of instructional systems to help workers in an industrial setting. Dr Damen covered the challenges involved in implementing each of these systems in their respective environments. In the case of the industrial system, the outstanding problem was the cumbersome implementation, which required a user to wear a set of wrist straps, a backpack and headgear for the system to function. In relation to activity recognition, Dr Damen highlighted how easily susceptible modern systems are to miscategorising user action due to the difficulties in distinguishing between successfully completed actions and incomplete actions.


Professor Adrian Hilton presented a novel method that was capable of taking a sequence of RGB or RGB-D images and reconstructing dynamic 3D scenes of deforming geometry. In the capture scenarios presented, Professor Hilton was able to reconstruct intricate 3D motions. This was done by accurately segmenting the relevant objects from the background in each frame, then attempting to reconstruct a 3D model. In the examples shown the backgrounds tended to be static, with no occlusions covering the subject, it would have been interesting to see how his system would have coped with such challenges. Nevertheless, the system yielded impressive results when processing image sequences containing highly fluid actions and the talk was humorously punctuated by the use of his system to put a pair of sunglasses on a dog.

The final talk of the day was given by Martin Runz, who presented a system capable of reconstructing multiple segmented objects within a single scene. The proposed system handled proficiently the movement of rigid objects within a scene; however, it showed there was still room for improvement as the system appeared to be incapable of handling non-rigid deformations.

The series of presentations gave me invaluable insights into the current state-of-the-art in RGB-D based systems. It demonstrated not only the capacity of current technologies, but also the limitations of such systems, with their creators openly discussing the flaws in their systems.

Roberto Dyke
Cardiff University
email: dykerm@cardiff.ac.uk

Security and Surveillance



This one-day BMVA symposium will be held in London on Wednesday 26 April 2017.

Chairs: Thomas Rogers (UCL), James Ferryman (Reading University) and Tim Ellis (Kingston University)

www.bmva.org/meetings


- 09:30 Registration and coffee
- 09:50 Welcome from the chairs
- 10:00 Where to look: intelligent video beyond surveillance: Professor Sean Gong, QMUL
- 10:45 Context-aware person re-identification via fusion of anthropometric and gait features: Athira Nambiar, Instituto Superior Técnico (Portugal)
- 11:15 Coffee Break
- 11:45 Recent methods for maritime object detection and tracking in visible and thermal imagery: Tom Cane, University of Reading/BMT Group Ltd.
- 12:30 Violent behaviour detection: Kaelon Lloyd, Cardiff University
- 13:00 Lunch
- 14:00 Biometric modalities at the border: Chris Hurrey, Intrepid Minds Ltd.
- 14:45 Bio-inspired object recognition framework for video surveillance: Ahmed Al-Obaidi, Staffordshire University
- 15:15 Coffee Break
- 15:45 Automatic object classification for 2D X-ray and 3D CT within an airport and border security context: Toby P. Breckon, Durham University
- 16:30 My text in your handwriting? Tom Haines, UCL
- 17:00 End of meeting.

Registration

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

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

Dynamic Scene Reconstruction



This one-day BMVA symposium will be held in London on 21 June 2017.

Chairs: Jean-Yves Guillemaut, Armin Mustafa, Marco Volino University of Surrey.

Keynote Speakers

Dr Michael Zollhöfer, Max Planck Institute Informatics
 Dr Jonathan Starck, The Foundry
 Professor Lourdes Agapito, UCL.

Call for Papers

Reconstruction of general dynamic scenes is motivated by potential applications in film and broadcast production together with the ultimate goal of automatic understanding of real-world scenes. With advances in hardware and the advent of virtual and augmented reality, dynamic scene reconstruction is being applied to more complex scenes. We welcome contributions to this workshop in the form of oral presentations, posters and demos. Suggested topics include, but are not limited to:

- Dynamic 3D reconstruction from single, stereo or multiple views
- Multi-modal dynamic scene modelling (RGBD, LIDAR, 360 video, light-field)
- 4D reconstruction and modelling
- 3D segmentation and recognition
- 3D/4D data acquisition, representation, compression and transmission
- Scene analysis and understanding
- Structure from motion, camera calibration and pose estimation
- Geometry processing
- Computational photography
- Appearance and reflectance modelling
- Scene modelling in the wild
- Applications of dynamic scene reconstruction (virtual/augmented/mixed reality, character animation, free-viewpoint video, relighting, medical imaging, creative content production, HCI).

