

BMVA News

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

Volume 28 Number 3
March 2018

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



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

BMVA News¹ is published every three months. Contributions on any activity related to machine vision or pattern recognition are eagerly sought. These could include reports on technical activities such as conferences, workshops or other meetings. Items of timely or topical interest are also particularly welcome; these might include details of funding initiatives, programmatic reports from ongoing projects and standards activities. Items for the next edition should reach the Editor by 10 June 2018.

Contents

Editorial: <i>In Research, as in Life, the Biggest Difficulty is Knowing When to Stop</i>	1
MIUA 2018	2
BMVC 2018 – Call for Papers	3
BMVA Computer Vision Summer School, 2018.....	4
Cognitively inspired eXplainable Perception-based AI	4
BMVA Workshop on Reinforcement Learning	5
BMVA Symposium: Robotics meets Semantics	6
Workshop on Robot Perception of Humans	6
Call for Expressions of Interest to Host BMVC in 2020.....	7
Book Review: E. R. Davies <i>Computer Vision</i>	8
Nominations for the BMVA Executive Committee	10
Nomination Form for the BMVA Executive Committee ...	10

Editorial: *In Research, as in Life, the Biggest Difficulty is Knowing When to Stop*

I imagine that many of us will have been in the situation that our boss, leader or head of department hasn't been living up to our original expectations (though on reflection this may not be too surprising as (s)he may well not have been trained for management vis-à-vis technical achievement). However,

the situation will be a double whammy, as not only is (s)he not doing a good job, but also (s)he will be holding a better person out. My point is that there is often an optimal period for people to hold particular jobs. Release them too early and they will not have had the chance to get properly started and prove themselves: release them too late and achievement levels will have been substantially reduced. Actually, this also goes for us: there is bound to be an optimum time to retire or move on. The difficulty is knowing quite when to stop. In fact, this also applies to tasks we undertake. When starting on a particular programme of research, we will adopt a particular strategy: obviously, we only attempt strategies that offer some reward and which present a clear path ahead from the very beginning. However, as has many times been said, “life is what happens while you're planning other things”. Practical factors and realities catch up with us and with the proposed strategy. Nevertheless, it frequently seems best to plough on regardless – or at least to negotiate obstacles in the best available way. Indeed, over the following days and months, it often seems that *one more push* is all that will be needed to achieve a breakthrough compatible with the original intentions of the project. We see this continually in other areas such as politics, where one more ounce of pressure or funding is expected to solve the problems of some aspect of education or national health. But the ‘one more push’ strategy is ill thought through and frequently (or even usually) ends up as a palliative.

I had all this brought home to me when I started work on image analysis. Following edge detection, edges seldom formed neat outlines of objects: there were often short breaks of one or two pixels and on many occasions significantly larger breaks. The very short breaks could often be eliminated by joining neighbouring ends, but more complex strategies were needed for larger breaks – e.g., allow joins only if the ends to be joined point in approximately the same direction – which also has to be the direction of the join. Clearly (in retrospect – after the long period from the early 1980s to today), this approach was flawed and indeed it tended to produce object outlines that wander around images joining up random shadows and parts of the background. It took me some time to realise quite what the data was trying to tell me. To cut a long story short, ad hoc joining was not clever but using model-based

¹ The British Machine Vision Association and Society for Pattern Recognition is a Company limited by guarantee, No. 2543446, registered in England and Wales. Registered Office: Granta Lodge, 71 Graham Road, Malvern, WR14 2JS. The Association is a non-profit-making body and is registered as charity No. 1002307.

methods such as the Hough transform was key. Nowadays we might say that we are even wiser and that probability based methods should be even better. Anyway, my point is that the ‘one more push’ way of working is on a par with death: what is more, we need to have safeguards against it. At the very least, we need to know when to put a stop on any proposed strategy. Actually, I had an even earlier inkling about all this. In his (1979) book on image analysis, Ernie Hall said about image enhancement “Much of the art of enhancement is knowing when to stop”: of course, image enhancement is a subjective activity – viz., improving the *appearance* of images (which contrasts with image restoration – a mathematical optimisation task), so what he said made very good sense.

