# 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 December 2018.

## Contents

## Editorial: *A Rip van Winkle View of Computer Vision*

Having reluctantly had to miss all conferences and meetings for the past two years, BMVC 2018 took me by surprise. I had not realised how substantial ongoing progress in the

---

[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.

subject had been in the interval: while 2 years is extremely short by Rip van Winkle standards (he is reputed to have gone to sleep for 20 years), in the current state of the subject it seemed like an age. Of course, I had been all too aware of the deep learning explosion of 2012 to 2015, but had assumed that its main effects had passed, though equally, I had realised that many people were capitalising on the new techniques, methods and databases and were using them in their own individual application areas. But the idea that the explosive development might be continuing virtually unabated had eluded me. In particular, whereas the original explosion had largely concentrated on processing individual images from enormous databases, by 2018 it seemed that this had progressed to processing *whole videos* in a similar manner. So (to take one example) we learned about 'anchor points', which are needed if extracts of videos are to be compared, and for me this was a rather new concept.

There were other such concepts and other emphases to learn about. Amongst the BMVC titles were more papers on person re-identification than might have been expected; a substantially enhanced interest in semantic segmentation; significant more attention to novelty and anomaly detection; much interest in zero-shot object detection (not to mention single-shot and few-shot detection); strong developments in adversarial learning and adversarial networks, including generative adversarial networks; enhanced interest in pedestrian detection; a new interest in facial expression recognition; and important new work on self-supervised learning, deep reinforcement learning and 'learning on the edge'.

Perhaps more important, some well-known and (one would have thought) indispensable techniques and subject areas seem to have come in for hugely depleted interest and relevance. One could argue that this is temporary and arises because of the immediate focus on new application areas that people have had to tackle, so they simply were not forced to look at methods that they 'ought' to be using. In fact, the attention being paid to such methods as edge detection, corner detection, line detection, Hough transforms, RANSAC, perspective, homography, invariants, and 3D vision as a whole – including the essential and fundamental matrices – all these seem to have dropped to a low level of attention, implementation and use. Clearly, in their rush to get new applications working, people are 'voting with their

feet' and are not bothering with uncalled for tasks. (I should add that 'voting with one's feet' is often an intuitive action rather than a fully thought through decision about the best way forward).

Actually, my Rip van Winkle view of the world may reflect something much deeper. I have long mused about the representation that the brain must be using in its analysis of visual scenes. Take the case of pedestrian surveillance: the first view the brain gets is the frontal one, but as time progresses and the pedestrian moves forward, it is quite easy to adopt a plan view of the motion and map the detailed path travelled by the pedestrian. But does the brain really do this? And does it need to do it? Surely this amounts to needless processing, as all the information needed about the pedestrian can be obtained from an ongoing frontal view. Indeed, the plan view of the pedestrian's path will be riddled with blank patches where occlusions have taken place and interpretive hypotheses are needed to make sense of them. So the answer to my musings is that detailed 3D analysis tends to be unnecessary and that all the required information can be obtained from the raw 2D data. What is needed to achieve this is a set of suitably trained deep neural networks that cut straight to 2D interpretation. In fact, we already knew this from the semantic segmentation work done from 2012 onwards – and still going strong.

Thus Rip van Winkle can lazily say that 3D vision per se is already dead. Likewise, edge, corner, line, ellipse and many rather basic detectors from yesteryear have already bitten the dust, because – when and if needed – they have been subsumed into end-to-end trained networks.

Many would say that setting all these standard techniques aside is an offence to science: how can all these hard-gotten gains merely be cast into the rubbish bin in favour of deep neural networks (DNNs) whose inner workings are incompletely known and therefore suspect. Yet, people are successfully solving new problems and confidence in DNNs is ever-growing. I have to add that Rama Chellappa's invited talk at BMVC reflected this. His approach was rigorous, confident and progressive: he advocated going ahead with research in this area and (of course) ensuring a proper abundance of training examples; and above all finding out quite where DNNs are lacking and instigating any necessary 'mitigation strategies'.

Meanwhile, a number of old learning strategies that once had their place in psychology or the wider field of AI are now becoming key approaches for computer vision: these include reinforcement learning (and deep reinforcement learning); self-supervised learning; and adversarial techniques. Note that generative adversarial nets (GANs) only arrived on the scene in 2014 and received strong attention in Rama Chellappa's keynote address.

As a parting shot, I continue to worry about dropping all our stalwart ideas of yesteryear. Should we pretend we were asleep from the 1970s to the 2010s, or else concur that all our ideas, ego-trips and torture must have been vital, as they enabled the evolution of the subject to develop in its own inimitable way? Either way, we should be comforted by the 'fallout' argument – that at least we now understand the subject a whole lot better.

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

## Election to the BMVA Executive Committee

Over the summer, elections were held to the BMVA Executive Committee. Up to six places were available but only five nominations were received. The following people are thus duly elected for the period 2018–2020:

- Marina Bloj
- Toby Breckon
- Adrian Clark
- Jingjing Deng
- Andrew Gilbert

Several PhD students have also contacted me in response to my plea for more involvement from them in the BMVA. The newly-elected ExCo will sort out how best to involve them in time for its next meeting, which is likely to be in early November.

Adrian Clark
BMVA Chairman

## BMVA Distinguished Fellow 2018: Majid Mirmehdi



Every year, the BMVA awards a single Distinguished Fellowship to a person in recognition of their outstanding contribution to Computer Vision in the UK. These contributions might be in recognition of significant advances in an important technical area, or in a addressing a particularly difficult or long-standing problem, or in guiding and supporting the research community. I am pleased to say that this year's Distinguished Fellow, Professor Majid Mirmehdi, has made contributions in all these areas.

Like me, Majid first started publishing in the area of parallel computer vision, figuring out ways of making vision algorithms run quickly on transputers. He then worked on texture and its use in probabilistic frameworks before moving on to explore the use of computer vision as an inherent part of larger systems, often in the medical domain – with papers on lung function, healthcare and assisted living published in the last couple of years. Back in the early

1990s, many of these applications would have seemed impossibly difficult, more fiction then aspiration, and the work of Majid has helped bring them to reality.

Despite these many and varied research contributions, our Distinguished Fellow is perhaps best known for his work in supporting the research community. He plays an important role in Bristol's Visual Information and Robotics Laboratories. He is editor-in-chief of the IET's Computer Vision journal, and an associate editor of Pattern Recognition and Pattern Analysis & Applications. He is an advisory editor of Scientia Iranica. Moreover, he has done these things while holding the significant roles of Graduate Dean and Faculty Graduate Education Director in his institution, so his commitment is clear for us all to see.

Majid has served on the BMVA's Executive Committee for approaching 20 years, chairing the BMVA from 2005 to 2008. He has been heavily involved in staging BMVC twice, most recently in 2013. On a personal note, I am grateful for the considered and thoughtful advice he has given me during my own tenure as BMVA Chairman. Finally, our Distinguished Fellow is also a Fellow of the IAPR, the parent body to which almost all national vision societies belong.

At a time when, given all he has done, our Distinguished Fellow could metaphorically put up his feet and rest on his laurels, he agreed to represent the BMVA's interests on the IAPR's Governing Board. This is typical of his commitment to the vision community.

Ladies and gentlemen, please join me in acknowledging the contributions of the BMVA's Distinguished Fellow for 2018, Professor Majid Mirmehdi.