This list is not exhaustive, so if you would like to present on anything related to dynamic scene reconstruction do submit an abstract. The work can be “in progress”, recently published, or novel research. We welcome submissions from both industry and academia, including interdisciplinary work and work from those outside of the mainstream computer vision community.

Submissions

Please submit your abstract at
<https://goo.gl/forms/wJckD5pOIWq4s7532>
 The deadline for submissions is 12 April 2017.

If you have any questions feel free to contact
m.volino@surrey.ac.uk

Registration

Book online at www.bmva.org/meetings
 £16 for BMVA Members, £36 for non-Members, including lunch and BMVA Membership.

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

BMVA Symposium: Computer Vision in Cancer



This one-day BMVA symposium will be held in London on Wednesday 11 October 2017.

Chairs: Greg Slabaugh and Constantino Carlos Reyes-Aldasoro (City University, London)

Keynote speakers: Julia Schabel (King’s College London), Nasir Rajpoot (Warwick) and Yinyin Yuan (Institute for Cancer Research)

www.bmva.org/meetings

Call for Papers

This BMVA one-day meeting will present state-of-the-art developments in Computer Vision applied to Cancer Data Analysis.

Recent progress in imaging hardware, acquisition techniques, and algorithmic processing of data have led to advances in detection, diagnosis, staging, treatment, and follow-up in cancer-related clinical workflows, as well as fundamental understanding of cancer modelling and dynamics. Cancer imaging includes varied modalities, and numerous scales including nano, micro, and macro.

This one-day meeting will be dedicated to technical advances that have the potential for clinical relevance, and seeks to bring together a collection of recently-developed approaches in this domain. We hope the methods presented will inspire future research, from both theoretical and practical viewpoints, to spur further advances in the field.

The topics include, but are not limited to:

- segmentation in cancer imaging: from cellular structures to cells to lesions to tumours
- tracking of cells in metastasis and migration processes
- registration of cancer images
- histopathology image analysis
- modelling of cancer cells, vasculature and cancerous processes
- image-based interventional techniques for cancer treatment
- tissue characterisation from images.

Submission Deadline

All those interested in presenting at this meeting are invited to submit a summary of their talk by 31 July 2017 [firm deadline]: see <https://goo.gl/forms/q3PzYHeyEtkCsYsL2> for submission details.

Invitation of Papers

Authors of selected papers will be invited to submit extended versions to be included in a special issue of IET Computer Vision.

Registration

Book online at www.bmva.org/meetings
 £16 for BMVA Members, £36 for non-Members, including lunch and BMVA Membership.

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

BMVA Computer Vision Summer School, 2017



UNIVERSITY OF
 LINCOLN

University of Lincoln, 3–7 July 2017

BMVA runs an annual Computer Vision Summer School aimed at PhD students in their first year, though it is normally also 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 2017 Summer School will take place at the University of Lincoln on 3–7 July. It will consist of an intensive week of lectures and lab sessions covering a wide range of topics in Computer Vision. Lecturers are researchers in fields from some of the most active research groups in the UK and abroad.

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.

For further details, see the summer school website and the list of speakers, or contact us by email:

<http://cvss.blogs.lincoln.ac.uk>
<http://cvss.blogs.lincoln.ac.uk/speakers/>
 email: cvss2017@lincoln.ac.uk

Dr Nicola Bellotto
 University of Lincoln
 email: nbellotto@lincoln.ac.uk

BMVC 2017 – Call for Papers

BMVC 2017 will be held at ICL, London on 4–7 September 2017: See the website below for the full details of the conference and the Call for Papers.

<http://bmvc2017.london/>

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).

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. Accepted papers will be included in the conference proceedings published and DoI indexed by the BMVA. Please note that BMVC is a single-track meeting with oral and poster presentations and will include two keynote presentations and one tutorial.

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.

