

BMVA News

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

Volume 23 Number 4
June 2013

Editor: Professor Roy Davies
Department of Physics
Royal Holloway, University of London
Egham, Surrey, TW20 0EX
Tel: +44(0)1784 435064
Fax: +44(0)1784 472794
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 September 2013.

Contents

Editorial: <i>How Recognisable is an Evolutionary Dead-End?</i>	1
BMVA Distinguished Fellow 2014 – <i>Call for Nominations</i>	2
BMVA Thesis Archive.....	2
Report on ISBI 2013	3
Report on FG2013.....	4
Reports on ICRA 2013.....	6
.....	8
Phil Torr is moving	10
Message from the BMVA Technical Meetings Organiser.....	10
Quantitative Image Analysis for Astronomical Applications.....	10
Vision in an Increasingly Mobile World: Review of a Recent BMVA Technical Meeting.....	11
MIUA 2013 – Call for Participation.....	11
Real-Time Image and Video Processing 2014	12
Travel Bursaries for International Conference Attendance	13
Announcements: <i>The Queen's Birthday Honours List</i> 2013.....	13
BMVC 2013 – Call for Participation.....	14
Fifth International Conference on Imaging for Crime Detection and Prevention (ICDP 2013).....	14

¹The 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.

Editorial: *How Recognisable is an Evolutionary Dead-End?*

It is part of human nature to assume that any institution you join has always been the way it was when you started: only subsequently do you learn the real state of affairs. Thus it was when I arrived on the scene in 1975, and our subject was called (visual) pattern recognition, which in those days was divided into three parts – statistical PR, structural PR and syntactic PR. In fact, at that time statistical PR was gradually turning from a difficult subject with many key papers and few books into a much more understandable one supported by a plethora of readable texts. Key to this development was the widely known volume by Duda and Hart (1973), which in many ways helped the subject to become fossilised. I say ‘fossilised’ advisedly, as there is a substantial core part that has changed little since then, though other parts have now been added – artificial neural networks (lately metamorphosed into machine learning systems) and the bag-of-words/dictionary approach, to name but two. I ought to add that support vector machines (SVMs) are nowadays vitally important and it may seem surprising that they only arrived upon the PR scene around 1996.

But what of *structural PR*? That started life as image processing, mainly using simple neighbourhood operations in digital lattices. Here Rosenfeld's early (1969) book was influential in the development of the subject, soon being replaced by the even more major 2-volume work by Rosenfeld and Kak (1982). Not surprisingly, a lot of the ideas were geometrical in nature and appeared – many of them in analogue form – in Part II of Duda and Hart's book, thereby giving valuable added insight into shape analysis, e.g., using grassfire transformations, convex hulls and Euler numbers.

The other member of the PR triplet, *syntactic* PR, was the prime province of King-Sun Fu of Purdue University (remembered by the IAPR in its biennial K.-S. Fu Prize). His subject was a natural one for decoding, explaining and thus recognising characters (in OCR) and other parsable entities such as fingerprints and ECG waveforms. The hallmark of syntactic recognition is the development of grammars that can regenerate the patterns being recognised. In fact, for some years after Fu's (1974) book was published, syntactic PR was quite strongly linked with structural PR. However, by now structural PR is what we all do by structural analysis of images and videos, using methods such as Hough transforms and RANSAC, and syntactic PR seems to have taken a back seat: but has it also become a fossilised dead-end topic?

While planning this editorial, I bore in mind that I had previously examined the archaeological terrain of computer vision and felt I needed a fresh analysis from the viewpoint of evolution. It is of course well-known that evolution embodies such concepts as missing links, evolutionary dead-ends, fossilised species, periods of rapid evolution, periods of relative stasis, accidental mixing of genomes – not necessarily as in “The Fly”, but more naturally as interbreeding of species such as humans and Neanderthals,² or even bacterial DNA becoming introduced into the human genome – as indeed appears to have happened in the case of mitochondrial DNA. I had thought it would be easy to identify examples of all these in our subject, but maybe one needs to look back 100 years or so later to make a good job of it. Actually, some instances of dead-ends include the once ubiquitous chain code, which dates from the time when each bit of information was an expensive commodity, whereas nowadays the emphasis is on robustness and high accuracy rather than on storage: i.e., we can now afford to do things properly without caring too much about hardware limitations. Similarly, use of centroidal profiles for shape analysis has long been eschewed for lack of robustness. Note also that ‘everyone’ is using SVMs instead of the nearest neighbour algorithm, which in spite of being within a factor 2 of the Bayes error rate seems to be in abeyance. And even invariance in the shape of cross-ratio based methods has disappeared from view – though narrower forms of invariance remain firmly embodied in SIFT and its ilk. (However, I have long had the feeling that people are curiously attached to fashionable names ...)

(continued on p.15)

²Apparently, Europeans and Asians have ~2.5% Neanderthal DNA: <http://discovermagazine.com/2013/march/14-interbreeding-neanderthals> (website accessed, 10 July 2013).

BMVA Distinguished Fellow 2014 – Call for Nominations

The BMVA Executive Committee seeks nominations for the *Distinguished Fellow 2014* award. This prestigious award is given to one person only each year in recognition to his/her services to the British Machine Vision community. The nominees must be distinguished researchers, based in the UK, who have contributed significantly to the field of research and the reputation of the British Machine Vision Community both nationally and internationally. Nominations, *with a few lines of rationale*, should be sent to Professor Roy Davies, Chair of the Distinguished Fellow Award Panel, by 1 September 2013.