Adrian Clark
BMVA Chairman

## BMVC 2018's Place in History!

BMVC 2018 has been a great success! This year, BMVC is organized at Newcastle upon Tyne. The conference is hosted by Northumbria University, a research-rich, business-focused, professional university with a global reputation for academic quality. The main conference venue is the Student Union building located at the core of Northumbria University in the city centre. The Conference Reception, 'Simple Dinner' and Conference Banquet are hosted at the Newcastle Civic Centre, which was first opened in 1968 with a reputation of rich history and authentic modernist architecture.

BMVC 2018 has made several major achievements. It attracted a total of 862 full paper submissions, which is the highest number in the history of BMVC – a sharp rise over 635 submissions in BMVC 2017 and 365 in BMVC 2016. Of the 862 submissions, 255 were accepted for presentation in BMVC 2018, which is another record high. The acceptance rate is 29.5% – one of the lowest in BMVC history. The accepted papers represent a truly international research community, with 13% of the papers being from the UK, 21% from Europe excluding the UK, 23% from North America, 37% from Asia, 3% from Australia and the rest from elsewhere in the world. This year there were 3 full

keynotes given by Rama Chellappa, Sven Dickinson and Shaogang (Sean) Gong. There were also 4 tutorials given by Vittorio Ferrari, Ivan Laptev, Abhinav Gupta and Zeynep Akata. BMVC has always had strong links with industry, and again we are very grateful to our industrial sponsors for supporting the event. Platinum Sponsors: Scape, Amazon, Microsoft, NVIDIA, SCAN, Sage, Apple, Facebook, Intel. Gold Sponsors: Disney Research, IET, Telensa, iniVation. Silver Sponsors: Gaist, Ocado, Snap Inc. Special Support: Springer. BMVC 2018 had a record high of over 510 attendants.

We wish to thank all members of the Organizing Committee, the Area Chairs, reviewers, emergency reviewers, authors, and the CMT and TPMS teams for the immense amount of hard work and professionalism that has gone into making BMVC 2018 a first-rate conference. See you again at BMVC next year!

Professor Ling Shao
Dr Hubert P.H. Shum
Dr Timothy Hospedales

## BMVC – Prizes and Awards

The following prizes and awards were presented at the Conference Banquet:

1. Best Science Paper: *Non-smooth M-estimator for Maximum Consensus Estimation*
   Huu Le, Anders Eriksson, Michael Milford, Thanh-Toan Do, Tat-Jun Chin, David Suter
2. Best Science Paper Honourable Mention: *Deep Network for Simultaneous Stereo Matching and Dehazing*
   Taeyong Song, Youngjung Kim, Changjae Oh, Kwanghoon Sohn
3. Best Industry Paper: *Automatic Semantic Content Removal by Learning to Neglect*
   Siyang Qin, Jiahui Wei, Roberto Manduchi
4. Best Industry Paper Honourable Mention: *Semantic Priors for Intrinsic Image Decomposition*
   Saurabh Saini, P.J. Narayanan
5. Best Student Paper: *QuaterNet: A Quaternion-based Recurrent Model for Human Motion*
   Dario Pavllo, David Grangier, Michael Auli
6. Best Student Paper Honourable Mention: *Learning on the Edge: Explicit Boundary Handling in CNNs*
   Carlo Innamorati, Tobias Ritschel, Tim Weyrich, Niloy Mitra
7. Best Organization Chairs: Edmond S.L. Ho, Kamlesh Mistry, Ammar Belatreche
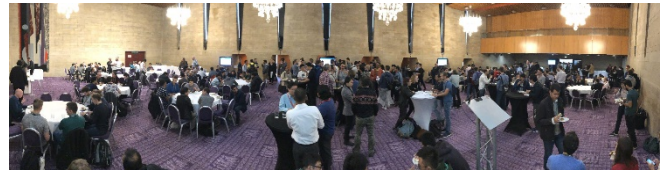8. Student Bursary: Zheng Xu, Tae Joon Jun, Fan Wang, Terence Morley, Naima Otberdout.

# Around and About at BMVC 2018



In his keynote, Shaogang Gong (QMUL) demonstrates the complexities of people search in surveillance videos.



Tellingly, Sven Dickinson (University of Toronto) starts his keynote with a rather remarkable optical illusion.



Close-up of keynote speaker Rama Chellappa giving his gripping talk on the how and why of deep learning networks.



Abhinav Gupta (Carnegie Mellon University) presenting on inductive visual localisation.



Panorama of the Civic Centre Reception room.



Jie Li, Zheming Zuo and Daniel Organisciak: three colleagues meet again.
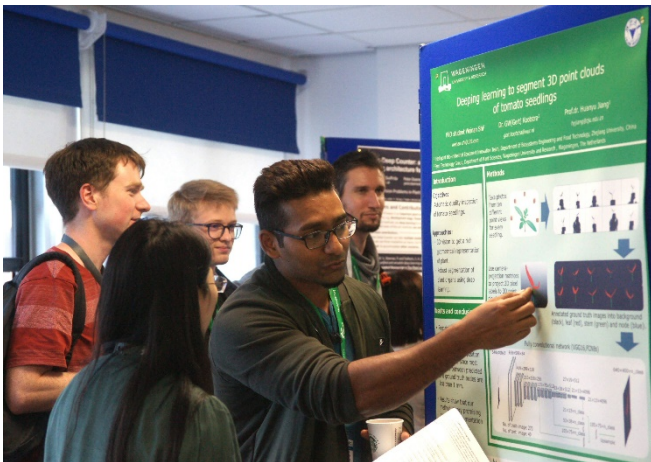


Yael Moses, Charles Attwood, Tim Cootes, Majid Mirmehdi and the photographer enjoy a less than chance encounter over coffee.
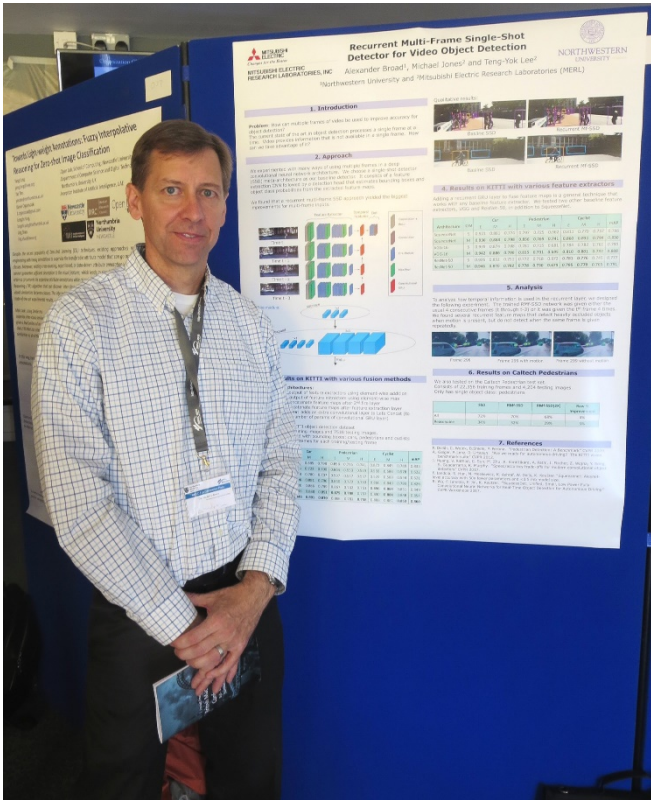
Michael Jones presented a fascinating poster on video object detection.