But how do you know when to stop? Just as an epidemic can be defined as an continuing expansion in the spread of a disease beyond what would a priori be expected, so knowing when to stop can be considered as the point where progress drops significantly below reasonable expectations. To guard against this, it is useful to have other irons in the fire, so one can make comparisons, or push the methods that currently have the strongest promise. E.g., in my example above, one could set Hough transforms against intuitive joining strategies; and then add RANSAC to the list; or look for solutions (objects) corresponding to the greatest a posteriori probabilities.

The trouble with this general approach is the investment of effort needed to bone up on and test all strategies that could reasonably be applied. So it is no surprise that workers tend to have a vested resistance against trying too many alternatives; and there are even people who seem to work with the same methodologies for most of their careers. After all, they have been successfully using morphology or relaxation methods or neural networks (to name but a few) for such a long time that they feel they can’t leave their area of expertise behind and delve into other mostly unrelated methods. Again, leaving their area of expertise behind means loss of practice at it, declining expertise with it and even a tacit admission that it was never the best thing since sliced bread. On the other hand, in the grand scheme of things, we can appreciate that different scientific methods come into and go out of fashion, or at least have their moments of success in particular areas. What is more, the people who promote different methods may themselves come into and go out of prominence, perhaps because their expertise is more or less practical, more or less theoretical, more or less mathematical or more or less probability based. And lately we have seen that people who are not sensitive to the possibilities of neural-based strategies (aka deep learning) may be putting themselves at risk. (Albeit, there is the possibility that people are putting themselves at risk by *avoiding* conventional scientific methods.)

Be all this as it may, I have recently been reminded in no uncertain terms that anyone in any profession (i.e., not just computer vision) can be subject to these problems. In fact, this probably applies more to people outside computer-based subject areas, as computer-based workers have to be more adaptable than those who work with real systems. For example, doctors may be experts in very specific areas such as kidney transplants, eye problems, breast cancer, and so on: thus, they may have vested interests in continuing with their own specific ways of doing things and may excusably ignore for years what is happening in cognate areas. With the pressures on the NHS, this may not be too surprising, for

it is far easier to continue with one’s old expertise and way of doing things than to take a break to explore other methodologies and procedures. But in the end, international pressures will tend to win out; and anyway, a minority of workers will not be changed by persuasion, which may result in their ending up in no-man’s land. After all, Albert Einstein never actually believed in quantum mechanics – recall that he said “God does not play dice with the universe”. But he and others of that ilk died and quantum mechanics survived ...

Postscript

In my last editorial I spent a little time reminiscing about my father’s choice of reading matter and recalling the story of the man whose eyes saw other scenes than those in front of him. The title of the work eluded me, but luckily, Professor Jenny Read of Newcastle University has informed me that this short story – by none other than HG Wells – is *The Remarkable Case of Davidson’s Eyes* and was written in 1895. For anyone who is interested, the full text can be found at

<http://www.online-literature.com/wellshg/2867/>

Unfortunately, over the many years since my youth, I had built up such great expectations for this story that there was no way the real thing could match them, imaginative though the story is. C’est la vie – ou peut-être, c’est la psychologie.

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

MIUA 2018

You are invited to attend the BMVA sponsored 22nd Medical Image Understanding and Analysis Conference (MIUA 2018) which will be held at Southampton 9–11 July 2018. We promise a topical programme of papers in medical understanding enriched by two special sessions (one on retina and the other on liver image analysis). Our programme will be further enriched by keynote speakers including Professor Alison Noble (University of Oxford), Professor Tom Wilkinson (University of Southampton) and Professor Anant Madabhushi (Case Western Reserve University, USA).

You can follow our tweets @miua2018 and find our event details online at:

<https://miua2018.soton.ac.uk/keynote.html>

We promise understanding and analysis of medical images of many parts of the human body. We expect to present deep learning and modern approaches. We look forward to welcoming you to Southampton.

Conference Chairs: Mark Nixon, Sasan Mahmoodi and Reyer Zwiggelaar.

Professor Mark Nixon
University of Southampton
email: msn@ecs.soton.ac.uk



BMVC 2018 website: <http://www.bmvc2018.org>

BMVC 2018 will be held in Newcastle upon Tyne, UK at Northumbria University. BMVC is the BMVA annual conference on machine vision, image processing, and pattern recognition. It is one of the major international conferences on computer vision and related areas held in the UK. Because of its increasing popularity and quality, it has established itself as a prestigious event on the vision calendar.

