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 1 December 2009.

Contents

Editorial: *Emulating the Human*

Amongst my earliest recollections when starting work on vision in the 1970s were applications such as remote sensing, fingerprint recognition, and a variety of instances of cytology and industrial inspection. While stereo imaging was of general interest, because of its obvious relevance to the human visual system, this and many other aspects of practical vision systems were hugely limited by the power of the computers then available, and the mood of the day was more on how to achieve the most by whatever shortcut methods could be found. In industrial inspection this meant that there was undue emphasis on thresholding techniques and binary processing, and indeed in many areas thinning algorithms were in vogue and new ones were frequently reported. While one would have liked to spend more time on the higher level aspects of vision, I for one found myself continually pushed back to the lowest level possible, as for example accuracy lost is never regained by further processing. Similarly, robustness has to be built in and is difficult to graft it on later. In those days my algorithm priority list already included detection sensitivity, accuracy, robustness, adaptability and speed of operation, and these priorities have hardly gone away in the subsequent decades. Of course, the one major change is that computers have become hugely more powerful, both in speed and storage, but modern applications involve correspondingly large images, video sequences and image databases, so even now we have by no means all the power we need to deal with these. Somehow we got by earlier ('cutting our cloth according to our means'), but we are still unsatisfied ('work expands to fill the time available' – which reflects human psychology rather than necessity).

Having just attended a seminar given by Dame Wendy Hall, in which she expounded on the Internet and its

two-way communication possibilities, I now realise that my mental model of the situation was inadequate. For there we have not just two-way communication but three-way links. Thus we have not just the sender and the sendee, but also the observer, the 'big brother' who will see all the traffic going or coming our way and will anticipate all our needs and supply not only them but more than we *think* we want. Am I just talking about Google – or about a future Google or Google look-alike? Anyway, for such a monster to operate perfectly it must *understand* all the data going to and fro, and from an imaging point of view this presages much more work for all of us, acting as technological enablers. To the best of my knowledge we are still in the infancy of image understanding. In fact, until robots can walk and talk and acquire information and mechanistic knowledge about the real world in the same way that babies to, I don't feel we will have cracked that one.

Naturally, we have come a long way since the 1970s – not least in production of 3D and motion interpretation techniques, which BMVA has been at the heart of. Somewhat oddly, however, we seem to have taken some steps back: look at the number of total rethinks of feature detection that have taken place in just the last few years – SIFT, SURF, GLOH, MSER, …; and think of the new approaches to noise suppression that have arrived in recent years. Such aspects are reflected in the prize papers at BMVC, not just by more run-of-the-mill work. And then there is the work on archive film defect detection, detection of forgeries in paintings, and many other modern applications that could not have been tackled seriously in earlier decades. Now we have even started to think the unthinkable – analysis of esoteric entities such as beauty and symmetry. Of course, some would say that symmetry is a mathematical concept highly amenable to traditional image analysis – and there are interesting papers that consider this carefully and rigorously (though this topic is not a 'done deal' yet). On the other hand, in nature many objects such as leaves are almost symmetrical, and to the human eye the lack of perfect symmetry can itself make for even greater levels of beauty (handmade pottery is more expensive than machine made, and a face that has a slight bilateral asymmetry can have an intriguing, enticing quality). Which moves us on to John Robinson's 'last stand', on the Fourth Plinth (see page 15). Actually, I'm sure it isn't his last stand: knowing John, giving beauty marks out of 100 cannot be his last word! But it does presage important further work ahead for all of us. Dare I suggest that as beauty is multi-faceted, it will be highly amenable to PCA?!

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

Announcements at BMVC

Results of the BMVA elections:

This year there were 6 nominations for 5 places. The following were elected to the BMVA Executive Committee:

- Professor E. R. Davies
- Dr. M. Everingham
- Dr. A. Fitzgibbon
- Dr. D. Makris
- Professor M. Mirmehdi

The main change is that M. Everingham has replaced P. Hall on the Committee.

Location of BMVC 2010:

The next BMVC will be held in Aberystwyth and will be organised by R. Zwiggelaar and his colleagues.

Dr. Neil Thacker
BMVA Secretary
email: neil.thacker@manchester.ac.uk

BMVA News Conundrum

The conundrum that appeared in the last issue of BMVA News asked how it is that space can be saved in a document by using double column instead of single column format. The solution was provided by Chris Trayner:

"It's because the saving in wasted space at paragraph ends more than compensates for the space lost in the inter-column gap."

Chris followed this with another conundrum for readers:

With journal papers, which are many pages long, you would expect that when typeset about half would come to an odd number of pages and half to an even. Thus about 50% should start on a left-hand page and 50% on a right-hand page. Why is the prediction so radically wrong for professional journals? (The solution appears on page 20.)

Chris Trayner
University of Leeds
email: c.trayner@leeds.ac.uk

Around and About at BMVC 2009

Rama Chellappa gives his invited lecture.

View of lecture theatre, with Rama Chellappa on the near right.

Alexei Efros launches into his invited talk.

Overview of the scene at Alexei Efros's invited talk.
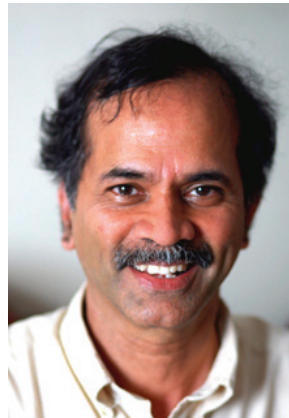
Close-ups of invited speakers Rama Chellappa (University of Maryland) and Alexei Efros (Carnegie-Mellon University).

Umar Mohammed (UCL) giving his paper "Visio-lization – generating novel facial images".

The conference banquet in full swing.

High table diners – left to right: Rama Chellappa (invited speaker), Andrea Cavallaro, Bernard Buxton, Phil Torr, Andrew Zisserman, Lourdes Agapito, Alexei Efros (invited speaker) and Andrew Fitzgibbon.

High quality light entertainment!

High-up view of the elegant octagon dining hall.

The octagon itself.

Ben Benfold receives his prize from Daniel Alexander.

Paul McIlroy receives his prize from Daniel Alexander.

Xiaosong Wang receives his prize from Daniel Alexander.

Andrew Zisserman makes his Distinguished Fellow address.
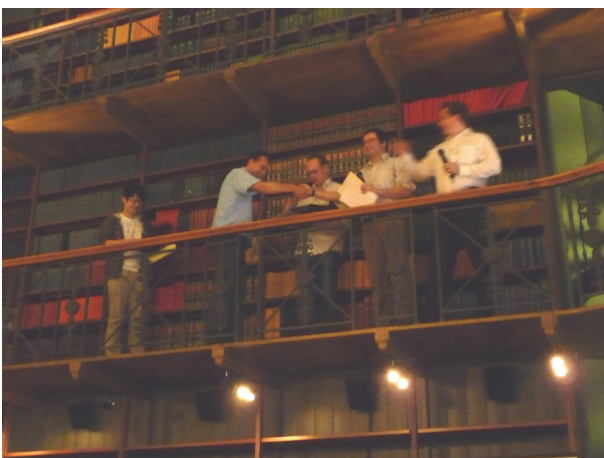
Francisco Estrada receives his prize from Daniel Alexander.

Another view giving some idea of the sense of occasion.

Andrew Zisserman about to receive his Distinguished Fellow award from Andrew Fitzgibbon.