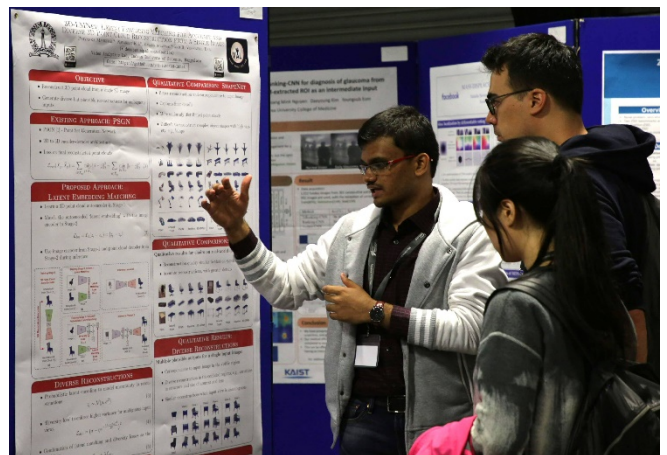

A deep question on segmenting tomato seedlings for Weinan Shi (facing away from the camera).



Chen Gao explaining his Instance-Centric Attention Network.



Aditya Arun showing Sven Dickinson his work on learning human poses.



Siddharth Mahendran explaining his work on pose estimation.



Priyanke Mandikal describes his 3D-LMNet for latent embedding matching.
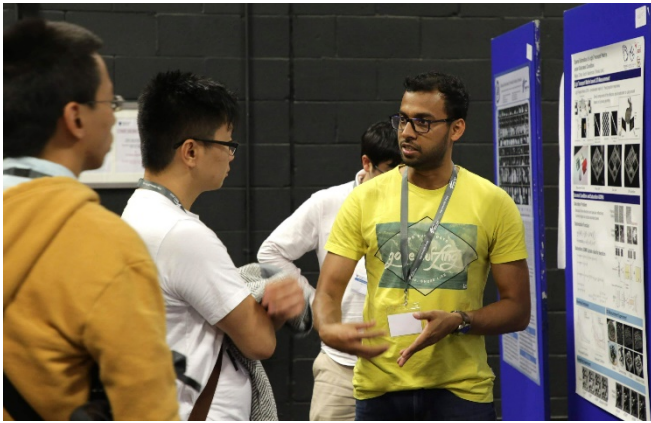
Pulak Purkait presenting his poster on synthetic view generation.



The RCO String Quartet signals that the banquet is almost ready to begin.
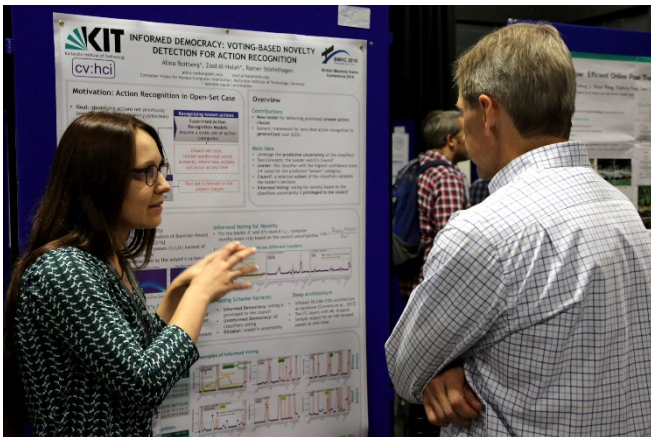


Alina Roitberg explains about informed democracy for action recognition.



Majid Mirmehdi, Adrian Clark, Christine Clark, invited speakers Sven Dickinson and Rama Chellappa, and last but not least, Ling Shao – all ready for the food to arrive!



Dimitrios Sakkos and others listen intently as Suriya Singh describes his work on self-supervised semantic segmentation.



Majid Mirmehdi, Kirill Sidorov and Adrian Clark

Darrel Mond (iniVation) and Charles Attwood (Thales) seemed happy with the feast!



Enjoyable conversations at every table!



Kirill Sidorov (Cardiff) and Peter Hall (Bath)



Huu Le receives his Best Science Paper Award from Ling Shao.



Taeyong Song with his Best Science Paper Honourable Mention Award.

Siyang Qin receiving his Best Industry Paper Award.



Dario Pavllo receiving his Best Student Paper Award.



Saurabh Saini receiving his Best Industry Paper Honourable mention award.



Edmond S.L. Ho, Kamlesh Mistry, Ammar Belatreche with their Best Organisational Chairs Awards.

Zheng Xu, Tae Joon Jun, Fan Wang, Terence Morley, Naima Otberdout with their Student Bursary certificates.



Majid Mirmehdi receiving his BMVA Distinguished Fellow 2018 Award from Adrian Clark.

May I take this opportunity to thank the following team of helpers for providing a wealth of photographs making it easy for me to compile this gallery of events from BMVC.

- Kevin Mccay
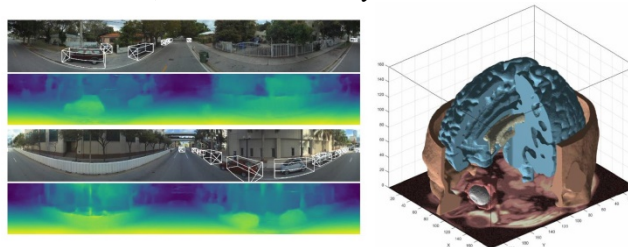- Daniel Organisciak
- Dimitrios Sakkos
- Edmond Ho.

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

## Upcoming Meetings

### Deep Learning in 3 Dimensions

This one-day BMVA Symposium will take place in London on Wednesday 20 February 2019.
www.bmva.weebly.com

Chairs: Chris Holder, Chris Willcocks, Grégoire Payen de La Garanderie, Durham University



### Call for papers

Deep learning has revolutionised the world of artificial intelligence in recent years, providing a huge boost to machine learning research as well as to real-world applications of such technology. In this meeting, we aim to explore the key challenges of combining deep learning with 3D vision. Areas of interest include, but are not limited to:

- Neural network architectures for processing 3D information
- Representations and encodings of 3D data for deep learning
- 3D classification and segmentation
- Models for generating 3D data
- 3D visualisation
- 3D reconstruction
- SLAM
- Volumetric biomedical imaging.

### Submission Deadline

We encourage submissions from students, academics and practitioners in the area. Anyone interested in presenting on the day should submit an abstract to https://goo.gl/r8mWE2 by 21 November 2018.

### Registration

Register online at www.bmva.weebly.com
£16 for BMVA members, £36 for non-members (both prices include lunch).

We have 4 exciting varied meetings planned in 2019 so far: check out the website www.bmva.weebly.com for all the details.

## Summary of Upcoming Meetings
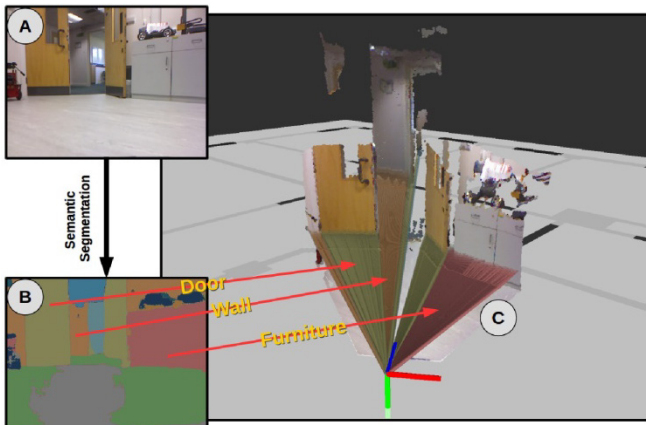
Andrew Gilbert
University of Surrey
email: a.gilbert@surrey.ac.uk

## Report on Robotics meets Semantics: Enabling Human-Level Understanding in Robots

Chair: Oscar Mendez, University of Surrey
Organiser: Andrew Gilbert, University of Surrey



This technical meeting was held at the BCS headquarters in London on 18 July 2018. The main focus of the workshop was the applicability of semantically-enabled vision for robotics and brought together experts in machine learning and robotics.