Professor Roy Davies
Chair, Distinguished Fellow Award Panel
email: e.r.davies@rhul.ac.uk

BMVA Thesis Archive

In order to promote and improve access to the large base of high quality PhD research undertaken in Computer Vision in the UK, the BMVA maintains an online repository. This provides a single source archive of all past, current and future PhD work undertaken in this area in UK academic institutions.

The service allows students to quickly and easily share the results of their work with the Computer Vision community, nationally and internationally, and it is a useful database for searching and reviewing previous PhD research work undertaken in the UK.

The real value of this service can only be realised if the UK community support the effort and so the BMVA would like to encourage all members of UK academic institutions to contribute material to the repository. Contributions are required to be in PDF format and supplements such as videos and images are welcome.

The PhD repository can be accessed through the main BMVA website (www.bmva.org). If you have any problems submitting your thesis to the repository please contact Aphrodite Galata.

Dr. Aphrodite Galata
University of Manchester
email: a.galata@cs.man.ac.uk

Report on ISBI 2013

The 10th IEEE International Symposium on Biomedical Imaging (ISBI) was held in Westin Hotel in San Francisco, California, on 7–11 April 2013. The objective of ISBI was “to integrate different imaging communities and to contribute an integrative imaging approach accessing all scales of observation” (the conference was titled “From Nano to Macro”).

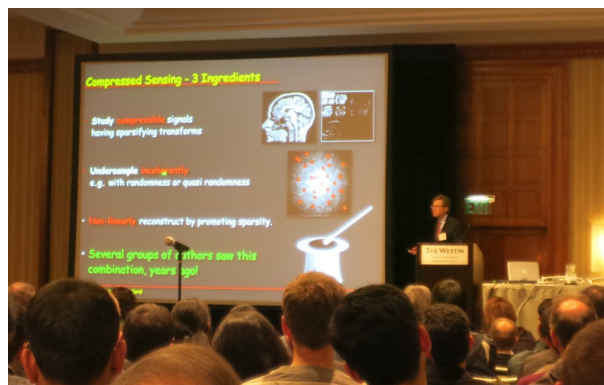
The meeting was multi-track, featuring preconference tutorials, followed by a strong scientific program of plenary talks and special sessions as well as oral and poster sessions of peer-reviewed contributed papers. Out of nearly 400 accepted papers, 140 were presented orally and the rest in poster form. The conference covered many aspects of biological and medical imaging, from mathematical image reconstruction for medical applications to integration of image-based modelling with genomic data for hypothesis generation and network inference.



Plenary session audience

The conference welcome address was given on Monday lunchtime and was followed by a talk by the first plenary speaker, John Sedat, on imaging in biophysics and biochemistry. The afternoon was filled with multi-track oral, and then poster sessions. The following days were structured similarly, with additional multi-track poster and oral sessions in the morning.

The conference hosted five plenary speakers. My favourite talk of the symposium was by David Donoho from Stanford University. He is a mathematician with fundamental contributions in theoretical and computational analysis. He started with an introduction to compressed sensing, pointing out how all the algorithms published in the field can be broken down into 3 steps (study signals, under-sample, reconstruct) and how each paper is a variation of one or more of the steps.



David Donoho's plenary talk

He next talked about this in the big picture perspective of vision, stressing that most computational research is very hard to reproduce, largely because the code is rarely made available. And thus, publications will not have as high impact as desired. As a solution, he invited people to publish their code along with their papers, and referred to a website, reproducibleresearch.net gathering information and links about reproducible research.

Another plenary session I liked a lot was by Simon Cherry – about making better PET scanners – where he went deeply into imaging physics and clearly drew out the technical challenges of building a longer PET scanner for increased 3D spatial resolution.

Several poster and multi-track talk sessions, 1.5 hours each, were held in parallel and involved various aspects of medical and biomedical imaging. The sessions covered a wide range of techniques used to analyze images of various biological structures and scales.

Out of the many interesting papers, I would like to highlight two here, both on efficient recognition of biological structures in 3D. The first, “Spatially Adaptive Random Forests”, by Geremia, Avache et al. introduces a general and efficient solution to the problem of classifying relevant structures in medical images. This task is challenging due to the large size of data volumes, and the severe overlap of object classes in feature space. To keep the computation time down, the data needs to be sampled, even though sampling reduces classification success. The authors propose a supervised learning algorithm that learns the optimal image sampling during training associated with the classification task, and demonstrates top performance in the context of multi-class brain tumour segmentation in 3D MR images.

The second, “3D Haar-Like Elliptical Features For Object Classification in Microscopy” by Amat and Keller, presents an extension to the Haar wavelet framework by adding orientation, illumination and scale information by assuming that the neighbourhood

surrounding points of interest in the image can be described with ellipsoids. Haar wavelets are commonly used for efficiency but they are not very discriminative. In this paper, the discriminative power is increased by incorporating edge and shape information into the features, and computation time is reduced by GPU-based implementation.

I gave my presentation on the Wednesday morning of the conference. I was allocated 18 minutes for my talk and Q&A in the “Registration 2” session. The session was attended by almost 100 people and my talk was on landmark localization in CT scans of variable fields of view. This was the first time I presented at a major conference. I enjoyed the whole experience and got some very useful questions and feedback from the audience.

I am very grateful to BMVA for providing the travel bursary that supported my attendance at this beneficial and interesting international conference. While lots of the research was not directly related to what I am doing, it was great to see how it fits into the bigger picture in imaging, and to pick up some ideas from other fields.

Meelis Lootus
Oxford University
email: meelis@robots.ox.ac.uk

Report on FG2013

Ni Hao! Hello from the 10th International Conference on Automatic Face and Gesture Recognition (FG) held in Shanghai, China from 22–26 April 2013. This conference is the premier international gathering for researchers in the field of image and video-based face, gesture, and body movement recognition. This year more than 300 submissions were received for the main conference of which 36 were accepted for oral presentations and 68 were accepted for poster presentations, with an acceptance rate of about 35%.