Bernard Buxton receiving his Distinguished Fellow award from Andrew Fitzgibbon.

Bernard Buxton making his Distinguished Fellow presentation.
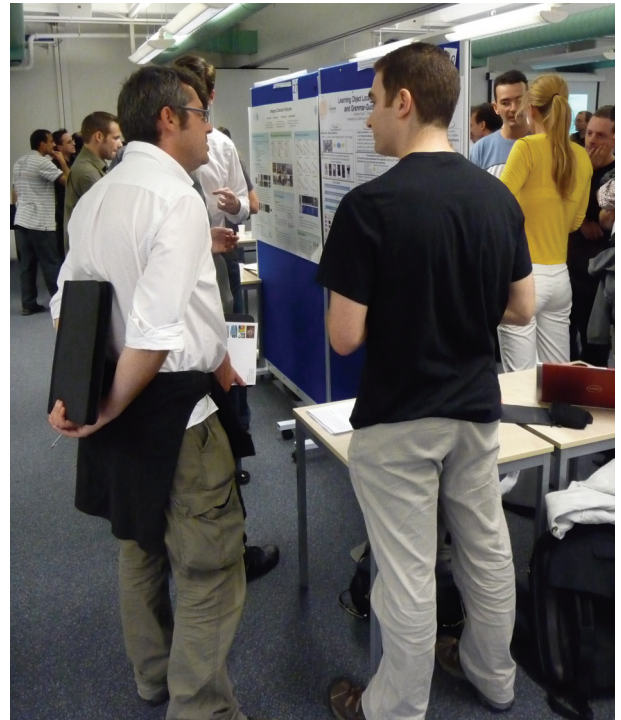
Richard Bowden in discussions about a poster.

Alastair Moore discusses his poster with Toby Breckon.

Francisco Estrada (University of Toronto), Andrew Zisserman and Simon Prince enjoying a quiet drink.
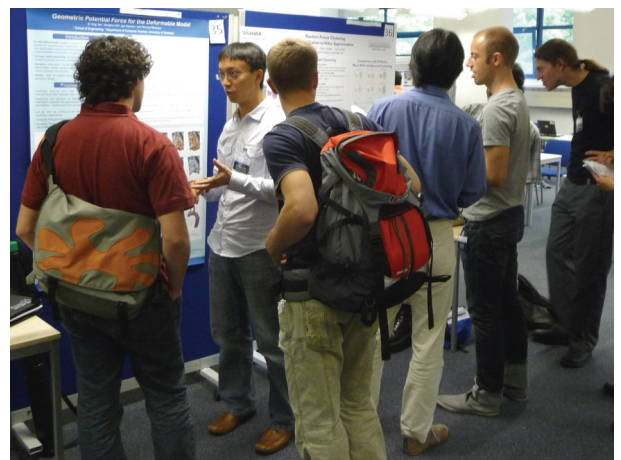
Xianghua Xie explains the details of his work to Emilio Maggio (extreme left).
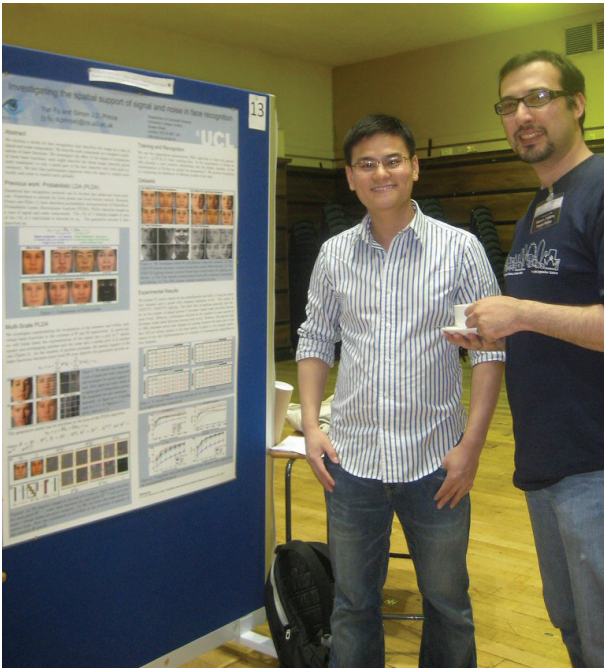
Yun Fu and Ioannis Tziakos pose next to Yun's poster.
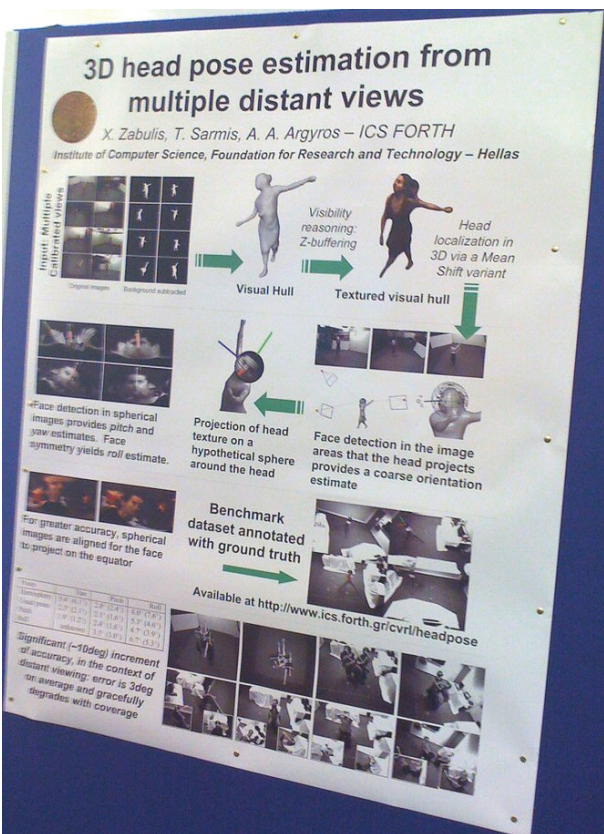
This elegant poster speaks for itself!

Views of the elegant conference banner designed by Andrea Cavallaro, which doubled as a sundial! (Credit must also be given to Alastair Moore for the conference logo, which appears near the top right of the banner.)
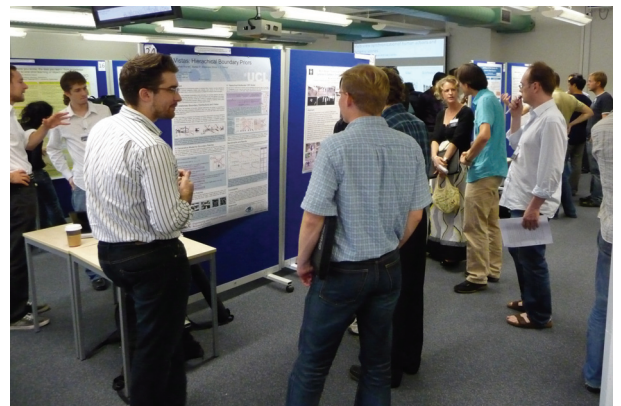
Finally, I would like to acknowledge Jania Aghajanian and Riccardo Mazzon who acted as highly productive photographers – and who made it possible to bring together all these views of the life of the conference.

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

Summary of Prizes Presented at BMVC 2009

Best Paper Award:

- Francisco Estrada, David Fleet and Allan Jepson – Stochastic image denoising.

CRS Best Industrial Paper Award:

- Xiaosong Wang and Majid Mirmehdi – HMM based archive film defect detection with spatial and temporal constraints.