Daniel Cremers (Technische Universität) opened with the first keynote presentation. He introduced direct methods in the context of multi-view reconstruction. The advantages of cameras over laser scanners were shown by reconstructing thin geometry such as a skipping rope but as an offline process. He next showed a real-time dense reconstruction algorithm which utilises photometric consistency. This led to direct SLAM methods, showcasing LSD-SLAM implemented on a drone. He finished by showing that by using deep learning to estimate depth in 2D images, a 'virtual stereo' pair can be created which allows accurate localisation for monocular sensors, sometimes matching the performance of stereo sensors.

Following was a talk given by Ondra Miksik (University of Oxford). He pointed out that sparse semantic scene descriptions are insufficiently detailed for autonomous cars, and a full semantic overlay is more ideal. Streamed data leads to flickering segmented areas, while batch processes are not real-time. He explained that the current state-of-the-art is limited, e.g., to indoor environments, to offline data, etc. He introduced Semantic Stereo Fusion, which extracts 2D features to be transferred into a densely-connected conditional random field, through evaluation from random forest classifications. Ondra finished by showing his progress with a SLAM-augmented deep reinforcement learning approach to the DOOM video game. Despite achieving an impressive performance, he pointed out that even the state-of-the-art is not very close to a human level.

After a short coffee break, Cyrill Stachniss (University of Bonn) gave the second keynote talk which covered two projects. The first was Rovina, to digitise Roman Catacombs for cultural heritage. He showcased the robotic platform which used Kinect-like cameras and an arc of cameras with bright LEDs, powered by two laptops. The caterpillar tracks allow the robot to ascend/descend staircases. The localisation and mapping using optimised ICP for a traversability analysis, and can exploit hand-drawn maps. Cyril pointed out the problems in reconstruction when the light source moves with the cameras. Nevertheless, the final reconstructions were of high quality, and can even be navigated in virtual reality. The second project was the Flourish project, a solution to weed management for crops with the motivation to minimise the impact on the ecosystem and to increase the yield. It is a more open task with many avenues to explore. Plants in a field can be semantically classified using a drone or wheeled rover and weeds can be identified. One environmentally friendly way to eliminate them is to use a nailgun-like stamping mechanism mounted on the rover. More deep learning is being investigated for weed distribution based on previously stamped areas, fertilisation levels, disease, etc. Aerial photographs from different times in the year can be difficult to align, whereas plant-level monitoring solves this with time-aligned point clouds of the plant locations from the rover.

Stuart Golodetz (University of Oxford) gave a talk explaining dense 3D reconstruction using a collaboration of multiple agents. His motivation was that existing work can be low-cost and interactive but with limited scale, or large-scale but high-cost and offline. His approach aims to take the best of both worlds. Each agent uses visual-inertial odometry which is locally accurate. Rather than expensively linking images between visits, the trajectories and 3D maps are used, using multiple estimates as samples. Synthetic RGB-D raycasts from one agent are rendered in the frame of a previous agent. In practice, a lab around 820m$^2$ is reconstructed in under 20 minutes, using android devices with WiFi for the data capture.

After lunch, an industrial talk was given by Brian Holt (Parkopedia). He outlined the autonomous valet parking work he is involved with. It is the most desirable advanced driver-assistance system, as people are generally comfortable handing over control for this. The premise is for people to leave their vehicle at a multi-storey car park, and for it to park itself. On return, the vehicle will leave its parking space and meet the user. Brian outlined four possible options: (1) All computation is performed on the

device. This will make vehicles expensive and complex, and auto-makers don't want to do this. (2) Sensors are fitted to existing car parks. This introduces the problem of determining who pays for the fitting and maintenance. (3) Add markers such as QR codes to the car parks. These must be clearly visible to be useful, and there is a risk that they will be in poor light or sprayed with graffiti. (4) Shared high-definition maps of the car parks. This is the chosen route, and the maps will be obtained using Lidar. There is still discussion about which format to use, and how to ensure that they work everywhere.

Rebecca Allday (University of Surrey) showed how she uses reinforcement learning for human-inspired robotic grasping in a supermarket context. Engineered grasps require known and accurate representations of the objects to be held, so this is not a scalable solution. Instead, she presents a data-driven approach using a Markov decision process, enabling previously unknown objects to be picked up. She highlighted the slow rate of data capture for robotics, which is tackled using simulation (ROS multi-master and Gazebo) and transfer learning. Another simulated method was utilising curriculum learning for a 2D robotic arm to reach a specified position.

Next, Joanna Olszewska (University of West Scotland) detailed the visual ontology classification project. This is work towards a standardisation of robotics, perhaps so that a unified communication system can be used between heterogeneous systems. She covered Spatio-Temporal Visual Ontology (STVO) for inter- and intra-object spatial reasoning and the 3D clock model for directional relations. Among the applications for this are the automated reasoning for static and dynamic scenes, for faces, and expressions. These are well-suited for human-robot interactions in unconstrained real-world scenarios.

Michael Wray (University of Bristol) explained how action descriptions may be ambiguous, and his methods to create unequivocal representations. The commonly used verb-noun labels are insufficient to fully describe some actions. An example is that the words 'push', 'pull', 'turn' and 'cut' can all be used to describe the action 'open'. Is it possible to ignore the noun and use multiple verbs instead? Using video-to-text retrieval, a network treated each verb score as a dimension in the embedding space. This allows actions between datasets to be matched based on their similarity despite having completely different ground truth labels. Another benefit of the multi-verb representation is that previously unknown actions can be described as a combination of others. By restricting the vocabulary, one can choose how coarse or fine the action description should be.

After a short break, the third keynote talk was led by Andrew Davison (Imperial College London). He introduced the term Intelligence Augmentation (IA) as the next evolution after SLAM, where vision complements artificial intelligence to allow devices to usefully interact with its surroundings. He gave an overview of how systems transitioned from sparse to dense map representations, and then with embedded semantic information. The newer techniques are making their way into products today, but they are still expensive and limited. Can SLAM be replaced with an end-to-end trained network? He explored the possibilities for network modularity, data capture and level of supervision for training. Next were the different ways one might visualise the various graphs in neural networks for

spatial AI, and where in the pipeline the different types of processing might take place. Andrew proposed that the degree of fidelity in the map representation should be controllable. CodeSLAM represents depth in a compact way with autoencoders. Reliable techniques exist to get depth from monocular images but the gross depth is incorrect. The depth map of CodeSLAM can be adjusted by changing a single axis of the 'code'.

Andrew Calway (University of Bristol) explained a localisation method which uses images to produce binary semantic descriptors, which are used for large-scale urban localisation. Where image-to-image database matching is difficult to scale up and not robust to image conditions, an alternative is to use semantic features in a similar way to human map-reading. A simple 4-bit descriptor is generated for given positions which looks for junctions and gaps between buildings, and this is applied to the images from Google Streetview and the 2D Open StreetMap. While a single descriptor has ambiguity in its matched locations, a set of consecutive locations quickly discriminates.