View of the Skyscrapers from the Bund

The program comprised three keynote presentations, twelve single-track oral sessions including three special sessions on “Frontiers in Face and Gesture,” and three poster sessions with summary spotlight presentations.



Majesty Plaza Hotel, venue of FG2013

The conference was hosted at the Majesty Plaza Hotel in the central district of Shanghai, a prime location in the heart of a city of 23M people. The city has seen massive growth and development since 1990, the history being detailed through photos, scale models and maps in the Shanghai Urban Planning Exhibition Centre. For some in our group the radically different culture came as a shock but we soon became accustomed to the offers of fake watches, massages, lasers, i-phones and tea. The locals were quite willing to try to communicate through sign language, or would listen patiently to awkward attempts of phrases from a book or cheat sheet. Evenings on East Nanjing Road were a special treat, with large groups of older women congregating to do T'ai Chi, many performing artists entertaining crowds, and karaoke on a grand scale with hundreds participating, including passers by.

The keynote lectures stood out as highlights, with a common focus of addressing real-world, often cross-disciplinary problems. The first keynote was by Professor Gregory Abowd on the use of Computer Vision to address problems relating to autism. He


described his journey as a researcher in ubiquitous computing, having a child who began to show, and was later diagnosed with autism, and how he saw that computers could help to automate many time-consuming tasks and assist in bringing earlier and better diagnoses. It was a fascinating insight into a very practical problem consisting of real-life data. What makes his work more interesting is the cross-disciplinary nature which requires a long term view, since it takes years to build up relationships across disciplines. As an example he showed how machine vision can be used to automatically determine child head orientation using a top-view camera. The child's name is called by an adult sitting at 90° to the child. Failure to respond to its name by turning its head towards the adult when called is widely recognised as a strong early sign of autism. Since most autism is diagnosed at 4–5 years of age, helping parents and medical professionals to diagnose and manage child behaviour earlier on would significantly help parents who often comment that they thought behaviour was different from as early as 1 year old. Perhaps the most difficult task for computer vision researchers is helping psychologists understand the tools that computer vision can provide and to encourage them to think about their problems in new ways.

The second keynote was on the topic of deception. Professor Judee K. Burgoon is well-known for her work on nonverbal and verbal communication, deception and computer-mediated communication. What made the lecture interesting was the approach of taking a broad, realistic view with natural data. Although the results of automated systems shown weren't very good, what was most interesting is that human 'experts' achieve an average deception detection rate of 53% – slightly more than chance! It is a sign that the field has matured that they're addressing the real problem but there is clearly much more research to be done before the problem can be considered solved.

The final keynote was by Professor Jessica Hodgins, who is at the forefront of research on computer graphics, animation, and robotics with an emphasis on generating and analyzing human motion. This third lecture was on modelling and synthesizing facial and whole-body animation. After reviewing the state of the art and the progress made during the past 20 years, recent work to make characters appear natural and lifelike was shown. She drew interesting parallels between face and whole-body animation, going into some detail on the commonalities and differences of both.

Four students from the University of Surrey supervised by Professor Richard Bowden presented work on upper body and hand-pose estimation, eye-gaze tracking, emotion recognition and sign language recognition.

There were lots of interesting talks, particularly around the topic of facial recognition. It was fantastic meeting all the people from the community which is rather smaller than vision in general making this a very focused conference. Our opinion was that although the contributions to other top conferences are of a very high quality, much of the work can be unrelated to an individual's research whereas at FG it seemed that the overall quality of the contributions was at a slightly lower standard but having such a focus meant that most of the work was relevant to all the participants.



Animated discussion during a poster session

The banquet was held at the Shanghai International Convention Centre overlooking the Bund, the waterfront buildings dating from the British colonial era. As with many Chinese meals, the dishes were varied and the traditional culture of eating from the same dishes of food reflected our shared trust. During the banquet a new award was unveiled for long-term impact in the field of Face and Gesture. It was awarded to William Freeman and Michal Roth for the 1995 paper entitled "Orientation histograms for hand gesture recognition". Although Bill couldn't be there to accept the award in person, he recorded a video thanking FG for the award and showing the original program running and the experiments. It was striking both how the field has progressed since then, but also how the ideas developed in that paper are still relevant today.



Local entertainment at the Banquet

In conclusion, FG2013 was a fantastic, well-organised conference, hosted in one of the most dynamic and diverse cities on earth. Everything from the food, the keynotes, presentations and posters was thoroughly enjoyable and it was great to be able to meet and exchange ideas with other researchers in this field. I would like to thank BMVA for providing a very generous travel bursary to allow me the opportunity to participate in the conference.

For more information about FG2013, please visit <http://fg2013.cse.sc.edu/index.html>.

Brian Holt
University of Surrey
email: b.holt@surrey.ac.uk

Report on ICRA 2013

From 6–10 May this year, the International Conference on Robotics and Automation (ICRA) was held in Karlsruhe, Germany. ICRA is one of the world's largest robotics conferences, with 873 accepted papers and almost 2000 attendees. This made it a truly large-scale conference, much bigger than any I have attended in the past. The conference took place at the Karlsruhe Kongresszentrum, a large conference venue situated very near the city centre. Spread over three floors, including six rooms for presentations, a floor for interactive sessions, two exhibition spaces, and a main lecture hall big enough to accommodate all the attendees, this was a very impressive location.



The conference venue

Participating in this conference was a very interesting experience, partly since robotics is not my main area of research. There were plenty of interesting presentations from a very diverse range of subjects, ranging from stability of bipedal robots, to designs for bio-inspired soft robots, to applications of robots in space exploration. A good many papers were also presented on computer vision topics, with the first session I attended being on recent advances in simultaneous

localisation and mapping, a topic I've been involved with for many years.