Best Student Poster Awards[2]

- Ben Benfold and Ian Reid – Guiding visual surveillance by tracking human attention.
- Paul McIlroy and Tom Drummond – Reconstruction from uncalibrated affine silhouettes.

Dr Simon J.D. Prince
University College London
email: s.prince@cs.ucl.ac.uk

---

[2]Sponsored by Elsevier, EPSRC and Springer

Why a bounding box is not enough if you want to shoot a giraffe!

You might be wondering what this article is about! Shooting giraffes, what does this have to do with computer vision?

That was the exact same question that went through my mind during the IEEE Computer Vision and Pattern Recognition (CVPR) conference this July. It was part of a talk about simultaneous detection and segmentation when the speaker pointed out that, I quote, "a bounding box is not enough if you want to shoot a giraffe". I would be surprised if, after this comment, there was anybody in the audience without a little smile on their face. It took a few days before I found out that this quote actually originates from David Forsyth of UIUC.

CVPR 2009 was held in Miami on 20–25 June. A typical day started at 9 am with two parallel sessions of four oral presentations each, followed by a poster session before lunch. Then the afternoon saw another session of talks, followed by a second poster session, and concluded with the final four talks of the day.

Break after a session of oral presentations. The organisers made sure there was enough coffee, cookies etc. for everyone.

Two days before the start of the conference, and another at the end, were entirely dedicated to workshops and short courses. I especially found the short courses to be very interesting. They covered topics such as Kernel Methods in Computer Vision and a full-day course on Computer Vision on GPUs.

In my opinion, the most important session was on the second day when the four authors who won awards presented their work. The Best Paper award went to Kaiming He, Jian Sun and Xiaoou Tang, and was entitled "Single image haze removal using dark channel prior". In this presentation the authors showed how haze in outdoor images can be removed. Haze is an additive artefact which depends on the distance of an object to the camera: the further away, the stronger the effect of haze. The main observation is that most local patches in haze-free outdoor images contain some pixels which have very low intensities in at least one colour channel. The authors use this observation to compute the so-called "dark channel prior" which indicates how much haze is in between an object and the camera. This dark channel prior is then used to remove the haze from the original image. In addition, the dark channel prior serves as an estimate of the depth map of the image.

Poster-session on the second day of the conference. This was one of three rooms where posters were presented.

Overall, I believe this work was a good choice for the best paper award, with the elegance of converting a simple yet powerful observation into state-of-the-art results for the problem of haze removal.

One of the hot topics at this year's CVPR was discriminative approaches for structured output learning. An example for such a structured output space is a graph which, given an input image, assigns labels to the individual patches such as "sky", "road", "tree", etc. This is in contrast to traditional classification approaches where the outputs are comparably simple, for example binary labels in the case of SVMs.

A popular way of learning these regressors involves Structured Support Vector Machines. I noticed at least five papers during the conference which use structured output spaces to learn a similarity measure between shapes which minimises classification loss (Chen et. al.); to learn structured latent variables to recognise human actions (Wang and Mori); and to incorporate a global connectivity prior which encourages objects to be segmented into only one part (Nowozin and Lampert).

Also very interesting was the work of Ce Liu, Jenny Yuen and Antonio Torralba on scene labelling. Given an unseen outdoor image, the task is to assign one of 33 labels to every pixel, e.g. "window", "tree", "sky" etc.

Traditional approaches perform this by learning a model for each class and then computing the label assignment with highest score. In this work though, the authors compare the given input image with a large database of annotated images using SIFT flow and compute the pixel labels for the given input image from these comparisons.

Another approach worth mentioning is that of Abhinav Gupta, Praveen Srinivasan, Jianbo Shi and Larry Davis. The authors show how activity recognition can be improved if the temporal sequence of events is known. This is true for baseball games broadcast on TV where a sequence of events starts with the pitcher throwing the ball, followed by the batter trying to hit the ball; if he succeeds then we either see the fielder running towards the ball, or the batter running towards a base. With this kind of prior knowledge, the authors show how activity recognition can be improved while simultaneously learning the storyline model.

For me, the most important event occurred on the last day: it was the last talk of the last session. This was when my talk was scheduled about automatic learning of British Sign Language using signing footage and subtitles taken from TV. It was not the best timing to have the last talk of the conference, since this meant that I had to rehearse when others went swimming in the Atlantic at the end of each day.

The CVPR 2009 conference hotel with 1504 rooms, one of the biggest hotels in Miami.

In case you are wondering, yes, there was a beach, right behind our hotel. Generally, the hotel was a treat for the eyes. However, some of the outside facilities were not exactly in line with PhD student salaries. For example, there were lovely pavilions with sofas and TVs. Our hopes to enjoy one of them were soon shattered though. How naïve we were, thinking that the outside facilities would be included in the room rate. Apparently, renting one of the pavilions for a day would have set us back 250–500 US dollars, depending on the location!

More information about the IEEE Computer Vision and Pattern Recognition (CVPR) Conference 2009 can be found on the conference webpage: www.cvpr2009.org.

I would like to thank the BMVA for a generous travel bursary which made it possible for me to attend this conference.

Patrick Buehler
University of Oxford
email: patrick@robots.ox.ac.uk

Report on CVPR 2009

The IEEE Conference of Computer Vision and Pattern Recognition (CVPR) 2009 took place in Miami, Florida on 20–25 June 2009. The Fontainebleau Hotel with prime beachfront location was chosen as the conference venue. This prestigious annual computer vision conference received 1464 submissions, of which 383 papers were accepted (61 for oral and 322 for poster presentation). More than 1000 participants from academia and industry participated in the conference.

The main conference (22–24 June) included a total of 16 oral sessions (dual-track) in the course of three days. Six poster sessions, together with various demonstrations and exhibitions were held in between the oral sessions, where individual and group discussions came to life. The main program this year also allocated a session called "doctoral spotlight" for students near graduation to showcase their work to the community and potential employers. In addition to the main program, a number of workshops and tutorials were held on the days adjacent to the main conference, covering topics such as tracking, kernel methods, biometrics, surveillance, video search engines, GPUs, etc.

The main conference commenced on Monday morning (22 June) with two oral sessions. I attended the one on "Image and video search". Some interesting talks in this session "Pose Search: retrieving people using their pose" and "Bundling features for large scale partial-duplicate web image search". The first presentation demonstrated a query-based pose retrieval system on unconstrained movies. The second talk introduced a new feature representation which improved on bag of words representation by exploiting geometric relationships among visual words. After the lunch break, I attended Babenko's talk on a multiple instance learning-based appearance model for robust real-time tracking. I was also able to catch up the talk given by Saffari on multi-class semi-supervised boosting method.

The long plenteous day ended with a reception in the Bass Museum, where all the delegates shared a good time under the artistic atmosphere.

The second day of the main conference consisted of four oral sessions. I found some of the talks very fascinating, such as the idea on using bag of regions to encode shape and scale information for object detection and recognition. In the afternoon, the conference gave the Best Paper award to the authors of "Single image haze removal using dark channel prior". The best student paper award went to Ce Liu with his paper entitled "Nonparametric scene parsing: label transfer via dense scene alignment". The conference also awarded two papers from CVPR 1999 the Longuet-Higgins Prize for Fundamental Contributions in Computer Vision. The papers are "Statistics of natural images and models" and "Adaptive background mixture models for real-time tracking". The PAMI TC Meeting was held after a wonderful banquet dinner. It was announced that CVPR will be held in San Francisco, California next year, whilst Colorado Springs and Rhode Island were selected as the sites for CVPR 2011 and 2012 respectively.