Finally, Horia Porav (University of Oxford) exhibited an approach to increase the robustness of algorithms such as visual odometry by warping images from one weather condition into another. The advantage to this is that it removes the need to find a universal descriptor which performs well in all conditions. A cycle-GAN architecture is used which does not require image sequences. This does not have to be pixel perfect and some visual inconsistencies are permitted so long as the feature descriptors can be regenerated. The best results are achieved by transforming poor conditions such as rain or night into a good day. Future work involves discretising conditions, for example, having snow at night.

Celyn Walters
University of Surrey
email: c.walters@surrey.ac.uk

# Report on MIUA 2018

https://miua2018.soton.ac.uk/

We very much enjoyed having the 22nd MIUA Conference at Southampton on July 9–11 this year. This is a conference under the aegis of BMVA and is much smaller than BMVC. Medical Image Understanding is not a major area of mine, and I must say I am very impressed by this conference series. This fine little conference is very high level and has a set of dedicated people who have helped it to run over the years. The program committee did an excellent job and were very much on time. We had 38 papers made up of 27 oral papers, 7 posters and 4 clinical abstracts. The papers came from 15 different counties with many from top institutions. The sessions at the conference were all well attended: having many people in attendance on the last morning is always the sign of a good conference. Another indication of its quality is the level of sponsorship it attracts (five in total), and that almost twice as many people registered for the conference as there were papers. There was much spirited discussion too

after each paper, which is often the case when Mike Brady is in attendance!


An attentive audience

We held the conference at the University of Southampton's Highfield campus, inside the Nuffield theatre complex. It took place in this summer's hot weather. In the UK only the biggest lecture theatres are air-conditioned so we had fun managing the heat, and with a combination of fans, water coolers and door chocks the temperature was kept to a good level. We dined in the University staff club accommodation and the food went down well. Especially the conference dinner. At my table every plate was empty for each course and that was not because the amounts were paltry: clearly, we all enjoyed the food!

We had oral and poster presentations and the posters had a two-minute oral on the podium. The programme was complemented by three invited keynote speakers: Professor Alison Noble from the University of Oxford; Professor Anant Madabhushi from the Case Western Reserve University, USA; and Professor Tom Wilkinson from the University Hospital Southampton NHS Foundation Trust. We were pleased to have all three and enjoyed their mix of research and practical application. It was great to see Alison Noble back as she originated the MIUA series. New to the MIUA series, we had two special sessions, one on Liver Image Analysis (organised by Mike Brady and Benjamin Irving) and the other on Ocular Imaging Analysis (Yalin Zheng). These were organised by Emma Lewis from the University Hospital Southampton. As in the previous MIUA there were Clinical Abstracts, and these were also organised by Emma.


At coffee with posters, clinical abstracts and sponsors

The following prizes were awarded to papers (with the poster and clinical abstract prizes adjudged by delegates):

1. Best Paper: M. Brudfors, Y. Balbastre, P. Nachev, J. Ashburner for *MRI Super-Resolution using Multi-Channel Total Variation*. The Wellcome Centre for Human Neuroimaging, UCL, London
2. Best Poster: V. Azzopardi, M. Guy, E. Lewis for *Identifying shape-based biomarkers for diagnosis of Parkinson's Disease from loflupane (123I) SPECT data*, University of Surrey, Guildford and University Hospital Southampton
3. Best Clinical Abstract: L. Vass, M. Fisk, M. Polkey, I. Wilkinson for *Compartmental modelling of 18F-FDG to assess Pulmonary Inflammation in Obstructive Sleep Apneoa/Hypopneoa Syndrome (OSAHs) Patients*, University of Cambridge, Barts Health NHS Foundation Trust, Royal Brompton and Harefield NHS Foundation Trust, Cambridge University Hospitals.


At the conference dinner

In parallel with his tireless work for the conference my co-chair, Sasan Mahmoodi, arranged for excellent sponsorship support. Our Gold sponsors were Mathworks and the MedIAN network. Our Silver sponsors were OCF, the Journal of Imaging and NVIDIA. We had presentations from Median and Mathworks in the conference programme, and Bulat Khusainov from Mathworks gave an excellent and well attended tutorial on Matlab and its use (especially on Deep Learning). I gave a tutorial on Journal paper writing, which is something we are doing in the conferences of the IEEE Biometrics Council, of which I am President.


Attendees on the last day

So there was much new to MIUA when it was held at Southampton. There was a quiz introducing it, as happened last year at MIUA 2107. There were special sessions and tutorial sessions. The delegates enjoyed a bottle of beer brewed locally for the conference. They also enjoyed the magician Greg Williamson who entertained us at the conference dinner. These are largely ideas derived in the

conference series run by my IEEE Biometrics Council. MIUA might be more occasional territory to me, but I enjoyed the conference, its ethos and spirit. Education should be at the highest level and fun: MIUA 2018 was certainly that.

I remain grateful to my two co-chairs, Sasan Mahmoodi (University of Southampton) and Reyer Zwiggelaar (Aberystwyth University), and to Emma Lewis, as mentioned previously – and to the MIUA steering committee. Equally, Vicky Bennett and our students kept the conference running well. Thanks for coming to Southampton, we enjoyed it too. Next year, people will enjoy the conference when it is held at the University of Liverpool – see the website https://miua2019.com/ .

Mark Nixon, BMVA Distinguished Fellow 2015
Co-chair MIUA 2018

# Report on the 22nd MIUA

This year MIUA 2018 was organised by the University of Southampton in collaboration with Aberystwyth University and supported by the BMVA: it took place between 9 and 11 July at the University of Southampton. The location of the conference was The Nuffield Southampton Theatre (NST) in the university campus. MIUA 2018 specifically targets research in theory and applications providing a forum for the members of this community. The conference proceedings will be within Springer's Communications in Computer and Information Science series, with later association with the Journal of Imaging.
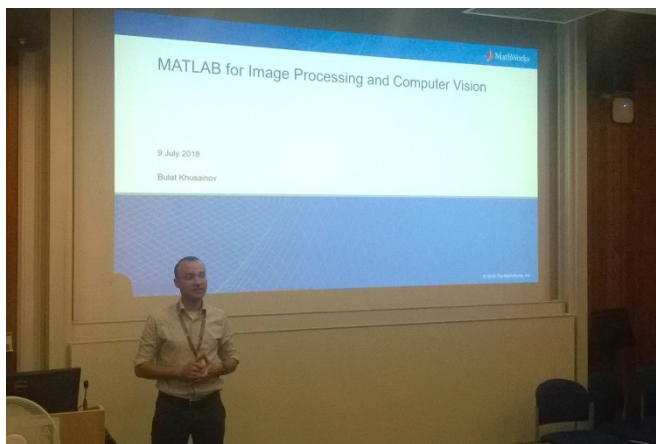

Matlab Presentation for participants.

The topics covered in the oral and poster sessions of the main conference were highly representative of the key interests of today's Medical Imaging community. The main conference offered 3 Keynote Lectures explaining the importance and greater advances of the field, and oral and poster paper deliveries organised in thematic sessions. The programme consisted of 38 papers, which were made up of 27 oral papers, 7 posters, and 4 clinical abstracts and included sessions on Deep Learning in Medical Imaging, Texture and Image Analysis; MRI: Applications and Techniques, Segmentation in Medical Images; CT: Learning and Planning; and Application of Medical Image Analysis.