Keynote speakers

Richard Szeliski, Facebook, USA
<http://szeliski.org/richardszeliski.htm>
 Pietro Perona, California Institute of Technology, USA
<https://www.vision.caltech.edu/perona.html>

Tutorial speakers

Michael M Bronstein, USI Lugano/Tel Aviv
 University/Intel
<http://www.inf.usi.ch/bronstein/>
 Lourdes Agapito, UCL
<http://www0.cs.ucl.ac.uk/staff/l.agapito/>
 Shimon Whiteson, University of Oxford
<http://www.cs.ox.ac.uk/people/shimon.whiteson/>

Workshop keynote

Jamie Shotton, HoloLens Science Team at Microsoft
 Cambridge, <https://www.microsoft.com/en-us/research/people/jamiesho/>

Important dates

Abstract deadline	Tuesday 25 April
Submission deadline	Tuesday 2 May [†]
Author notification	4 July
Camera-ready deadline	18 July
Author registration deadline	18 July
Conference tutorials	Monday 4 Sept
Main conference	Tu–Th, 5–7 Sept
Workshops	Thursday 7 Sept

[†]There will be no further extension.

Submission guidelines

Paper submission and registration both use CMT. Paper registration is performed by registering as a user with CMT and entering a paper title and abstract. This will result in the allocation of a “paper ID”, which indicates that registration is completed and should be used in preparation of the review copy. You will be able to make edits and upload new paper drafts until the final deadline. Submitted papers should be prepared according to the published specification for formatting and style. Please be sure to read both the formatting instructions and policies before submission.

See the following page for major details:
<http://bmvc2017.london/paper-submission/>

Chairs and organisers*General chairs*

Tae-Kyun Kim, ICL; Stefanos Zafeiriou, ICL

Program chairs

Krystian Mikolajczyk, ICL; Gabriel Brostow, UCL

Advisory board

Andrew Davison, ICL; Maja Pantic, ICL;

Yiannis Demiris, ICL

Publicity chair

Giovanni Maria Farinella, University of Catania, It

Tutorial/workshop chairs

Yannis Panagakis, ICL; Tania Stathaki, ICL

Sponsorship/demo chairs

Ben Glocker, ICL; Stefan Leutenegger, ICL

Local arrangements chairs

Rigas Kouskouridas, ICL; Zhiyuan Shi, ICL;

Anastasios Roussos, ICL

Website chairs

James Booth, ICL; Patrick Snape, ICL.

Sponsorship

If you are interested in sponsoring BMVC 2017, would like to discuss a custom package, or require more information, please contact the sponsorship chair.

More information on the sponsorship packages is available at <http://bmvc2017.london/sponsorship/>

Dr Tae-Kyun Kim

ICL

email: tk.kim@imperial.ac.uk

Report on S+SSPR 2016, Mexico

The joint IAPR International Workshops on Structural and Syntactic Pattern Recognition (SSPR 2016) and Statistical Techniques in Pattern Recognition (SPR 2016) will be held in Mérida, Mexico from 30 November to 2 December. The location of the conference is in Mérida as it is one of the oldest cities in the Americas, with Mayan, Spanish and French influences.

S+SSPR 2016 is a joint biannual event organized by Technical Committee 1 (Statistical Pattern Recognition Techniques) and Technical Committee 2 (Structural and Syntactical Pattern Recognition) of the International Association of Pattern Recognition (IAPR). Following the trend of previous editions, S+SSPR 2016 was held in close proximity to, and shortly before, the International Conference on Pattern Recognition (ICPR).

The main conference offered an abundance of keynote presentations organized in thematic sessions. A passion for research and vision for ever greater advances of the field in the future characterized all keynote talks.

Professor Hamid Krim (North Carolina State University, Raleigh) presented the topic of machine learning with convexity, nullity and sparsity. Hamid’s research interests are in statistical signal and image analysis and mathematical modeling, with a keen emphasis on applied problems.

Richard Wilson (University of York) presented his work entitled “A Jensen-Shannon Divergence Kernel for Directed Graphs”. The speaker introduced a series of graph entropy computations with the Jensen-Shannon Divergence and extended the von Neumann entropy to directed graphs. Furthermore, he established the effective graph kernel on tested datasets.

Edwin Hancock (also from the University of York) presented two talks. One was on graph entropy from closed walks and cycle functions and the other was on thermodynamic characterization of temporal networks.

Edwin analysed the networks or graphs structure from physical point of view and presented the statistical structural characterization in graph-theoretic methods.



Edwin Hancock: Graph Entropy from Closed Walks and Cycle Functions; Thermodynamic Characterization of Temporal Networks.

Professor Mário Figueiredo from Instituto Superior Técnico (IST) was given the Pierre Devijver Award at the conference. He presented his state-of-the-art research on unsupervised selection and grouping of correlated features. Mário developed new methods for fitting the mixture models using the EM algorithm. Mário's group continues to formulate new challenges for image processing and analysis, machine learning and optimization. He is now a Fellow of the IEEE and of the IAPR, and since 2014 he has been included in the Thomson Reuters' Highly Cited Researchers list.