There have been 28 successful BMVCs since 1985. They were organized in many major universities in the UK, such as Cambridge (1987), Oxford (1990), Birmingham (1995), UEA (2003), Edinburgh (2006), Leeds (2008) and Imperial College London (2017). The scale of the conference has increased significantly in the past decade. In particular, the 2017 conference attracted 450 attendants and 635 full paper submissions, which was a record high. 188 high-quality research papers were accepted, in which around one-third of the workers came from Europe, one-third from Asia and one-third from North America. Apart from international researchers, the conference also attracted a wide spectrum of international companies, which sponsored the conference and used it to demonstrate their products.

Important dates

- Conference date: 3–6 September
 - Tutorial: 3 September
 - Main conference: 4–6 September
 - Workshop: 6 September (afternoon)
- Submission deadline: 30 April
- Acceptance notification: 2 July.

Call for papers

Authors are invited to submit full-length high-quality papers on image processing and machine vision. Papers covering theory and/or application areas of computer vision are invited for submission. Submitted papers will be refereed on their originality, presentation, empirical results, and quality of evaluation.

All papers will be reviewed doubly blind, normally by three members of our international programme committee. Please note that BMVC is a single track meeting with oral and poster presentations and will include three keynote presentations and three tutorials.

Topics include, but are not limited to:

- Statistics and machine learning for vision
- Stereo, calibration, geometric modelling and processing
- Face and gesture recognition
- Early and biologically inspired vision
- Motion, flow and tracking
- Segmentation and grouping
- Model-based vision
- Image processing techniques and methods
- Texture, shape and colour
- Video analysis
- Document processing and recognition
- Vision for quality assurance, medical diagnosis, etc.
- Vision for visualization, interaction, and graphics
- Object detection and recognition
- Shape-from-X
- Video analysis for action and event recognition
- Illumination and reflectance
- Deep learning for vision
- 3D computer vision
- RGBD analysis.

Submission details

Paper submission and registration are handled on CMT:
<https://cmt3.research.microsoft.com/bmvc2018>

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

<http://www.bmvc2018.org/faqs.html>

Formatting your paper

A complete paper should be submitted according to the formatting style. The length should not exceed 9 pages excluding the references. Any appendices must precede the reference list.

Submitted papers

Papers submitted for review should be anonymous and instead include the paper ID as allocated by CMT upon registration of the paper. The addition of line numbers should be included, in order to facilitate the reviewing process.

If you have submitted papers to previous BMVCs, note that the 9 pages are not measured from the abstract. They are measured from the top of the page; from the 10th page onward only references are allowed. This is to avoid ambiguity in measuring the precise page length. The final accepted papers are allowed some extra space to add author information, acknowledgements where appropriate, minor changes to address referee comments and the like.

The above details are taken from the full set of instructions and information on the website below. *It is vital to refer to the full instructions before submitting your paper.*

<http://www.bmvc2018.org/submityourpaper.html>

Keynote speakers

Professor Rama Chellappa, Professor Sven Dickinson, and Professor Shaogang Gong. Details of the keynote speakers can be found at:

<http://www.bmvc2018.org/keynotespeakers.html>

Tutorial speakers

Professor Vittorio Ferrari, Dr Ivan Laptev and Dr Abhinav Gupta. Details of the tutorial speakers can be found at:

<http://www.bmvc2018.org/tutorials.html>

Dr Hubert PH Shum
Northumbria University
email: hubert.shum@northumbria.ac.uk

BMVA Computer Vision Summer School, 2018



BMVA runs an annual Computer Vision Summer School aimed at first-year PhD students and other early-stage researchers. This year it will be held at the University of East Anglia on 2–6 July. Despite the title, students from non-UK universities are welcome to attend, as well as students from UK universities. Places are limited to ensure good interaction in lab classes.



The Summer School will consist of an intensive week of lectures and lab sessions covering a wide range of topics in Computer Vision. Lecturers are researchers in fields from some of the most active research groups in the UK and abroad. In addition to the academic content, the Summer School provides a networking opportunity for students to interact with their peers, and to make contacts among those who will be the active researchers of their own generation.