New to the MIUA series, this year the programme was enriched by two Special Sessions: one on Liver Analysis and the other on Ocular Imaging Analysis.

In addition to the main conference, two one-hour tutorials, MATLAB and Journal Paper Writing Club were organised to further explore the key interests of the community in a more informal setting.


Journal Paper Writing Club by Mark Nixon.

Invited speakers presented 3 Keynote Lectures on 9, 10 and 11 July in the following order.

Alison Noble, the first chair of MIUA 1997, from the University of Oxford presented "How Machine Learning Is Changing Ultrasound Image Analysis" with a particular focus on raising the profile of ultrasound imaging as a first class data type by understanding the interplay of ultrasound device design, clinical acquisition, and downstream image analysis and computer vision. She also talked about applying machine learning to ultrasound to advance automatic 2D, 3D and video ultrasound analysis and how to extract the useful clinical content in ultrasound images and videos.


Mark Nixon introduced Alison Noble.
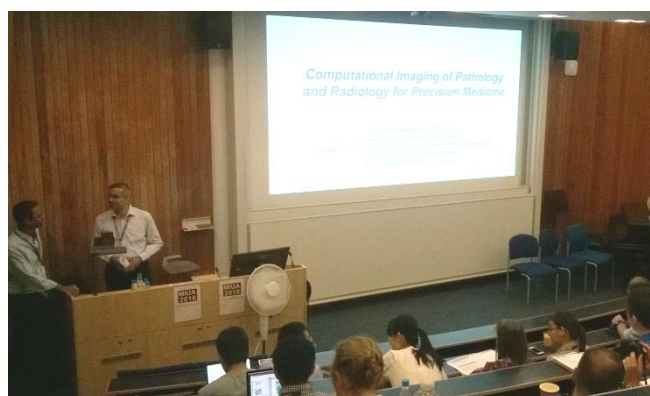

Nasir Rajpoot introduced Anant Madabhushi.

Anant Madabhushi – a well-established figure in the Computational Imaging community from Case Western Reserve University (USA) and director of the Case Western Reserve University Centre for Computational Imaging – presented "Computational Imaging and Artificial Intelligence: Implications for Precision Medicine" with an emphasis on deep learning versus hand-crafted features. He started his talk by explaining the importance of learning and addressing previous works in computer analysis of cell images and took participants on a very nice journey all the way to up-to-date research in computational imaging.

Tom Wilkinson from Faculty of Medicine and Honorary Consultant at University Hospitals Southampton talked about "Image Analysis in Respiratory Disease" focusing on understanding the mechanisms which contribute to the vulnerability to and impact of respiratory infections in patients with chronic lung disease. He also explained how to establish a cohort of deeply phenotyped patients with COPD – a disease of global importance which will become the third greatest cause of worldwide mortality by 2020 – to determine the mechanisms underpinning the dynamics and susceptibility to infection and to develop new vaccines for this patient population.
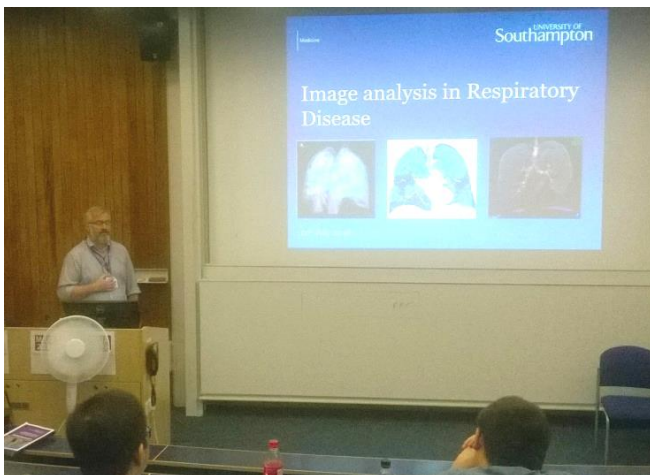


Tom Wilkinson presented "Image Analysis in Respiratory Disease" to the participants below:



A collection of 7 posters were presented in conjunction with Poster highlights and breaks. Each poster presenter talked about their work for 2 minutes. Posters spanned a variety of data analysis methodologies and clinical applications in cancer. Numerous interesting projects were discussed in oral presentations. The most memorable research problem shown was about smiling and the talk title was "What is in a smile? Initial Results of Multilevel Principal Components Analysis of Facial Shape and Image Texture": in this talk Damien J.J. Farnell concluded that strong changes in mouth, exposure of teeth, and increased prominence of the cheeks – in terms of shape and image texture – occur when we smile.



Conference Dinner and MIUA 2019 invitation.

After the successful conference over 2 days, it was time to network with others and enjoy the well-earned conference dinner and the social event. This year's MIUA 2018 welcome reception and conference dinner were held at Hartley Suite – university Staff Club. Before the conference dinner participants were offered drinks and a magician played his tricks. We were very excited and tried to understand and solve his magic but nobody succeeded. After the conference dinner the awards were presented. The paper with the title "MRI Super-Resolution using Multi-Channel Total Variation" by M. Brudfors et al. was given the 'Best Paper' award and the Best Poster award went to "Identifying shape-based biomarkers for diagnosis of Parkinson's Disease from Ioflupane (123I) SPECT data" by V. Azzopardi et al. Finally the bar opened and started serving alcoholic beverages. The participants indulged in beer, networked in an informal medium, and the sounds of toasting constantly rang through the air. The conference chairs also offered participants a free bottle of beer which was brewed locally and labelled the 22$^{nd}$ MIUA.

All in all, MIUA 2018 was a great success, and a memorable experience academically. After months of hard work with the satisfaction of a job well done, the organisers could pass the MIUA 2018 baton to the organising committee of MIUA 2019. Yalin Zheng talked about next year's conference – MIUA 2019 and invited everyone to Liverpool. The announcement of MIUA 2019 was made at the closing ceremony. A lot has been promised to prospective participants: a beautiful setting, excellent cuisine at a stylish restaurant, an easy commute for UK participants, and of course the best and the latest research of the Medical Image Understanding and Analysis world. I would strongly encourage researchers working in the field to consider submitting to MIUA 2019.



The poster session covered a wide variety of topics in cancer image analysis and modelling and provided excellent opportunities for scientific exchange and detailed discussions.

Above and below: The best paper and the best poster awards were presented by Mark Nixon to M. Brudfors and V. Azzopardi.



I would like to thank BMVA for sponsoring my conference trip to Southampton and helping me to get very good experience for my PhD.

Cefa Karabağ
City, University of London
email: cefa.karabag.1@city.ac.uk

## Report on International Conference on Image Analysis and Recognition (ICIAR) 2018

In the last few years, the employment of deep learning methodologies in pattern recognition and computer vision tasks has grown exponentially. Therefore, it was no surprise that the 15[th] International Conference on Image Analysis and Recognition, hosted in the coastal Portuguese city of Póvoa de Varzim, boasted a diverse range of deep learning methods by researchers from around the world. The conference, which ran from 27 to 29 June, aimed to highlight novel work within natural image enhancement, restoration and reconstruction, biomedical image segmentation, and human activity recognition.

Amir Atapour-Abarghouei, from Durham University, addressed the problem of hole filling in depth images, obtained from active or stereo sensing, for the purposes of depth image completion in an exemplar-based framework.