The final day of the main conference kicked off with oral session "Video analysis" during which I presented my paper entitled "Multi-camera activity correlation analysis". After the lunch break, the day continued with a number of interesting talks presented by leading experts from computer vision and pattern recognition. Some interesting talks include the presentation on a method for multi-object tracking in crowded scenes given by Yuan Li from Nevatia's group, University of Southern California.

The focus of the conference this year included object detection, object recognition, motion, tracking, segmentation, face, gesture, video analysis and activity understanding. This research trend clearly reflects the increasing importance of image and video understanding due to growing needs in commercial and the consumer markets. For the domain of visual surveillance, research has evolved from analysing single camera view with not-too-crowded scenarios to more complex scenarios with multiple cameras featuring crowded activities. I found some new datasets released in the conference may be useful for benchmarking purposes. One of them is Caltech Pedestrian Dataset for pedestrian detection, which is available at http://www.vision.caltech.edu/Image_Datasets/CaltechP edestrians/. There is also a new action recognition dataset named as "Hollywood 2", which can be obtained from www.irisa.fr/vista/actions/hollywood2. Besides that, a large-scale image database has been released under www.image-net.org.

CVPR 2009 was a great success and the participants had certainly enjoyed every part of it including the great summer weather in the Sunshine state. I would like to express my gratitude to the BMVA and the School of EECS, QMUL for funding the trip.

Loy Chen Change
Queen Mary University of London
email: ccloy@dcs.qmul.ac.uk

Report on CARS 2009

The 23[rd] Computer Assisted Radiology and Surgery Conference (CARS) was held in Berlin on 23–27 July 2009. CARS is an interdisciplinary conference, bringing together people from academia, the medical community as well as the healthcare industry. Its main focus is research and development for computer assisted systems and their applications in radiology and surgery. Topics covered include image guided interventions, medical imaging, molecular imaging, image processing and display, Computer Aided Diagnosis (CAD), surgical simulation, surgical navigation and robotics, as well as PACS applications.

The conference program comprised a mixture of plenaries, orals, posters and tutorials spread over five days. CARS started with the usual welcome from the organisers, followed by four highly enjoyable lectures from keynote speakers. Michael Vannier from the University of Chicago talked about the transition from patient specific modelling to obtaining model-based evidence in clinical applications. Bill Buxton from Microsoft Research offered us a peek in the future of multimedia and human computer interaction. He also envisaged use of Microsoft Surface technology in an application of medical imaging, diagnosis and consultation. Next, Liz Beckman from Lanmark Medical and Stephen Golding from Oxford University, gave a joint lecture demonstrating how Darwin's theory of evolution applies in computer assisted radiology and therapy. Finally, Michael Vannier once again got onto the podium and closed the conference opening session with a lecture on Computed Tomography (CT) imaging of Meresamun, a 3000 years old mummy which has never been unwrapped.

There were some interesting applications of machine vision to medical images. Matsuda et al. used GPU-based acceleration of image filtering to yield a speed-up of ~50 times over conventional CPU architecture that was suitable for real time interactive display. Another GPU-based implementation by Kelly et al. involved Voronoi-based probability maps for region

segmentation in abdominal CT images. A patient specific atlas model was applied to the pancreas in 3D CT image by Shimizu et al. Their algorithm uses Maximum A Posteriori (MAP) probability to label voxels as pancreatic tissue based on their shape atlas which is warped to the image. Bashar et al. designed a method of stitching together several endoscope images using Fourier phase correlation in the matching process.

Moreover, various clinical applications of image processing and analysis techniques were presented in the 11th International Workshop on CAD, which formed part of the CARS conference. The breast CAD session started with two invited lecturers, Gillian Newstead from Chicago who told as about the current status and future potential for CAD in breast MRI, and Toula Destounis from Rochester who presented a retrospective study showing the ability of CAD to detect cancer on prior mammograms. Next, Rangaraj Rangayyan from Calgary presented results on the detection of architectural distortion in prior mammograms by combining Gabor filters and texture analysis. A 3D fuzzy technique for tumour detection from whole breast ultrasound was presented by Ruey Chang from Taiwan. In a special interactive session, a commercial CAD system was put head to head with expert radiologists in an attempt to classify mammograms. The CAD system provided useful help to the human expert, and even pointed out cancerous cases that were initially overlooked by the radiologists. Ground truth results were provided by biopsies.

In the Thoracic CAD session, invited lecturer Kensaku Mori presented an automated analysis method for lung nodule follow-up based on non-rigid registration by mutual information minimisation. Moreover, in the CAD for MRI session, Vipin Chaudhary from Buffalo presented a method for the detection of abnormal intervertebral discs based on a Gibbs model of disc position and appearance. The workshop closed with a highly interesting live demonstration of CAD research prototypes.

We represented UCL with two oral presentations involving machine vision to aid in clinical research and diagnosis. Irving Dindoyal worked on segmentation of 3D foetal cardiac data using multi-region interacting level sets to reconstruct missing intra-cavity boundaries. The algorithm was validated using a physical phantom. Sofia Michopoulou presented a CAD system for characterising intervertebral disc degeneration severity from MR images of the lumbar spine, utilising atlas-based segmentation and texture analysis.

All conference papers were published in the International Journal of Computer Assisted Radiology and Surgery (Vol. 4, No. 4, 2009) and can be found online at www.springerlink.com.

An industrial fair was present at the conference to show off the latest medical technology. These included products ranging from electronic whole body vibration massagers, rapid prototyping model building to software for aiding medical visualisation and physical measurement of features within the image. Due to the expensive hardware on display, security was very tight, with smart suit bouncers on the lookout for suspicious people without name badges. The products on display complemented the academic content of the conference talks and posters.

The organisers put together a well structured scientific program by running up to four key research areas as parallel oral sessions. Sessions were lively, with varied questions following each presentation and strong participation in panel discussions. The posters were grouped according to subject and were always visible throughout the conference. Concluding, this conference not only presented the latest research in computer assisted radiology and surgery, but also promoted the exchange of ideas between clinical experts, engineers and physicists, while offering an insight into the application of research results in clinical practice. CARS 2010 will be held in Geneva, next June. We would like to thank the BMVA for sponsoring our participation in this outstanding event.

Irving Dindoyal and Sofia Michopoulou
University College London
emails: dindoyal@medphys.ucl.ac.uk
        s.michopoulou@ucl.ac.uk

Report on ISBI 2009

The Sixth IEEE International Symposium on Biomedical Imaging (ISBI'09) was held on 28 June–1 July at the Park Plaza Hotel in Boston, Massachusetts, USA. Just like last year's edition in Paris, the conference attracted a highly international group of participants. This year, 527 contributed papers were submitted for review, of which 338 were accepted. Given the large number of submissions, the competition was strong and the quality of the papers presented high. Each day, the conference started off with a poster session in the morning, followed by a plenary talk and multitrack oral presentations. The contributed papers covered topics as diverse as computer aided detection, image segmentation, reconstruction, and registration, super-resolution image processing, electron-microscopy, image denoising, strain imaging, etc.