Professor Sudeep Sarkar from University of South Florida was invited to give a talk about pattern theory for representation and inference of semantic structures in videos. Sudeep has more than 25 years of experience conducting and directing fundamental and applied research in computer vision, image processing and pattern recognition. He holds four US patents and has published more than 200 high-impact journal and conference papers, as well as notable books and book chapters.



Sudeep Sarkar: Pattern Theory for Representation and Inference of Semantic Structures in Videos.

When writing about a successful conference, one cannot neglect to mention the social event. This year the conference organized a tour at Sotuta de Peón where there was an ancient farm full of Mexico culture. Sotuta de Peón is the restoration project of a landmark located in the heart of the ancient henequen zone in the Yucatan state, and gives a true glimpse of what was once a fully operational Henequen Hacienda in the grand style and tradition of this period. When touring Sotuta de Peón, the guests witnessed the step-by-step process of the transformation of the Henequen (or Green Gold as it was known throughout the region), from plant to fibre and from fibre to finished product.



Tour: Main House in Sotuta de Peón.

The tour also included a visit to a typical Mayan House and involved the opportunity to see the antique machinery that has been rescued from the ravages of time and rebuilt to precise operational capacity. The underground rivers with crystal clear water form a natural swimming pool surrounded by lush gardens of tropical flowers and brilliant green trees and indigenous plants. Naturally, the participants indulged in beautiful views and Mexico food.



Conference Meeting Room: Intercontinental Presidente "Villa Mercedes" Hotel in Mérida

Overall, S+SSPR 2016 was a great success and 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 baton to the organising committee of S+SSPR 2018. In 2018 the conference will be held in China. A lot has been promised to prospective

participants: a beautiful setting, excellent Chinese cuisine at a stylish restaurant and of course the best and the latest research in statistical and structural pattern recognition. Nowadays, S+SSPR has been ranked ‘A’ in computer science conferences by the Computing Research and Education Association of Australasia, CORE. I would strongly encourage researchers working in the field to consider submitting to S+SSPR 2018.

Let me conclude by thanking BMVA for sponsoring my conference trip to Mexico. Should you have any questions about S+SSPR 2016, please don’t hesitate to contact me.

Jianjia Wang
 University of York
 email: jw1157@york.ac.uk

Report on ACCV16 Assistive Vision Workshop

This year the 13th Asian Conference on Computer Vision (ACCV16) was held in Taipei on 20–24 November. The conference took place in the Taipei International Convention Center (TICC), a perfect place for such events located next to Taipei 101, the tallest and largest green building in the world. This biennial conference and its workshops provide a high level forum in computer vision and related areas that gathers researchers, developers and practitioners from diverse regions of Asia, Europe and North America.



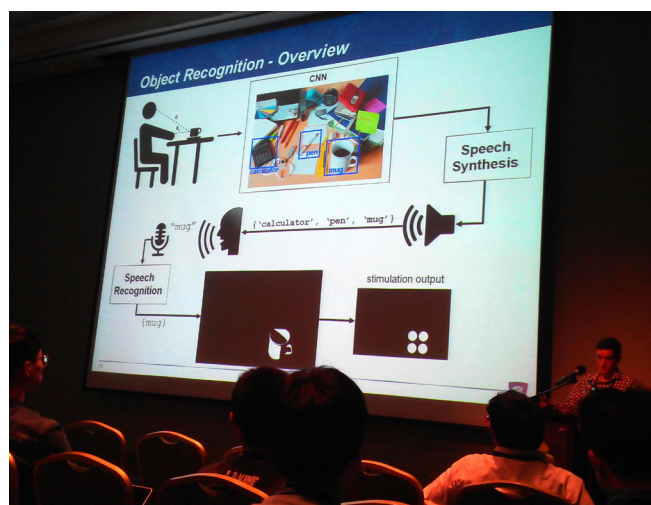
The workshop on Assistive Vision was held on 20 November 2016. This workshop organized by Chetan Arora, Vineeth N Balasubmanian, CV Jawahar, Vinay P Namboodiri and Romanathan Subramanian, focuses on researches that use vision technologies to improve the day-to-day lives of physical and cognitively impaired people.