Using colour (RGB) and depth (D) information available from a commonplace RGB-D image, patch prioritization terms are utilized for target patch ordering to enable improved distribution of complex texture and linear structures within depth completion. Evaluations demonstrate the efficiency of the proposed method compared to other state-of-the-art completion techniques.
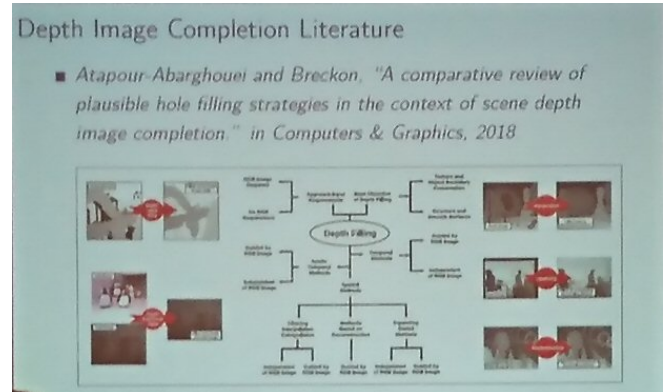


Amir Atapour-Abarghouei: Extended Patch Prioritization for Depth Filling Within Constrained Exemplar-Based RGB-D Image Completion.

Tien D. Bui, from Concordia University (Canada) discussed the latest developments and effectiveness of sparse representation and low-rank approximation techniques in large-scale data processing problems. It was highlighted that such techniques can also be used in unsupervised and dictionary learning to uncover high-order relations in the data and to train deep neural networks.
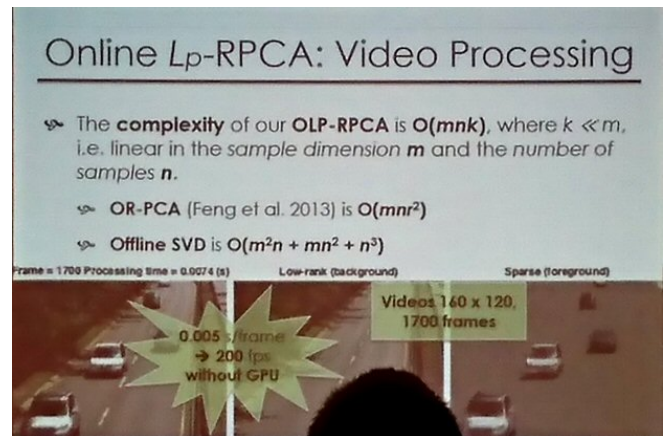


Tien D. Bui: *Lp* Norm Relaxation Approach for Large Scale Data Analysis: A Review.

The first and final day of the conference show-cased a diverse range of poster presentations that ran throughout the main afternoon sessions. Different areas of research included facial expression recognition using artificial neural networks, sign language recognition based on 3D convolutional neural networks, nerve structure segmentation from ultrasound images using an SVM classifier, and prediction of damage in SAM images of semiconductor devices (to name just a few).

A particular highlight included the Grand Challenge on Breast Cancer Histology Images. Breast cancer diagnosis and prediction are critical for addressing mortality rates among women, and Bahram Marami, from The Mount Sinai Hospital (USA), received first prize for his work towards

region identification in breast histopathology slides. This innovative tool combines recent developments in deep learning for image classification, trained using different data subset sampling. Consequently, the proposed approach produces a high-performing detector with robustness to data variations.
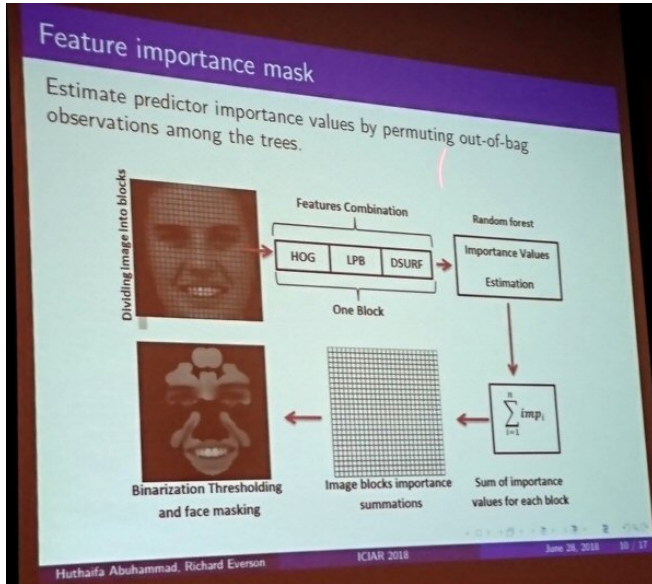


Huthaifa Abuhammad: Emotional Faces in the Wild: Feature Descriptors for Emotion Classification.

Building on the growing popularity of recognition in images, Huthaifa Abuhammad, from the University of Exeter, presented a new approach for facial expression recognition of seven emotions. In fact, this method yields superior results, particularly for the perceptually hard-to-distinguish emotions such as 'fear', 'anger' and 'sadness'. Using random forests to train texture features (HOG, D-SURF and LBP) from static images, this approach identifies the most important feature types and locations for emotion classification.

Naoki Kato, from Keio University (Japan), discussed a novel approach for re-identification of individuals from videos (e.g., matching persons captured from different cameras). A fundamental task in video surveillance and applicable to human tracking systems, this novel approach exploits convolutional neural networks to learn the similarity of persons observed from a video camera. Experimental results reveal that this model achieves new state-of-the-art identification rates on the iLIDS-VID dataset and the PRID-2011 datasets of video clips.

Guodong Zeng, from the University of Bern (Switzerland), addressed the problem of automated skin lesion segmentation in dermoscopy images and was awarded first prize for best paper. Using Multi-Scale Fully Convolutional DenseNets (MSFCDN) and the ISBI 2017 *Skin Lesion Analysis towards Melanoma Detection* Challenge dataset, this proposed approach achieves quantitative accuracy scores that are better than existing methods such as U-Net and FC-DenseNets.

With heightened synergy between academia and industry, the second day of the conference saw an open discussion by a panel of researchers and industry leaders, about the application of artificial intelligence and machine learning, particularly in healthcare (e.g., diagnosis of radiographical images) and transportation (e.g., self-driving cars). The discussion, which was led by technology enthusiast and journalist, Hugo Correia from Global Media Group, explored the optimism and concern in relation to the rapid-growth of artificial intelligence-based tools in Computer Vision fields including detecting human behaviour, diagnosis in medicine and autonomous vehicles.
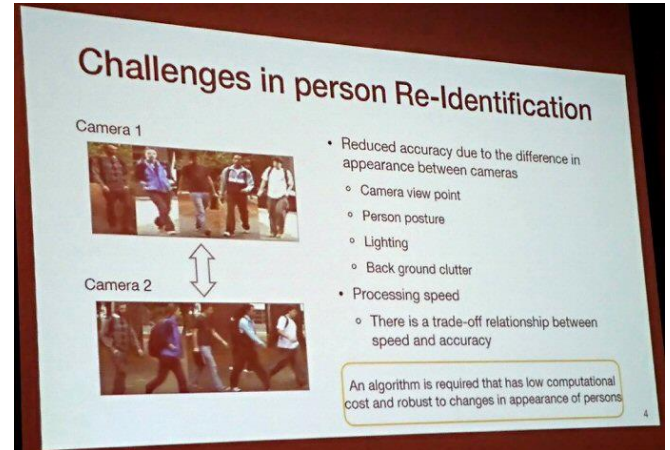


Naoki Kato: Video-Based Person Re-identification by 3D CNNs and Improved Parameter Learning.