Inside the huge conference centre

One of the highlights of the conference was the excellent series of plenary talks – hour-long keynote lectures by some of the biggest names in robotics research. On the first day we were treated to a fascinating lecture by Robert Wood of Harvard School of Engineering and Applied Sciences, describing his group's work on creating robotic bees. Especially interesting was their fabrication method for such tiny machines, which ordinarily requires precise craftsmanship that takes years of training. Instead, they have developed a method involving printing and etching thin layers of material, encoding in the shape the ability to fold itself into the right shape when pressure is applied. This is a powerful technique, enabling complex shapes on very small scales to be created reliably and repeatably – an essential criterion when the resulting robotic bees have a tendency to crash land.

Another fascinating plenary talk was given by Aude Billard of EPFL, Lausanne. The subject was their recent efforts to get robots to catch objects; having decided that catching tennis balls was too easy, her group are developing methods for robots to catch thrown tennis rackets, half-full bottles of water and other complex objects. The results are most impressive, and one of the few instances where videos of robots need to be slowed down to be watchable.



A robot made in Karlsruhe

Rather than regular poster sessions, ICRA featured ‘interactive sessions’. In these, presenters were given a large display screen and a table for a laptop, allowing them to display the content of their choice. This was how I presented my paper, and compared to regular poster sessions I found it to be quite an interesting and liberating experience. Rather than making one static poster, I was able to show different poster pages to highlight different aspects of the work, and even to show videos. Some of the highlights from these sessions included videos of an octopus robot swimming through water, and a demonstration of a snake robot able to climb up vertical pipes (and peoples’ legs).

An interesting aspect of the conference, unlike those I’ve attended before, was the large number of exhibitions and displays. These included two large spaces populated by companies exhibiting their wares, which gave us all an excellent opportunity to look around and see some of the cutting edge developments

in robotics. This included a stand by Aldebaran Robotics, makers of the well known Nao robot, one of whom was casually playing on a swing set up in their booth; and at least one Reem humanoid robot, from PAL Robotics. There were also a number of small wheeled robots trundling around the floor space throughout the conference, including one carrying a tray of Toblerones.

As well as the wonders of the conference, the city of Karlsruhe was a very pleasant place to spend several days. One of the prominent features of the city – and clearly visible from any map – is the large circular park surrounding the magnificent Karlsruhe Palace. Just behind the palace are the botanical gardens, and beyond that, within the outer ring of the circular feature, a large forest. Karlsruhe is also famous for its pyramid, standing in the main market square – a rather unexpected sight in a German city. The pyramid is located over the vault of Margrave Charles III William, the founder of the city, and was built in the early 19th century. As well as all this, there was a very nice zoo next door to the conference centre, featuring at least four elephants and countably many penguins – to which the attendees of the conference apparently had free access.



The Karlsruhe Pyramid

Overall the conference was a very enjoyable experience. Not only was it a good opportunity to present my work to a slightly different kind of audience, but I learned a lot about the fascinating developments in robotics. I would like to express my gratitude to the BMVA for funding me to attend this conference.

Osian Haines
 University of Bristol
 email: haines@cs.bris.ac.uk

Report 2 on ICRA 2013

The IEEE International Conference on Robotics and Automation (ICRA 2013) is the annual flagship robotics conference sponsored by the IEEE Robotics & Automation Society. It is the world's largest robot research conference which was organised this year by Karlsruhe Institute of Technology in Karlsruhe, Germany. This was the first time the conference was held in Germany. Every year this conference brings together loads of academics and industrialists from across the world to share ideas, concepts and technologies. ICRA 2013 was held from 6–10 May 2013, consisting of two days of workshops/tutorials, three days of technical presentations with six plenary talks, exhibitions, technical and non-technical tours and multiple robotic challenges. More than 800 oral and poster papers were presented in the main conference. Along with the main conference, there were 36 workshops and tutorial sessions held on the first and the last day of the conference. The conference spread across a wide spectrum of robotics research including medical robotics, interactive/manipulative robotics, industrial robotics, vehicles and automation, computer vision, scene analysis, SLAM, human robot interaction, rehabilitation robots, underwater robotics and space robotics, to name just a few. There were six plenary talks along with many invited speakers in the different workshops. In all, the conference was a nice platform for any researcher to understand the current state of art robotics research as well as gaining information on how different disciplines like computer vision, machine learning, semiconductor manufacturing and biology are being applied to robotic research.

The conference started on 6 May with multiple workshop sessions. An interesting workshop was organised by Mark Toussaint (University of Stuttgart) on Autonomous learning. In the workshop, Siddhart Srinivasa (CMU) gave a talk on weakly supervised learning for object discovery used in robotics applications. He also talked about lifelong learning for robots in the context of object discovery. Their method tries to bridge the robotic community where domain knowledge is important for particular problems and the computer vision community where we assume a more general setup. He was using domain information as a metadata in their optimisation. Another interesting workshop was on topics like surgical vision for robotics, multi-agent systems, and nano-robotics. There were technical tours to the MPI in Tübingen and the Fraunhofer IPA in Stuttgart, and to robotics labs at the Karlsruhe Institute of Technology. The workshop ended with a social gathering which included a welcome speech by the IEEE RAS president and the organising committee. This was followed by a spectacular firework display and a quadrotor UAV show in the night from

Roland Siegwart's team at ETH Zurich using quadcopters from Ascending Technologies.