I was given the opportunity to present my work on the incorporation of anatomical priors into Bayesian tomographic reconstruction during an oral session on the third day. The title of my paper was "Robust incorporation of anatomical priors into limited view tomography using multiple cluster modelling of the joint histogram". I very much enjoyed giving the talk, as I received useful questions and feedback at the end. After the talk I also made new contacts with researchers at University College London (UCL) and the University of Leuven in Belgium, who are working on similar topics. While their work focused on diffuse optical tomography and micro CT, they shared my interest in the use of information theoretic similarity measures to regularise poorly conditioned reconstruction problems.

The plenary talks at the conference were highly inspirational. On the first day, Tom Miller, CEO of the Workflow and Solutions division at Siemens, gave the opening talk entitled "Potential and possibilities: from the very critical to the very difficult". He delivered a fascinating account of the history of medical imaging, and sketched a picture of the exciting developments that lie ahead in the field of medical imaging and device design, most notably their role in enabling personalised medicine. He also commented on the impact of the current economic crisis on medical research, and how it is forcing companies like Siemens to assign their reduced research budgets much more critically.

The second day saw Professor Ralp Weissleder, Director of the Center for Molecular Imaging and Center for Systems Biology at the Harvard Medical School, give a plenary talk on recent advances in molecular imaging. He reported on a number of exciting advances obtained in his lab, such as the development of novel magnetic nanoparticles for MR imaging, enzyme activatable probes for minimally invasive laparoscopy, and clinical testing of long circulating polymers for angiogenesis imaging.

On the third day, Professor Xiaowei Zhuang, Professor of Physics, Chemistry and Chemical Biology at Harvard University, gave a talk on nanoscopic imaging of biomolecules, cells and tissues. Using sparse fluorescence techniques, Professor Zhuang managed to break the diffraction limit on the resolution of optical microscopy, which is classically limited to a few hundred nanometers, and applied the resulting super-resolution technology to the imaging of cell biology and neurobiology. The technique is called stochastic optical reconstruction microscopy (STORM), and uses photo-switchable fluorescent probes to temporally separate the otherwise spatially overlapping images of individual molecules, allowing the construction of high-resolution images. During the talk, Professor Zhuang showed spectacular three-dimensional, multicolour fluorescence images of molecular complexes, cells, and tissues with ~20 nm lateral and ~50 nm axial resolutions.

Congratulations go to the organising committee of the conference, and all those who helped with tasks ranging from the reviewing of papers to the provision of logistic support for the conference. The Park Plaza Hotel was a terrific location to hold the conference, as it provided superb facilities and lecture rooms, as well as immediate access to the centre of Boston for those wishing to go sightseeing after the conference. Particularly noteworthy were the reception and dinner held on the second day of the conference in the Prudential Tower Skywalk Observatory, from where one can usually enjoy sweeping 360 degree views of Greater Boston and beyond. Unfortunately, however, the view that day was impeded by low-hanging clouds which shrouded the top of the tower in a thick grey mist. Nonetheless, the food and atmosphere were great, providing a wonderful opportunity to relax with friends and colleagues after a day's worth of poster and oral presentations.

The 7th ISBI will be held in Rotterdam, The Netherlands in 2010. The deadline for the initial submission is 2 November 2009, and the notification date is 15 January 2010. Further details can be found at:

http://www.biomedicalimaging.org/

Finally, my thanks go the BMVA for supporting my attendance at the conference with a travel bursary.

Dominique Van de Sompel
University of Oxford
email: dominique@robots.ox.ac.uk

Report on OHBM 2009

The 15th Annual Meeting of the Organization for Human Brain Mapping (OHBM) was held in San Francisco, USA this year on 18–23 July. The conference is the main international forum for all those involved in brain mapping to meet, with the multidisciplinary nature of the field on display in the form of neuroscientists, psychiatrists, computer scientists, physicists and psychologists to name but a few – the common goal being the better understanding of both the function and structure of the human brain.

This year more than 2000 posters were presented over the course of four days, including my own, entitled "Using iterative wavelet cluster analysis to produce individualised reference regions for characterisation of drift in pharmacological MRI", which was a practical

application of a method I first presented at MIUA 2008 in Dundee. At first glance a conference of this size can seem a bit daunting, but good organisation and a well laid out program made it all manageable, with the posters split into 12 categories and around 70 sub-categories. Indeed, the conference website allowed one to create a personalised timetable of the posters and sessions they wanted to attend. The workshops and oral sessions were multi-track, usually having at least three occurring in parallel, with everyone meeting up again for the symposia and keynote talks.
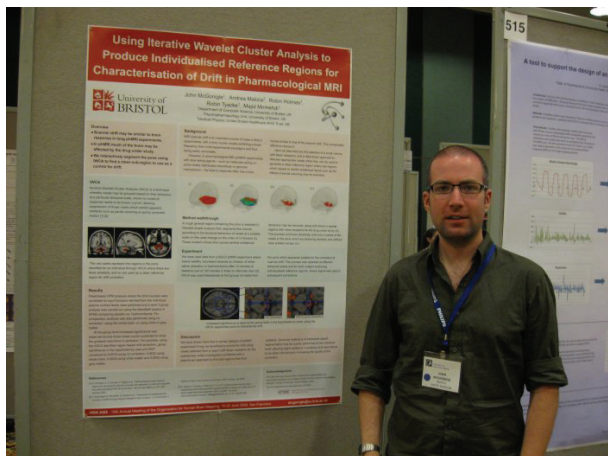
Kicking things off bright and early on the Friday morning was a statistics workshop entitled "To be Bayesian or frequentist or not". An interesting debate took place, the pros and cons and the myths of each approach being discussed with regard their application in functional brain imaging. Sadly, Keith Worsley, who was the world's premier statistician and mathematician in imaging neuroscience, died earlier this year. Always a popular figure at OHBM, he was sorely missed.

Due to the breadth of expertise attending, the keynote talks were kept accessible and enjoyable. The most prestigious of these is the Talairach Lecture, which this year was given by Patricia Churchland on "Morality and the social brain". Of particular interest to me were Tom Mitchell's keynote "Machine learning and brain imaging" and Kamil Ugurbil's "Frontiers in brain imaging at ultrahigh magnetic fields". There were several oral sessions on Modelling and Analysis, with those on Computational Neuroanatomy and Machine Learning featuring talks with a computer vision theme. The two that stood out most for me were "Particle swarm voxel clustering for multivariate fMRI mapping" by Malin Björnsdotter Åberg and "Modeling the relationship between cortical geometry and cytoarchitectonics via image registration" by Thomas Yeo.

A popular topic this year was the brain's default network – what's going on when we're not trying to do anything – with several hundred posters and a workshop dedicated to it. Also, and in common with recent conferences in several different fields, small-world networks made an appearance with Paul Laurienti chairing the symposium "The brain as a small-world network: from micro- to macro-scale". Ed Bullmore gave a memorable talk in this on "Fractal scaling of small-world functional networks", exploring some of the scale-free network organisation in functional brain connectivity.

A favourite of the attendees was a poster on brain activity in a dead north Atlantic salmon (as a humorous lesson on why one should use multiple comparisons correction in brain imaging). Following the conference proper I attended the Advanced fMRI course which focused on some of the more recent developments in analysis techniques, especially data driven methods, which is the area of my current research.

San Francisco was a great host city, with lots to see and do, including the famous cable cars and iconic Transamerica building. It was easy to tell how long attendees had been in the city to acclimatise beforehand by their level of sunburn – the fresh sea breeze making us forget to apply our sun cream on a walk across the Golden Gate Bridge! As one of America's more "European" cities the food was excellent, with plenty of good restaurants to be found. Next year's meeting will be closer to home in Barcelona, Spain and promises to be just as interesting in this fast moving field. The BMVA generously provided a large portion of the funding to enable me to attend this exciting conference, for which I am most grateful.