The Assistive Vision workshop included three really interesting and inspiring invited talks. Manohar Paluri from Facebook presented the assistive vision topic from a company point of view. The concept ‘scale’ was the key of the presentation, where ‘scale’ represents the huge diverse amount of data contained on Facebook and the complexity of dealing with it. Manohar Paluri also introduced some of the recent projects that Facebook is implementing for assistive living, such as ‘automatic alternative text’, a tool that reads/describes images for blind people. The second invited talk was presented by Kris Kitani from Carnegie

Mellon University. Kris Kitani introduced the topic of understanding the interaction with things from a first person viewpoint, remarking on the importance of hand detection and hand position in order to understand the scene better. He also presented some of the projects of his research group such as ‘NavCog’, an iPhone navigation application for blind people. Finally, Min Sun from NTHU presented their work anticipating accidents in Dashcam videos. Min Sun presented their accident dataset recorded in Taiwan, which includes 620 videos and their Dynamic-Spatial-Attention RNN model for predicting accidents.



The topics covered by the presenters of the workshop represented some of the current approaches and technologies for assistive living. For example, the work of Wei-Ta Chu and Tsung-Han Yang demonstrated a good method for applying patterns to images to help Dichromats to recognise colours. Lili Tao from Bristol University presented their work on calorie counter using RGB-Depth data. This work is a small part of their Sphere project, a sensor based platform for assisting living at homes. Hamed R.-Tavakoli presented an interesting fixation prediction method based on unsupervised hierarchical models. Xiaoli Sun, invited from



the main conference, presented their work in saliency detection using diversity-induced multi-view matrix decomposition. Finally, Derek Rollend presented their work for face and object detection for a retinal prosthesis, showing some really interesting examples of blind people

using the prosthesis being able to distinguish different objects from a table.

Despite the fact that one of the presenters was absent due to sickness, the workshop time was covered. The last part of the workshop was dedicated to a panel discussion where two of the invited speakers and two of the paper presenters answered different questions from the audience. This final discussion was very interesting, covering different questions about computer vision and its future. They also talked about the challenges to overcome during the research, such as privacy/ethics problems and the huge and diverse amounts of data. Also, a comparison between academic and industrial research was provided, where the unequal resources of the industrial and academic sectors were remarked upon. On the other hand, Manohar Paluri mentioned the importance of both since academy provides novel ideas and industry exploits their resources to improve existing methods.



At the end of the first day of the conference, the welcome reception was held in the Courtyard by Marriot Taipei, providing another opportunity to network with all the conference assistants while tasting different types of food, desserts and drinks. In general, the success of the ACCV Assistive Vision Workshop was due to the quality of the selected papers, to the researchers themselves and to the organizations that allowed people to share knowledge and experience on the computer vision area.

Let me conclude by thanking BMVA for sponsoring my conference trip to Taiwan. Should you have any questions about ACCV16 Assistive Vision Workshop, please don't hesitate to contact me.

Juan Manuel Fernandez Montenegro
Kingston University
email: J.Fernandezmontenegro@kingston.ac.uk

MIUA 2017

The 21st Medical Image Understanding and Analysis Conference, MIUA 2017, will take place in Edinburgh on 11–13 July 2017.

The MIUA series of conferences is organised in the UK for communicating progress within the biomedical image analysis research community. Its goals are the dissemination and discussion of research in medical image processing and analysis: it also aims to encourage the growth and raise the profile of this multi-disciplinary field. These conferences

feature keynote speakers, tutorials, workshops, and oral and poster presentations.

Honorary Guest Speaker

Professor Sir Michael Brady is to open the Conference.

Invited speakers

- Ingela Nyström: Professor in Visualisation and Head of the Division of Visual Information and Interaction at the Department of Information Technology, and Director of the Centre for Image Analysis in Uppsala University, Sweden; Vice-Chair of the Council for Research Infrastructure in Sweden
- Jinah Park: Professor in the School of Computing at the Korea Advanced Institute of Science and Technology (KAIST, Daejeon, South Korea), Head of the Computer Graphics and Visualization Research Laboratory and Head of KAIST Research Group of Future Emerging Technology on Medical Imaging.
- Daniel Rueckert: Professor of Visual Information Processing and Head of the Biomedical Image Analysis Laboratory at Imperial College London; Fellow of the Royal Academy of Engineering, MICCAI and IEEE Societies.