Quadrotor show by ETH Zurich team

The main conference started on Tuesday 7 May. There were talks on scene analysis, sensing for grasping, SLAM, medical robotics, human robot interaction and computer vision. An interesting plenary talk was given by Robert Wood (Harvard University) on small-sized large-scale flying robots. He spoke about developing robots the size of houseflies and the common challenges associated with that. The main problem is the lack of the off-the-shelf technology to use for building the 'bots. After the plenary talk there was another session of talks and interactive presentations. The conference also had exhibitions where the industry leaders in robotics like Kuka, Schunk, Clearpath and many more were showcasing their latest products. The day ended with a 'welcome' reception and the ICRA midnight show at the concert hall.



Firework display during the opening ceremony

The second day of the conference started with talks on object recognition, rehabilitation robots and kinematics. There was a plenary talk by Rodney Brookes (Emeritus Professor, MIT) who talked about bringing robots to the household. He is the CTO of Rethink robotics which makes Baxter robots for human-machine interaction. These robots can be used in assembly lines along with humans to improve industrial productivity. As of now most of the industrial robots are actually used in the automobile industry while most of the other industries rely on human labour. Rodney Brookes explained how robots can be used alongside humans in the assembly

line. His previous company iRobot made the cheap home vacuum cleaner which was instrumental in making household robots more affordable. Rodney Brookes talked about creating smarter, more adaptable, low-cost robotic solutions which can help manufacturers to improve efficiency, increase productivity and reduce their need for off-shoring. The day finished with a banquet dinner and a special music show by the IEEE RAS band.



Industrial Exhibits at ICRA 2013

The last day of the main conference was again a packed day with multiple talks on SLAM, multi-legged robots, nanorobots, biologically inspired robots, Aerial/Marine/Space robots and multiple interactive sessions. The day had also a plenary talk by Michael Black on recent neurocomputing work at MPI Tubingen with collaborators from Brown University and Stanford University. His talk covered the neural activities in the brain which result in general human motion. They even tried to do this by learning the neural signal patterns in a specific area of the brain so that human motion can be controlled and simulated. He showed experiments on studying brain activity for human gaze control and for more general motion in animals. The last day of the main conference was also notable for a ceremony where different stalwarts in the field received awards for their contributions in the field. Ruzena Bajcsy from UC Berkley received a Lifetime Achievement award for her pioneering work in early machine vision. Other people among the awardees included Radu B. Rusu from OpenPerception for his contribution to society in the form of open source libraries like PCL and ROS.

The final day of the conference saw another round of multiple workshops and tutorial sessions. One interesting tutorial was on Object Perception with Point Cloud Library (PCL) with speakers talking about different 3D features in the library, registration of point clouds, object detection and people recognition. Other important workshop was on Vehicle Autonomy in Urban Transportation systems. This workshop mainly focussed on registering point clouds from LIDAR data, detection of moving vehicles in the road, pedestrian detection, and perception and mapping of road scenes. Another workshop on long-term autonomy focussed on how the robot can learn its environment better over time. It consisted of talks on creating large databases of object models, environment maps on an internet scale and showed how to use existing knowledge through cooperation by Mohanarajah Gajamohan, (RoboEarth project, ETH Zurich). There were also talks by Wolfram Burgard (University of Freiberg) on pedestrian-like navigation through crowded city centres and about SeqSLAM (Niko Sunderhauf, Chemnitz University of Technology) where challenges were discussed about performing SLAM on a 3000 km track under varying weather conditions.

Overall, the conference was a nice overview of robotic research going on in academia and industry across the world. Details about the conference can be found at <http://www.icra2013.org/>. I would like to thank BMVA for providing me support to attend the conference.

Sunando Sengupta
Oxford Brookes University
email: ssengupta@brookes.ac.uk

Phil Torr is moving ...

I have been reliably informed that Philip Torr will be leaving Oxford Brookes after nearly 10 years to take up a position as Professor of Engineering at Oxford University (St Catherine's College).

He will setting up a new computer vision group there and says that he'll be looking for bright people both as PhD candidates and post-docs.

I am sure that the very many people who know Phil will want to join me in wishing him many more years of success and ground-breaking work at his new professional home!

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

A Message from the BMVA Technical Meetings Organiser

The British Machine Vision Association organises a series of one-day technical meetings (symposia) in London. Each meeting is chaired by a member of the UK academic computer vision community and addresses a specific topic in Machine Vision. The meetings usually have 6–10 speakers and are attended by 50–100 people.

The program for 2013–2014 is currently:

- 9 Oct: Quantitative Image Analysis for Astronomical Applications, *Chair*: Neil Thacker
- Feb 2014: Student Meeting, *Chair*: Simon Hadfield

Note that there is room for more meetings, especially late this year and next. This is not an onerous task and involves (i) organizing a program of speakers (by call for papers or invitation), (ii) chairing the meeting on the day. All other arrangements (finances, food, registration, room-booking, etc.) will be taken care of by myself.

Please visit the BMVA meeting webpage for up-to-date information on scheduled meetings, programs and CfPs:


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

Andrew Gilbert
BMVA Technical Meetings Organiser
email: a.gilbert@surrey.ac.uk

Quantitative Image Analysis for Astronomical Applications

This BMVA Technical Meeting will be held in London on 9 October, and the Chair will be Dr. Neil Thacker (neil.thacker@manchester.ac.uk).

Call for Participation



Images from space history

We invite potential speakers and attendees to discuss the application of image analysis techniques to astronomical data. The event will bring together imaging science experts and planetary science researchers to explore the importance of quantitative data analysis methods in the interpretation of planetary surface images. There will be an emphasis on identifying quantitative planetary science applications and data analysis techniques which can provide scientifically meaningful outputs, e.g., applications where well-defined measurements must be taken in the presence of noisy data, requiring methods capable of giving quantitative error assessments.