John McGonigle
University of Bristol
email: mcgonigle@cs.bris.ac.uk

Report on ICASSP 2009

The 34[th] IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP) took place in April in Taipei, Taiwan this year. It was held in the Tapei International Convention Center, near the imposing tower Taipei 101, which until this year has been officially the world's tallest building. This large annual event, the flagship conference of the IEEE Signal Processing Society, attracted more than 1700 attendees this year, with 1261 papers presented from an initial pool of 2703 submissions. The impressive conference programme included 60 lecture sessions, 72

poster sessions and a number of plenaries, talks, discussions, tutorials and other sessions, which left attendees spoilt for choice, so that it was only possible to experience a fraction of the riches on offer.

This year saw the trial of a new experiment by the conference organisers, by grouping presentations under four "Thematic Symposia", each a sort of conference within a conference. These were entitled "Signal Processing for 4G Wireless", "Network Distributed Signal Processing", "Immersive Communication" and "Multimedia Search and Retrieval". Each had its own overview talk and list of associated lecture and poster sessions. Amongst the plenary talks were "Signal processing bringing on market growth" by Kazuo Murano (Fujitsu Laboratories), and "The evolution and trends of the semiconductor" by Ming-Kai Tsai (Media Tek Inc.). These gave a useful and wide-ranging industrial perspective on the value of signal processing research.

Amongst the interesting tutorials at the beginning of the week were "Distributed video coding for low cost multimedia communications systems" by Anil Fernando (University of Surrey), and "Distributed processing in smart cameras" by Hamid Aghajan (Stanford University, USA) and Andrea Cavallaro (Queen Mary, University of London). Amongst the lecture and poster sessions were a wealth of presentations of interest to those in the fields of image processing and computer vision, including presentations on distributed video coding, video quality evaluation, image enhancement, stereovision, probability-based detection methods, trajectory clustering, object tracking and dimensionality reduction. Very effective examples of image inpainting were given in the presentation "Image inpainting via sparse representation" by Ben Shen et al. (Tsinghua University, China), showing strikingly realistic replacement of objects with automatically-generated background. A number of presentations were given on stereo matching. One paper, "A variational framework for simultaneous motion and disparity estimation in a sequence of stereo images" by Wided Miled et al. (Telecom Paris Tech, France) provided a satisfying marriage between the two essentially equivalent areas of stereo matching and motion estimation, which have for so long been treated as separate research areas. An immensely clear presentation on manifold learning was given by Shannon Hughes (Princeton University, USA), in her presentation "Connecting spectral and spring methods for manifold learning". One presentation with strongly suggestive hints of future potential was "A fully affine invariant image comparison method" by Guoshen Yu et al. (Ecole Polytechnique, France), presenting a promisingly-named "Affine SIFT" method for fully affine invariant image comparison.

Taipei itself is a large city of busy streets and low concrete apartment blocks with roof gardens, spreading across a wide plain surrounded by the beautiful mountains of north Taiwan. Taipei has a reputation for excellent Chinese food, with which the conference excelled itself, having perhaps as many different dishes on offer during the week as poster sessions. In the welcome reception, we were shown displays of regional culture in the form of local folk music and an impressive display of kung fu acrobatics. To accompany the conference, a number of local sightseeing tours were available, including coach trips to the nearby mountains and to the National Palace Museum, which houses the world's largest and finest collection of the possessions of the past emperors of China. Finally, a visit to Taipei would not have been complete without a trip to the top of Taipei 101, in which the lift (the world's fastest) ascends so fast that one's ears pop.

I am grateful to the BMVA for awarding me a grant towards visiting this conference, where it was a great honour for me to receive one of the Best Student Paper Awards for the paper "Distance blurring for space-variant video coding".

Timothy Popkin
Queen Mary University of London
email: t.popkin@elec.qmul.ac.uk

Report on BMVA/EPSRC Summer School 2009

The BMVA/EPSRC Summer School on Computer Vision was hosted by Kingston University, on 5–10 July 2009. It was attended by 45 students, mainly PhD students from British Universities.

Summer School attendants with Dimitrios Makris and Roy Davies at the end of the official dinner.

As part of the continuous effort of BMVA to maintain an up-to-date high quality programme, a lecture on "Shape and Appearance Models" by Tim Cootes was added and a lecture on "Multiview Geometry" by Lourdes Agapito was re-introduced.

First prize for best poster was awarded to Sam Johnson from the University of Leeds …

… while the second prize went to Spyridon Bakas from Kingston University …

… and the third prize to Josiah Wang (here between Dimitrios Makris and Neil Thacker) from the University of Leeds.

Sarah Taylor, PhD student from the University of Lancaster, said the following on behalf of the attendants: "The BMVA summer school is the best I have ever attended. Not only was it extremely informative and educational, it was very enjoyable. It was fantastic meeting so many people working in a similar area who could offer advice and support in my research. I made some amazing friends who I am still in touch with now. I fully recommend the BMVA summer school to all!"

Dr. Dimitrios Makris
Kingston University
email: d.makris@kingston.ac.uk

Engineering, Image and Art

A public lecture by John Robinson on the Fourth Plinth, 7 am, 4 August 2009.

35,000 people applied to participate in Antony Gormley's *One & Other*. With 2,400 slots to fill, that's about a 1 in 15 chance of selection. I expected longer odds so I applied and forgot about it. But someone has to win a lottery and three weeks before 4 August, my allotted day, the call came through that I was in.

One & Other puts members of the public on the fourth plinth in Trafalgar Square. Each person has an hour-long stint to do with as they wish. There's no choice about the day or hour (mine was 7–8 am), nor about the weather. The show goes on even in torrential rain. The event is video streamed and Sky Arts presents a weekly summary show.

I decided to give a talk. That wasn't my first thought, but I love to share my enthusiasms and I talk for a living, so I shelved dreams of extreme performance art and persuaded the University of York to bill my plinth stint as a public lecture. Because people would come and go, both in Trafalgar Square and online, I decided to present a string of topics loosely connected by the theme of optimisation rather than try to build ideas progressively as you would in an ordinary lecture. The title *Engineering, Image and Art* signalled my intention to capture the excitement of engineering in general, to emphasise my own area – image processing – and to link to art.

As you can see from the picture, it didn't rain, though I had to steady my visual's easel from time to time. I included three items on images and vision. A few minutes in, I talked about compression, and managed a digression on Lenna. Midway there was a slot about movie scene analysis, where I forgot to pay the tribute

to British computer vision research that I'd intended. Then towards the end of the lecture I returned to image analysis and presented some original research.