There was definitely a balance for highlighting effective techniques that are not at present driven by deep learning. Alfredo Ruggeri, from the University of Padua (Italy) and co-founder of Resono Ophthalmic, an ophthalmic instrumentation company, discussed techniques from classical image processing and analysis, machine learning and pattern recognition and statistical analysis. The aim of these techniques is to provide ophthalmologists with a quantitative description of the main clinical parameters used in their diagnostic procedures. There was an interesting take on the idea of having full knowledge about image features that exists in classical image analysis as opposed to deep learning methodologies.

Towards the end of this particularly interactive conference, a bus tour brought to life the historical changes in Póvoa de Varzim and nearby Porto. This, coupled with authentic Portuguese banquet (Bacalau) that was served on a boat, further encouraged a lively discussion amongst researchers. Moreover, not only was the tour and banquet exceptionally well-organized but ended coincidently as locals started celebrating their national St. Peter's Day in the streets of Póvoa de Varzim.



Members of the conference having a banquet on a boat in Porto.

To say that the 15[th] International Conference on Image Analysis and Recognition offered a rich variety of cutting-edge research and innovations in computer vision would be

an understatement. By exploring the importance of classical and machine learning techniques in combination with ground-breaking deep learning methods, it was the opinion of several researchers (spanning multiple generations) that this conference 'opened their eyes' to further innovations in their field of interest, and the subsequent impact of their research contribution.

Hykoush Asaturyan
University of Westminster
email: h.asaturyan@my.westminster.ac.uk

# Inception Institute of Artificial Intelligence: A Bold Initiative to Fuel AI Research



The Inception Institute of Artificial Intelligence (IIAI), headquartered in Abu Dhabi, the capital of the United Arab Emirates, is an international research organization dedicated to achieving breakthroughs in fundamental and applied AI. IIAI consists of an elite team of handpicked scientists and engineers, led in their dedication to excellence by Oxford University trained CEO and Chief Scientist Dr Ling Shao, the former Chair Professor of Computer Vision and Machine Learning at the University of East Anglia.

Since its establishment earlier this year, IIAI has attracted over 60 members from leading universities and research organizations, including Oxford, MIT, ETH Zurich, Tsinghua University, Peking University, the Australian National University, the National University of Singapore, and New York University. IIAI continues to expand quickly, with plans to take on hundreds of researchers and engineers in the near future. In addition to our full-time team, IIAI will also accept a large number of graduate student interns from world-renowned universities.

> Artificial Intelligence: At the Heart of a Happier, Healthier, and More Productive Global Community

At IIAI, we believe AI is essential to tackling some of humanity's most challenging and pressing problems. To help our society reach its full potential, IIAI is leading fundamental and applied research across numerous domains, pushing the boundaries of AI capabilities. In fundamental research, we focus primarily on machine learning, with a particular emphasis on deep learning. Within this domain, we are developing more efficient and intelligent learning models, including unsupervised or self-supervised learning, few-shot or zero-shot learning, and lifelong learning.

Because we believe in the profoundly positive impact AI can have on society, IIAI is also investing extensively in applied research, with the aim of fundamentally transforming the use of AI across domains. IIAI's current applied AI research is focused on two primary areas: smart cities and healthcare. For instance, to accelerate the evolution of smart cities, we are developing a world-class, intelligent video-analytics platform powered by thousands of computing clusters to revolutionize object re-identification and activity recognition. In the healthcare sector, we are leveraging medical imaging and electronic health records to optimize the early detection, diagnosis, and targeted treatment of ailments, including diabetic retinopathy, a complication of the type 1 or type 2 diabetes prevalent across the Middle East, and breast cancer, which is most commonly detected by digital mammography.

> From Big Data to Cloud Computing, IIAI's Resources keep it at the Forefront of Research and Innovation.

To support its research, IIAI leverages powerful computing resources, including massive parallel-processing clusters with hundreds of NVIDIA DGX-class GPU servers, HPC resources consisting of thousands of CPU nodes, and petabytes of SSD-based data storage. In addition to these physical resources, IIAI also works with a wealth of big data from various domains, from text to audio, still images to full-motion video. The combination of our technological power and substantial data assets is instrumental in supporting not only the AI research being conducted, but also accelerating its application to multiple, wide-ranging fields and industries.

IIAI is rapidly making an impact on both the national and global levels, publishing numerous papers in top journals and at international conferences. Notably, four articles have been presented at this year's IEEE Conference on Computer Vision and Pattern Recognition (CVPR 2018), over ten articles were presented at the 2018 European Conference of Computer Vision (ECCV 2018), and two articles have been published in IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI). Furthermore, in order to promote the development of fundamental research in AI, IIAI has sponsored several international conferences, including ECCV, the International Conference on Machine Learning (ICML), and the Conference on Neural Information Processing Systems (NIPS).

To support its mission of expanding AI research across disciplines, IIAI is forming long-term strategic partnerships with numerous internationally renowned universities and organizations. This enables important academic research and knowledge to be combined with the experience and extensive domain expertise particular to industry, establishing the foundations for the future of AI, while at the same time promoting AI in the region – and beyond. Thus far, IIAI has initiated the establishment of joint laboratories with the University of Edinburgh, the University of Amsterdam, and New York University Abu Dhabi. These collaborations include joint doctoral programs, exchanged researcher visits, as well as joint research and development. In addition to academic collaborations, IIAI also maintains a close relationship with various research institutions globally, from Europe to Asia. Research within these collaborations is based on a broad spectrum of data resources, with the aim of promoting the rapid development of AI and data science for practical applications in a variety of fields.

> Join IIAI and be part of a one-of-a-kind Institute in a Visionary and Dynamic Nation.

The IIAI office is located in an international financial center, the Abu Dhabi Global Market, which is surrounded

by a wide range of businesses and entertainment facilities, as well as beautiful beaches and endless sunshine. In addition, IIAI's strategically central location at the intersection of the East and West offers the unique opportunity to work in a highly international community, which brings together a variety of perspectives and backgrounds while creating a diverse and open culture that is warm and welcoming to those from all backgrounds and nationalities.

IIAI is currently recruiting a select group of expert researchers and engineers highly specialized in the fields of deep learning, computer vision, natural language processing, and medical image analysis. Well-qualified researchers and engineers who are passionate about AI, driven towards excellence, and have a spirit for adventure are encouraged to contact IIAI.



Job portal:
https://www.jobs.ac.uk/enhanced/linking/inception-institute-of-artificial-intelligence/
IIAI website: http://www.inceptioniai.org/
contact: jobs@inceptioniai.org

## BMVA Computer Vision Summer School

This event took place on 2–6 July 2018 at UEA. As usual it was very well organised and people who attended (including the lecturers!) enjoyed the experience.

The first photograph shows students paying a high degree of attention to studying the posters during a poster session, this being an excellent way to learn what they should also be capable of!







The second photograph shows BMVA Chair Adrian Clark presenting the award for Best Poster to Julien Schroeter, who received a certificate together with a copy of the 5th edition of Roy Davies's 2017 book on Computer Vision. The last photograph shows a number of delegates holding up their all-important certificates of attendance.

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

**STOP PRESS**

As was announced immediately following the awards and presentations at the BMVC 2018 banquet, BMVC 2019 is to be hosted by Cardiff University!

Potential authors and delegates should keep their eye on the official website: https://bmvc2019.org



British Machine Vision Conference 2019

9–12 September 2019

Cardiff being my home town and having many interesting sights – many more than when I grew up there – I am very happy indeed to endorse this choice! – Ed.