Participants from planetary science backgrounds are encouraged to bring specific questions which can potentially be answered through the study of surface images. Participants from imaging science backgrounds are encouraged to bring quantitative statistical techniques and algorithms which could find utility in answering such questions. This meeting may be of interest to a wide audience, including those involved in the analysis of any type of astronomical image data, e.g., space telescope imagery, crowd sourced (Zooniverse) projects, etc.

Call for Participation

All those interested in presenting material at this meeting are invited to submit a paragraph describing their topic to Neil Thacker for consideration by 28 July.

Andrew Gilbert
BMVA Technical Meetings Organiser
email: a.gilbert@surrey.ac.uk

Vision in an Increasingly Mobile World: Review of a Recent BMVA Technical Meeting

In May we had a very successful Meeting run by Toby Breckon (Cranfield University) on Vision in an increasingly Mobile world, attended by 80 people from Industry and Academia.

My first technical meeting as a non-student/real world worker ...

Kicking things off was Dina Damen³ presenting the work of Pished Bunnun. A demonstration of real-time detection of fairly featureless objects was fun which included a highlight for me: ‘hallucinations’ is the new term for false positives! I wish I’d used that for my thesis ... out with the false-positive rate, in with the hallucination rate. Matthew Brown⁴ presented an in-depth history of AutoStitch for iPhone (coming soon to Android). It was interesting to see the approximations and tricks that had been implemented. Though, on the tube afterwards, I was kicking myself for not asking if it made any money. Jose Rivera-Rubio⁵ discussed his work on assisting the visually impaired understand their surroundings. In particular the generation of a database of images that can be used to test recognition algorithms.



Dina Damen (Bristol University) presenting on real detection of featureless objects

Renato Salas-Moreno⁶ presented his work on SLAM++ which moves Andrew Davison’s work up to the level of objects. This work is/was also presented at CVPR 2013.

Following a rather splendid lunch (worth the £5 – in the past I’ve nipped out for a burger but this time I should have brought a doggy bag) we entered the afternoon

sessions. First up was Ioannis Katramados⁷ who discussed his experience of crowd-funding through Kickstarter⁸ – timing seems to be very important. This was followed by Will Maddern from the Oxford Mobile Robotics Group who presented an overview of the RobotCar that has been in the news recently⁹. Having seen a similar presentation about the Google car at CVPR last year it looks to me like Google are not going to have everything their own way.

A short break for tea and biscuits, then Calum Blair¹⁰ presented an analysis of CPU vs. FPGA vs. GPU in detection of people using HOG – Which is fastest? Which uses most power? Can you split the task up between processors to optimise the performance? Calum’s presentation is available on his website.

Finally, Rui Hu¹¹ presented her work on using a free-hand sketch to retrieve images from a database – if I draw a house I should get back images of houses. Interestingly this was extended into the realm of video – can a sketch retrieve relevant videos?

So, as I sat on the tube to Cockfosters, I would say it was a good day. I would recommend these one-day technical meetings – the talks are always interesting; the location is superb; the food is pretty good; and you can take in the sights and sounds of London afterwards.

Greg Flitton
Cranfield University
email: greg@flitton.net

MIUA 2013 – Call for Participation

The 17th Annual Conference on Medical Image Understanding and Analysis will be held in Birmingham on 17–19 July 2013. The conference chair is Professor Ela Claridge of Birmingham University.

MIUA is the principal UK forum for communicating research progress within the community interested in medical and biological image analysis. This successful series began in Oxford in 1997. The meetings are designed for the dissemination and discussion of research in this rapidly expanding area, which is notable for the range of research communities involved. The meeting aims to encourage the growth and raise the

³www.cs.bris.ac.uk/~damen/

⁴www.cs.bath.ac.uk/brown/

⁵www3.imperial.ac.uk/people/jose.rivera10

⁶www.doc.ic.ac.uk/~rfs09/

⁷www.cosmonio.com

⁸www.kickstarter.com

⁹<http://mrg.robots.ox.ac.uk/robotcar/>

¹⁰<http://home.eps.hw.ac.uk/~cgb7/>

¹¹<http://personal.ee.surrey.ac.uk/Personal/R.Hu/>

profile of this multi-disciplinary field by bringing together the various communities.

MIUA is a single-track conference with oral and poster presentations and an excellent programme of keynote lectures from leading international experts in molecular imaging, translational research and machine learning for quantitative image analysis. We are grateful to the following three speakers who have agreed to give keynote lectures at the conference:

- Professor Boudewijn Lelieveldt (Leiden University Medical Centre, Netherlands)
- Professor Daniel Rueckert (Imperial College London)
- Professor Milan Sonka (University of Iowa, USA).

Pre-conference workshop

This year's pre-conference workshop (17 July) on optical molecular imaging will be given by Dr. Hamid Dehghani, one of the leading experts in the field and co-developer of the NIRFAST package for modelling of light propagation in tissue (<http://www.nirfast.org/>). The tutorial will be particularly beneficial for research students and early career researchers.

Conference dinner

The venue for the conference dinner (18 July) is the Edgbaston Cricket Ground pavilion. There will be an opportunity to have a guided tour of the ground and its facilities.

Important dates

Author registration deadline:	30 May
Early registration deadline:	13 June
Conference workshop:	17 July
Main conference:	18–19 July

Conference registration

Registration for MIUA is now OPEN! You may complete your registration via the on-line registration shop at:

<http://events.cs.bham.ac.uk/miua2013/registration.php>

Registration includes the Molecular Imaging tutorial on 17 July, tea/coffee and lunch on 18 and 19 July, the conference dinner on 18 July, conference proceedings and access to all sessions.