At York we have been working on face image description. We have multiple kernel probability models that we slice using conditional density estimation, then read off the values of parameters like age, gender and ethnicity. Even a single kernel model works quite well when it is properly regularised, as we will be reporting at ICIP in November. Our industrial partner has been keen for us to extend the estimator to attractiveness. We had only begun to think about this in mid-July, but the *One & Other* opportunity was too good to miss. I would be part of an artwork right in front of the National Gallery where two of Rubens' renditions of the famous

story about human beauty, *The Judgement of Paris*, are on display. The idea of a robot trundling around a gallery saying, "I don't know if it's art but I know what I like" was an appealing way to wind up multiple threads in the talk. So in the fortnight before the lecture I trained our face describer with judgements of human beauty and 'apparent trustworthiness'. Then I downloaded 200 pictures from the National Portrait Gallery available on Wikimedia (the Gallery and Wikimedia are currently in dispute about these), and ran the estimator on them all. Now in face description the character of the test set is even more determinative of experimental success than it is in face recognition, so with the sitters all being white and predominantly male, I can't claim that this experiment gives generalisable results. And because application to paintings means that

John Robinson braving the elements on the Fourth Plinth.
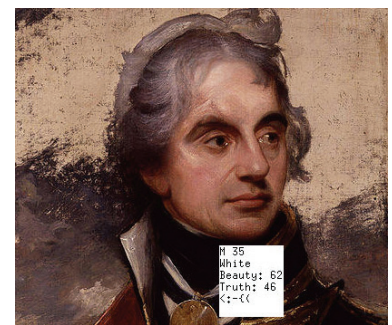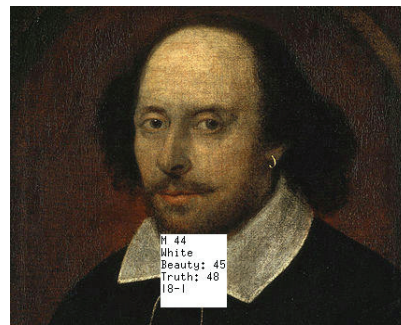
Examples of pictures with their Beauty and Truth estimators.

you are not only (or even primarily) judging the subject but rather the painter's ability to capture something of them, it is all too tempting to explain away mediocre performance. Yet even with those caveats I was delighted at how effective the primitive Beauty and Truth (trustworthiness) estimator was on those 200 portraits. As I said in the lecture, it gave plausible judgements on about 80% of them. Here are some examples. The tags also show the sex, age, ethnicity and expression (in the emoticon) estimates from our standard system. But the Beauty and Truth ratings are scores out of 100 for my hurriedly-trained parameters. The whole lecture is archived at:

http://www.oneandother.co.uk/participants/John_R.

The portrait judgement section, about three-quarters through, includes these pictures and more, and even manages to drag Lenna in again. I expect it (and this BMVA News column) will be the only presentation of these results because we are now training attractiveness parameters more carefully and systematically. But it was certainly worth having the first version going for the lecture, as it generated a lot of email feedback. Overall, the lecture was an intriguing experience – certainly one I'd do again the next time a plinth comes along.

Professor John Robinson
University of York
email: jar11@ohm.york.ac.uk

Image Analysis for Digital Pathology

Call for Participation

One-day BMVA technical meeting in London, UK on 7 December 2009

Chairs: Derek Magee (University of Leeds), Stephen McKenna (University of Dundee), Nasir Rajpoot (University of Warwick)

Digital image acquisition systems for pathology slides, tissue micro arrays (TMAs) are rapidly becoming ubiquitous in Pathology labs around the world. Add to this the availability at affordable prices of computers that are increasingly more powerful and offer progressively higher storage capacity. The combination of these two factors is spurring a revolution in Pathology, ultimately bringing it into the digital information era, with its own benefits, promise, and challenges. Digital acquisition coupled with traditional 'omics' methods results in massive amounts of data, most of it in the form of images. More than ever before, there is a need for the image processing and analysis community and the pathology community to interact in order to harness the power of this information-rich data.

This meeting aims to provide a platform for:

* Clinical researchers working in different areas of digital pathology for purposes such as diagnosis, prognosis, drug discovery, and biomarker discovery.
* Engineers and scientists working in image analysis, computer vision, machine learning, and pattern recognition.

In addition to inviting speakers at the forefront of this newly emerging area, we are seeking contributions describing recent work at the confluence of image computing and digital pathology. Potential topics include, but are not limited to, system architectures, image coding, viewing, segmentation, classification, multi-modal registration, modelling of important structures such as glands, etc.

Please submit an extended summary of about one A4-sized page (no more than two pages) in length. Please send contributions by email attachment (PDF preferred, 2Mb max) to Nasir Rajpoot by 30 October 2009.

Dr. Nasir Rajpoot
University of Warwick
email: nasir@dcs.warwick.ac.uk

Articulated Human Motion

One-day BMVA technical meeting at the British Computer Society, 5 Southampton Street, London, UK on 23 October 2009.

Chairs: Dimitrios Makris (Kingston University), Aphrodite Galata (University of Manchester)

10:00 Registration and coffee
10:35 Welcome and Introduction
10:45 Pose Estimation using Low-Level Motion. Ben Daubney, David Gibson and Neill Campbell (Swansea University, University of Bristol)
11:15 Activity Independent 2D Pose Recovery from a Single Uncalibrated Video. Paul Kuo, Dimitrios Makris and Jean-Christophe Nebel (Kingston University London)
11:45 Combining Discriminative Appearance and Segmentation Cues for Articulated Human Pose Estimation. Sam Johnson and Mark Everingham (University of Leeds)
12:15 Lunch
13:15 2D Shape-Skeleton Models for Human Tracking in Unconstrained Monocular Sequences. Jesús Martínez del Rincón, Gregory Rogez, Carlos Orrite Uruñuela (Kingston University London, Oxford Brookes, University of Zaragoza)
13:45 Articulated Structure from Motion using Metric Projections. Marco Paladini, Alessio Del Bue, Marko Stoši, Marija Dodig, João Xavier, Lourdes Agapito (QMUL, ISR – Instituto Superior Técnico, Universidade de Lisboa)
14:15 Invited Talk: Randomized Trees for Human Pose Detection. Gregory Rogez, Phil Torr (Oxford Brookes)
14:45 Tea
15:15 A Spatio-Temporal Extension of Laplacian Eigenmaps for Nonlinear Dimension Reduction of the space of human motion. Michal Lewandowski, Jean-Christophe Nebel (Kingston University London)
15:45 Real-Time 3-D Human Body Tracking using Learnt Models of Behaviour. Shaobo Hou, Aphrodite Galata (University of Manchester)
16:15 Tracking Known and Unknown Human Activities. John Darby, Baihua Li, Nicholas Costen (Manchester Metropolitan University)
16:45 Closing remarks and finish

Registration Form: 23 October 2009 Meeting

Please return this form to BMVA Secretary, Royston Parkin, 95 Queen Street, Sheffield, S1 1WG, Tel 0114 272 0306, Fax 0114 272 6158 or via email to BMVA@roystonparkin.co.uk. The meeting is free to members of the BMVA but a charge of £20 is payable

by non-members. A sandwich lunch is available at a cost of £5 and should be booked in advance. When registering please enclose a cheque for the appropriate amount made payable to "The British Machine Vision Association". Book early to avoid disappointment.

NAME: …………………………………………
ADDRESS: …………………………………………
                      …………………………………………
tel: …………………………………………
email: …………………………………………

BMVA Member: YES/NO
Lunch: YES/NO
Vegetarian: YES/NO

Dr Dimitrios Makris
Kingston University
email: d.makris@kingston.ac.uk

PR Technologies for Anti-Terrorism Applications

Call for Papers for a Special Issue of IEEE T. SMC.