Tutorials

Dr Constantino Carlos Reyes Aldasoro: Processing and Analysing microscopy images with MATLAB
Dr Reyes Aldasoro is Senior Lecturer in the Department of Electrical and Electronic Engineering at City University London; a Senior Member of the IEEE; and author of *Biomedical Image Analysis Recipes in MATLAB: for Life Scientists and Engineers* (ISBN 978-1-118-65755-3, Wiley-Blackwell).

Dr Konstantinos Kamnitsas: Deep Learning on the analysis of radiological images
Dr Kamnitsas is a Post-Doctoral Researcher in the Department of Computing at the Imperial College London; a former Software Engineer and Microsoft Research Intern; and author of DeepMedic.

Registered attendants will be able to book a 1-day trip to the Scottish Highlands on Friday 14 July. To find how to make the reservation, email MIUA.2017@ed.ac.uk.

More information is available at the conference website:
<https://mua2017.wordpress.com>

Conference topics

- Big data processing
- Clinical and scientific evaluation of imaging studies
- Computer-aided pathology
- Computer-aided radiology
- Computer-assisted surgery
- Data compression and anonymization
- Data fusion
- Decision support
- Discovery of imaging biomarkers
- Human computer interaction

- Image enhancement
- Image interpretation
- Image-guided intervention
- Image formation and reconstruction
- Image perception
- Image registration
- Image segmentation
- Intelligent imaging systems
- Machine learning in imaging
- Modelling and simulation
- Motion analysis
- Multi-modality image analysis
- Pattern and feature recognition
- Protocol development and standardization
- Quantitative image analysis
- Shape analysis
- Software development
- Super-resolution algorithms
- Statistical methods in imaging
- Systematic testing and validation
- Texture analysis
- Time series analyses
- Virtual reality
- Visualisation.

Early Registration Deadline

Please take special note that the early-bird registration deadline is 12 May 2017.

The official MIUA contact address is:

MIUA.2017@ed.ac.uk

For further information about registration and other aspects of the conference, see the conference website:

<https://miua2017.wordpress.com>

Víctor Gonzalez-Castro

University of Edinburgh

email: victor.gonzalez@ed.ac.uk

8th International Conference on PR Systems

11–13 July 2017, Madrid, Spain

Introduction

The 8th International Conference on PR Systems (ICPRS-17) is an annual event that follows ICPRS-16, a continuation of the successful Chilean Conference on PR. In 2017 it is organised by two Madrid universities (UPM and UC3M) and sponsored by the Chilean Association for PR (ACHiRP), the Spanish Association for Artificial Intelligence (AEPIA), the Spanish Association for PR and Image Analysis (AERFIA) and the Vision and Imaging Professional Network of the IET, which will publish the proceedings. Papers deemed to be of the required standard and presented at the conference, will be indexed by INSPEC and by IEEE Xplore and Scopus. All paper submissions will be submitted via ConfTool to be peer-reviewed by an international panel of experts. Authors of excellent papers will be encouraged to submit extended versions for consideration in IET

research journals (*Computer Vision, Image Processing, and Biometrics*). For more information, please visit the conference website: <http://www.icprs.org>

Call for Papers

Interested authors are invited to submit papers describing novel and previously unpublished results on topics including, but not limited to:

- Special session on emotions in groups and crowds
- Artificial intelligence techniques in PR
- Bioinformatics clustering
- Biometrics (including face recognition)
- Computer vision
- Data mining and big data
- Deep learning and neural networks for PR
- Document processing and recognition
- Fuzzy and hybrid techniques in PR
- High performance computing for PR
- Image processing and analysis
- Kernel machines
- Mathematical morphology
- Mathematical theory of PR
- Natural language processing and recognition
- Object detection, tracking and recognition
- PR principles
- Real systems, applications and case studies of PR, e.g., health, environment, transportation
- Robotics
- Remote sensing
- Shape and texture analysis
- Signal processing and analysis
- Social media and HCI
- Statistical PR
- Syntactical and structural PR
- Voice and speech recognition.

Key Dates

The Conference:	11–13 July
<i>Final</i> extended submission deadline:	23 April
Notification of acceptance:	21 May
Camera-ready papers: [†]	4 June

[†]max 6 pp.