Registration fees

Full:	£290 (early rate)	£310 (late rate)
Student:	£130 (early rate)	£150 (late rate)
Single day:	£130	

Accommodation (£35 per night, en suite room + breakfast) is not included in the fee but can be booked via the on-line registration shop.

Further information

For further information, see the conference website <http://www.miaa.org.uk> or contact us by email at miua2013@cs.bham.ac.uk

Professor Ela Claridge
University of Birmingham
email: e.claridge@cs.bham.ac.uk

Real-Time Image and Video Processing 2014



Conference EPE115

The SPIE Photonics Europe Conference EPE115 on Real-time Image and Video Processing will be held on 14–17 April 2014 in Brussels (Belgium). The conference chairs are well known in the area of Real-Time Vision and are:

- Nasser Kehtarnavaz (University of Texas at Dallas, USA)
- Matthias F. Carlsohn (Computer Vision and Image Communication at Bremen, Germany)

Background to the conference

Real-time image and video processing involves algorithmic, hardware, and software aspects of making an image or video processing system to operate in real-time. The SPIE Real-Time Image and Video Processing conference has been the only conference dedicated to real-time (RT) aspects of image and video processing. This conference has been providing a field catalyst bringing together scientists and researchers from industry and academia working in RT image and video processing to present recent research results pertaining to new RT algorithmic, hardware and software approaches as well as RT system designs and applications.

Submission topics

Papers addressing RT issues are solicited but not limited to the following topics:

- RT and video processing algorithms
- RT mobile or embedded image/video processing systems
- RT image and video processing hardware including FPGA, DSP, GPU, GPP, ASIC, SoC, and SiP implementations
- RT software optimizations and related design paradigms
- RT image and video processing via parallel processing and related computer architectures
- RT computational photography, augmented reality, and 3D applications
- RT image and video processing for stereo and multi-view camera systems and multi-modal imaging sensor fusion
- RT depth acquisition methods for 3D digital imaging and video, e.g., RT RGB+D (i.e., Kinect-based)
- RT image and video compression and coding for storage and broadcasting applications including HDTV, HbbTV, UltraHDTV, 4K-imaging
- RT image and video trans-coding and streaming technologies for future internet applications and emergent image resolutions and upcoming standards
- RT image and video processing applications including digital, cell-phone, smart and system cameras
- RT image and video processing for automatic visual inspection and machine vision
- RT HDR image processing
- RT image and video processing for CCTV applications, intelligent surveillance and security including biometric imaging
- RT image and video processing for robot vision, autonomous systems, human-machine and machine-machine interaction
- RT image and video processing for multi-dimensional image analysis, spectral and hyperspectral imaging, and remote sensing
- RT image and video processing for immersive systems using VR, AR and mixed reality
- marker-less RT tracking of objects in imaging and video animations for VR, AR and MR systems and applications
- RT aspects in medical imaging.

See p. 15 for details of the Programme Committee. For further information, refer to the conference website <http://spie.org/x12290.xml> or contact Professor Dr. Carlsohn by email.

Professor Dr.-Ing. Matthias F. Carlsohn
email: matthias.carlsohn@t-online.de

Travel Bursaries for International Conference Attendance

The BMVA provides a number of travel bursaries for student members of the BMVA who are research students at UK institutions to present their work at significant international conferences within the BMVA's remit. The maximum amount of a bursary is £750. In return for the bursary, students are asked to write a conference report for BMVA News – most of the conference reports you read in these august pages are from students who have received bursaries – or do some work for the BMVA website. Details on eligibility and the application procedure are outlined on the BMVA website.

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

Announcements: *The Queen's Birthday Honours List 2013*

Readers' attention is drawn to two outstanding people from our field who have received honours in the June 2013 Queen's Birthday Honours list.

The first is **Sir Nigel Shadbolt**, who has received an ordinary knighthood¹² for services to science and engineering. During his 30-year career he has worked on computer and web science, psychology, cognitive Science, computational neuroscience and AI, and has strongly contributed to the emerging field of Web Science. Thus his work is highly relevant to machine vision.

The second is **Professor Alison Noble**, who has received an OBE for services to science and engineering. Professor Noble is Professor of biomedical engineering at the University of Oxford, and a Fellow of St. Hilda's College Oxford, and has worked on biomedical image analysis, in particular on the application of machine learning to medical imaging, including cardiology, obstetrics and perinatal care.

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

¹² To my mind there is nothing ordinary about an 'ordinary knighthood': it apparently means that in chivalry terms, it is not quite on a par with the Order of the Bath or the Order of the Garter. However, if anyone wishes to provide further clarification, I would be very happy to include it in a future issue of BMVA News. – Ed.

BMVC 2013 – Call for Participation



The 24th British Machine Vision Conference (BMVC 2013) will be held at the University of Bristol (<http://www.bristol.ac.uk/>) during 9–13 September 2013. The host City of Bristol is located in the South West of the UK, two hours from London.

BMVC is one of the major international conferences on computer vision and related areas. It is organised by the BMVA. It is a single-track meeting with oral and poster presentations. This year the conference received 439 submissions. Early Registration is now open and the early bird rate applies until 20 July 2013. Please visit the conference website for details:

<http://bmvc2013.bristol.ac.uk/registration>

BMVC 2013 will feature keynotes by Andrew Zisserman (Oxford) and Frank Dellaert (Georgia Tech), as well as tutorials by Adam Coates (Stanford) and Andreas Krause (ETH).

Conference reception and banquet

The conference reception will be in the Great Hall of the University Tower (MWB) and will feature a presentation by Aardman's Head of Technical Team Tom Barnes on the latest visual animation projects made in Bristol. The traditional banquet will be held at Bristol's Great Council Hall next to the city's 12th Century cathedral. See you at BMVC 2013!