The tragic events of 11 September in New York and the subsequent terrorist attacks around the world that aimed to inflict random and heavy casualties among non-combatants and for no reason other than causing death and mayhem. Modern terrorists have access to IT technologies, and to dual use materials that can be turned easily into crude, but nevertheless effective, weapons of mass murder. Traditional law enforcement technologies are becoming less than appropriate to deal with 21st century terrorists. Emerging counter-terrorism technologies are attracting significant attention as well as investment of time and effort among researchers and industrial developers. Pattern recognition plays an increasingly important role in the fight against modern form of global terrorism. Pattern recognition technologies can be used for surveillance, airport and border security, detection of hazardous chemicals in environment, are all but few examples. In recent years, there has been a rapid growth in the volume of publications on cutting edge pattern recognition and system integration technologies for anti-terrorism applications.

The main aim of the proposed special issue is to promote further research on pattern recognition technologies to address the increasingly complex challenges that arise within the context of counter-Terrorism applications. It is hoped that the special issue will attract wide interest in the broader research communities spanning pattern recognition, image processing, behaviour analysis, and other related areas. The topics to be considered include:

Ad hoc integrated network technologies for emergency responders; Automatic surveillance and analysis systems of terrorist activities; Biometrics; Crime scene analysis; Detection of unattended packages; Explosives Detection; Hazardous chemicals detection; Human abnormal behaviour understanding; Landmine Detection; Protection against Cyber-terrorism; Weapon Detection.

Manuscripts should conform to the standard guidelines of IEEE Trans. SMC, Part C. Instructions for formatting papers appear in the "Guide for authors" at http://www.ieee.org/web/publications/authors/index.html. Prospective authors should submit an electronic copy of their complete manuscript to Dr. Jinshan Tang or one of the other Guest Editors.

Important dates

Manuscript submissions due: 1 Jan 2010
Final manuscript due: 1 Nov 2010
Planned publication: Spring 2011

Dr. Jinshan Tang
Alcorn State University
email: jtang@alcorn.edu

Book for Review

The following book is on offer for review. As always, it will be sent out on a first come–first served basis, so please contact me immediately if you would like to review it. (If you are doubtful, go for the *quick view* option, and then return the book to me if you would rather not do the review.)

Rodri Davies, Carole Twining, Chris Taylor
Statistical Models of Shape: Optimisation and Evaluation
Springer-Verlag, London, 2008
314 pages, 57 Figures, £55
ISBN: 978-1-84800-137-4

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

CVPR 2010 – Call for Papers

CVPR 2010 will be held at the Hyatt Regency in San Francisco, California, 13–18 June 2010. For further information and updates, see: http://www.cvpr2010.org

Programme Chairs: Trevor Darrell, David Hogg, David Jacobs.

Papers in the main technical program must describe high-quality, original research. Topics of interest within computer vision and pattern recognition include, but are not limited to, the following:

- Applications of Computer Vision
- Colour and Texture
- Computational Photography and Video
- Early and Biologically-Inspired Vision
- Face and Gesture Analysis
- Illumination and Reflectance Modelling
- Image and Video Retrieval
- Image-Based Modelling
- Medical Image Analysis
- Motion and Tracking
- Object Detection and Categorization
- Object Recognition
- Performance Evaluation
- Segmentation and Grouping
- Sensors
- Shape Representation and Matching
- Shape-from-X
- Statistical Methods and Learning
- Stereo and Structure from Motion
- Video Analysis and Event Recognition
- Vision for Graphics
- Vision for Robotics

Important dates

Paper Registration Deadline: 12 November 2009
Paper Submission Deadline: 19 November 2009
Camera-Ready Deadline: 31 March 2010

N.B. Submission dates are fixed: no extensions will be given.

Paper submission

All submissions will be handled electronically via the conference website http://www.cvpr2010.org. Reviewing will be double-blind. In submitting a manuscript to CVPR, authors acknowledge that no paper substantially similar in content has been or will be submitted to another conference or workshop during the review period (19 November 2009–24 February 2010).

Tutorials, workshops and demos

In addition to the main technical program, the conference will include workshops, short courses and demonstrations. Proposals should be submitted to the appropriate chair:

Workshops: Silvio Savarese (silvio@eecs.umich.edu)
Short courses: Serge Belongie (sjb@cs.ucsd.edu)
Demos: Yuri Ivanov (yivanov@merl.com)

Dima Al Damen
University of Leeds
email: scs2dad@leeds.ac.uk

Report on AIAA Guidance, Navigation, and Control Conference

The AIAA Guidance, Navigation, and Control Conference was held on 10–13 August at the Hyatt Regency McCormick Place in Chicago, Illinois, USA.

Hyatt Regency McCormick Place

Each year this conference attracts hundreds of people from the field of aerospace and this year's edition was no exception. Over 600 papers were presented by researchers from industry, the military and academia, on a wide array of topics including flight simulation, aircraft dynamics and control, flight testing, unmanned vehicles, trajectory optimisation, system identification, formation flying, and air traffic management. The conference was organised into 121 sessions that included the paper presentations and four workshops.

The first three days of the conference started with a plenary session. The first plenary was presented by Graham Tomlinson, a lead test pilot in the F-35 Joint Strike Fighter (JSF) Program. His talk was about the

simulations and flight tests conducted on the F-35 Lightning II and, previously, on the Harrier. Particular emphasis was made on the Short Take Off and Vertical Landing (STOVL) capabilities of the two aircraft. The speaker of the second plenary was Professor Gareth D. Padfield from the University of Liverpool in the UK and his presentation concentrated on Atmospheric Flight Mechanics. The third plenary was a historical take on different navigation techniques that eventually culminated in the development of satellite navigation and the invention of the Global Positioning System (GPS). This plenary was presented by Richard Easton on behalf of his father Roger Easton who used to work at the Naval Research Laboratory and is considered to be the principal inventor and designer of GPS.

Chicago skyline

I presented my paper on the second day of the conference, during the session "Aircraft and Spacecraft Sensor Systems". My presentation was about the detection of generic obstacles around large commercial aircraft on ramps and taxiways through the use of computer vision techniques. The presentation went well and positive feedback was obtained from the audience.

Although the conference was not directly related to machine vision, a number of sessions were almost entirely dedicated to computer vision-based systems. Some of the applications presented were: autonomous vision-based landings of unmanned aerial vehicles

(UAVs) and helicopters, stereoscopic vision-based spacecraft relative state estimation, air-to-ground target tracking using optical flow estimation, aircraft navigation using online mosaic image construction, and monocular 3D obstacle detection. This proves the ever-growing role of machine vision in aerospace applications.

Even though the conference kept me busy most of the time, I had the opportunity to discover some parts of the beautiful city of Chicago. A visit to the John Hancock Observatory (known to the locals as 'Big John') allowed me to admire the city from a height of around 1000 ft. Back on the ground I walked through the Millennium Park, went past the Jay Pritzker Pavilion (where an orchestra was performing outdoors), and strolled along the shore of Lake Michigan, before returning to the vibrant atmosphere of the city centre. The locals recommended that I try the famous Chicago-style pizza, which I did. I was not disappointed.

The next AIAA Guidance, Navigation and Control conference will take place on 2–5 August next year, at Toronto, Ontario, Canada. I would like to conclude by thanking the BMVA for kindly providing me with a travel bursary to attend this year's conference.

Jason Gauci
Cranfield University
email: j.gauci@cranfield.ac.uk

Solution to Conundrum (page 2)

Most editors prefer to start papers on the right-hand side. So if you start a paper thus with an odd number of pages, you try to follow it with another odd-numbered one to get the 'page parity' right as soon as possible.

This suggests that only about 1/4 of the papers should start on the left. I did once count a journal and it came close to that.

Chris Trayner