Main chairs and organisers

- Technical Chair: Professor Pedro Gómez V, Universidad Politécnica de Madrid, ES
- General Chair: Professor Sergio A Velastin, Universidad Carlos III de Madrid, ES
- Tutorials Chair: MSc Daniel Palacios-Alonso, Universidad Politécnica de Madrid, ES
- ACHIRP Chair: Dr Marcelo Mendoza, Universidad Técnica Federico Santa María, CL
- Local Arrangements: Professor José Manuel Molina, Universidad Carlos III de Madrid.

Professor Sergio A Velastin
University Carlos III Madrid Spain
email: sergio.velastin@ieec.org

Lourdes Agapito promoted to a Chair!



Professor Lourdes Agapito, with (right) Professor John Shawe-Taylor who introduced Lourdes at her inaugural and Professor Andrew Zisserman who gave the vote of thanks.

Following Lourdes' promotion to a Chair in 3D Computer Vision, and her inaugural last October, it seemed appropriate to summarise the nature of her work and to include a short biography.

Lourdes' recent work on the capture of 3D models of the world from video²

As humans we take the ability to perceive the dynamic world around us in three dimensions for granted. From an early age we can grasp an object by adapting our fingers to its 3D shape; we can understand our mother's feelings by interpreting her facial expressions; or we can effortlessly navigate through a busy street. All of these tasks require some internal 3D representation of shape, deformations and motion. Building algorithms that can emulate this level of human 3D perception has proved to be a much harder task than initially anticipated. While some degree of success has been achieved when the scene observed by a camera is static or 'rigid', inferring the 3D geometry of the vivid moving real world is still in its infancy. This challenge has fascinated Lourdes throughout her research career. In this lecture she will show progress from her early systems which captured sparse 3D models with primitive representations of deformation towards our most recent algorithms which can capture every fold and detail of hands, faces and clothes in 3D using as input video sequences taken with a single consumer camera. There is now great short-term potential for commercial uptake of this technology, and Lourdes will show applications to robotics, augmented and virtual reality and minimally invasive surgery.³

A short biography

Professor Lourdes Agapito obtained her BSc, MSc and PhD (1996) degrees from the Universidad Complutense de Madrid (Spain). She held an EU Marie Curie Postdoctoral Fellowship at The University of Oxford's Robotics Research Group before being appointed as a Lecturer at Queen Mary, University of London in 2001. In 2008 she was awarded an

ERC Starting Grant to carry out research on the estimation of 3D models of non-rigid surfaces from monocular video sequences. In July 2013 she joined UCL Computer Science as a Reader (Associate Professor) where she leads a research team that focuses on 3D dynamic scene understanding from video. Lourdes is Program Chair for CVPR 2016, the top annual conference in computer vision; in addition she was Programme Chair for 3DV'14 and Area Chair for CVPR'14, ECCV'14, ACCV'14 and Workshops Chair for ECCV'14. She has been keynote speaker for CVMP'15 and for several workshops associated with the main computer vision conferences (ICCV, CVPR and ECCV). Lourdes is Associate Editor for IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI), a member of the Executive Committee of the British Machine Vision Association and a member of the EPSRC Peer Review College.

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

The Southampton–York Natural Scenes Dataset

The Southampton-York Natural Scenes (SYNS) dataset is a newly released public database with the following link:

<https://syms.soton.ac.uk/>

It includes more than 90 scenes, surveyed at randomly selected locations within 25 diverse indoor and outdoor scene categories. Each survey includes:

1. spherical LiDAR range data;
2. high-dynamic range spherical imagery; and
3. a panorama of stereo image pairs.

The database is described in the following paper:

W.J. Adams, J.H. Elder, E.W. Graf, J. Leyland, A.J. Lugtigheid and A. Murry (2016). The Southampton–York Natural Scenes (SYNS) dataset: statistics of surface attitude. *Scientific Reports*, 6, no. 35805.

<http://www.nature.com/articles/srep35805>

Professor Wendy Adams
University of Southampton
email: w.adams@soton.ac.uk

²Thanks are due to UCL Department of Computer Science for providing this summary, which is taken from Professor John Shawe-Taylor's introduction to Lourdes' inaugural presentation.

³For those who would like to hear more, here is a link to an article by Prof. Paul Curzon (a former colleague of Lourdes) about Lourdes' fascinating presentation: <http://www.cs4fn.org/vision/3Dmodelsinmotion.php>

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.org/executive_committee

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 five candidates. The results will be announced at BMVC and in *BMVA News*.

Adrian Clark
BMVA Chair
email: 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