Important dates

Author Notification of Acceptance:	1 July 2013
Camera-ready Paper & Author Registration due:	25 July 2013
Conference Start:	9 Sept 2013

Conference Chairs

Tilo Burghardt, Walterio Mayol-Cuevas and Majid Mirmehdi (Bristol University).

Information

For more information please visit:
<http://bmvc2013.bristol.ac.uk>

For enquiries please contact:
bmvc2013@cs.bris.ac.uk

Fifth International Conference on Imaging for Crime Detection and Prevention (ICDP 2013)

This conference will take place in Kingston University, London on 16–17 December 2013. It is organised under the auspices of IET's Vision and Imaging Network, with sponsorship by the IEEE's Social Implications of Technology Society and the BMVA. It follows on from the successful IDSS (Intelligent Distributed Surveillance Systems) events held in 2003 and 2004 and ICDP 2005, 2006, 2009 and 2011 to bring together researchers, industry, end-users, law-enforcing agencies and citizens groups to share experiences and explore areas where additional research and development are needed, identify possible collaboration and consider the societal impact of such technologies.

Full papers (up to 6 pages) are invited on all aspects of Imaging Surveillance technologies, from academia, industry, NGOs and others, to be selected for oral presentations or posters through a peer-review system. Topics include:

- Biometrics
- Cognitive Systems Engineering
- Gesture and posture analysis and recognition
- Human Machine Interfaces
- Information fusion
- Learning systems
- Metadata generation, video database indexing
- Multi-camera systems

- Robust computer vision algorithms
- Social implications of surveillance technologies
- Surveillance Systems and solutions
- Wireless communications and networks for video surveillance, video coding, compression.

Accepted papers will be published on the IET's Digital Library and indexed by Inspec, provided that at least one author registers and presents/displays the work. Authors of exceptional papers will be invited to submit extended versions to be considered for publication in the IET Computer Vision Journal. There will be delegate fee discounts for authors, students and members of the IET and sponsoring organisations.

Key dates

Paper submission:	1 September
Notification of acceptance:	25 October
Camera-ready:	11 November

For further information please see the conference web site (<http://www.icdp-conf.org>) or contact Sergio Velastin.

Professor Sergio A. Velastin
University of Santiago de Chile, Chile
email: sergio.velastin@iee.org

RTIVP 2014 Programme Committee

- Mohamed Akil, École Supérieure d'Ingénieurs en Electronique et Electrotechnique (France)
- Ahmed Bouridane, Northumbria University, Newcastle (UK)
- Chang Y. Choo, San José State Univ. (USA)
- Roy Davies, University of London (UK)
- Eran Edirisinghe, Loughborough University, Loughborough (UK)
- Barak Fishbain, Israel Institute of Technology (Israel)
- Johannes Fürtler, Kapsch TrafficCom AG (Austria)
- Mark Gamadia, Apple Inc. (USA)
- Sergio R. Goma, Qualcomm Inc. (USA)
- Christos Grecos, Univ. of the West of Scotland (UK)
- Reinhard Koch, Christian-Albrechts-University of Kiel (Germany)
- Mehrube Mehrübeoglu, Texas A&M Univ. Corpus Christi (USA)
- Antonio Nunez, University of Las Palmas, Gran Canaria (Spain)
- Amos Omondi, Zed Consultants, Toronto, Canada
- Antonio Plaza, Univ. of Extremadura (Spain)
- J.A. Madeiras Pereira, INESC-ID (Portugal)
- Volodymyr Ponomaryov, Instituto Politécnico Nacional (Mexico)

- Luis Salgado, Univ. Politécnica de Madrid (Spain)
- Jorge Santos, European Commission, Brussels (Belgium)
- Sergio Saponara, Univ. di Pisa (Italy)
- Mukul V. Shirvaikar, The Univ. of Texas at Tyler (USA)
- Athanassios Skodras, University of Patras, Patras (Greece)
- Stephen Stilkerich, EADS (Germany)
- Pedro Santos, Fraunhofer IGD (Germany)
- Benny Thörnberg, Mid-Sweden University, Sundsvall (Sweden)
- Thor Vollset, Tordivel S.A., Oslo (Norway)

Editorial

(continued from p.2)

So I am having to hedge about to identify evolutionary dead-ends, fossils, missing links, and so on. This brings to mind what happens in scale-space representations, as one goes down in scale: some lines split, and new ones form (or conversely on going up in scale, some lines disappear and others merge). In the case of merging, any description of the merged form will probably have to be altered. Thus it is that evolution will cause dead-ends and transmogrifications, with the result that for PR and computer vision, methods will become increasingly hybrid and unlike their earlier forms, and may well appear in some cases like entirely new approaches. Yet when examining PhD theses or prospective new publications, one often wonders whether there is anything really new there: it all sounds too familiar. But even with this familiarity, one can't necessarily discern the antecedents; and where one can, how much credit can one give for the novelty residing in a new recognition structure when it's the way it's put together that is new, not the components. Chemistry is putting elements together in new ways to form new compounds. And computer vision is now largely about putting MatLab algorithms together to make even more impressive algorithms. Which implies that we now have all the components necessary to achieve this. What a wonderful new world we live in wherein all potentially useful components and all possible image data reside on the Internet, so there is no longer any need to get one's hands dirty: rather, we only have to be inventive enough to try some of the probable 10^{10} combinations to earn our PhDs, CBEs or Nobel Prizes. A pity that we are now excluded from using new pictures from our own cameras – which is yet another sign of the times. But I'm not throwing my camera away just yet: the pendulum could be swinging too far and the fashion could change ...

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