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

<http://cvss-uea.uk>

Twitter: @BmvaCvss

Important Dates

Early registration: 11 May

Late registration: 22 June

Summer school: Monday 2 July – Friday 6 July

Note that onsite accommodation is only guaranteed for early registrations.

Dr Michal Mackiewicz
University of East Anglia
email: M.Mackiewicz@uea.ac.uk

Meeting Report: Cognitively inspired eXplainable Perception-based AI



This one-day BMVA technical meeting was held at the BCS in London on 7 February 2018. It was organized and chaired by Associate Professors Serge Thill (University of Plymouth, UK) and Maria Riveiro (University of Skövde, Sweden).

Explainable AI (XAI) refers to AI systems that behave or provide the necessary information so that their inner workings become comprehensible to the human user. Due to the fact that AI systems are increasingly present in everyday society, from simple computer systems to agents such as autonomous vehicles or social robots, several researchers have noted that it is critical to understand how human users perceive such systems. Therefore, the main goal of this highly interdisciplinary one-day BMVA workshop was to bring together contributions from the fields of robotics, AI, cognitive science, machine vision and HCI to outline the current state-of-the-art of XAI and talk about future challenges.

This workshop was supported by the BMVA as well as the research projects Dreams4Cars (H2020 grant agreement no. 731593) and BIDADF (Big Data Analytics Framework for a Smart Society, The Swedish Knowledge Foundation).

Keynotes

We were delighted to have three amazing keynote speakers in this workshop, Dr Alessandra Sciutti (Istituto Italiano di Tecnologia, Italy), Assistant Professor Brad Hayes (University of Boulder, Colorado, USA) and Professor Yiannis Demiris (Imperial College, London) who provided their views on XAI from a robotics and human-machine collaboration perspective.

Alessandra Sciutti started the workshop, presenting her group's current research aimed at defining which features of human and robot motion allow for natural mutual

understanding, referencing to low-level kinematics properties (as biological motion) and higher-level cognitive aspects (as intention reading and attribution). Alessandra reflected on the fact that we are still far away from fully interactive robots, but we are moving to inclusion of robots in society, though many challenges need to be solved. Alessandra elaborated on many relevant issues for XAI, like intuitiveness, implicit signals revealing internal states of collaborating partners, prospective signals, predictability and communication of intent. Brad Hayes started the afternoon session, talking about robot collaboration and teamwork. He addressed relevant aspects of his research with supportive robots for XAI, such as communication of shared expectations, interpretability of models, reasonable behaviour, learning by demonstration, and so on. Another interesting discussion brought up by Brad was, how do we arrange for the human to have a theory of mind about the robot? Yiannis Demiris closed the speakers' sessions talking about personalized robot assistance and human modelling in areas like robotic wheelchair assistance and learning. He focused on robots that help and provide support, rather than on capability replacement, e.g., for children with disabilities or the elderly. He suggested that explanations will be more understandable if presented closer to the internal models of the user; thus, we need to create personalised mental models of the users in each case. He coincided with Alessandra and Brad by introducing important aspects like intention recognition.

Theoretical aspects of XAI

Tarek Besold (City, University of London) characterises four notions of explainable AI, focusing on: (1) *opaque systems* (which offer no insight into their algorithmic mechanisms), (2) *interpretable systems* (users can mathematically analyse the algorithmic mechanisms), (3) *comprehensible systems* (which emit symbols enabling user-driven explanations of how a conclusion is reached) and finally, (4) *truly explainable systems* (where automated reasoning is central to output crafted explanations without requiring human post-processing as a final step of the generative process). Henrik Svensson (University of Skövde) suggested the application of the situated and distributed cognition perspective to XAI. He provided several examples from the area of autonomous cars, and from his own and his colleagues' research in the Dream4Cars project. Tove Helldin (University of Skövde) provided a brief introduction to evaluation and metrics for explainable AI, providing lessons learned from other areas like HCI, robotics, information visualization and recommender systems.

General applications of XAI

Elin A Topp (University of Lund) started the morning session after Alessandra talking about how to generate meaningful clarification requests. She presented her research with industrial robots and her work on extracting knowledge from experts to teach robots to do certain tasks. She discussed relevant issues for XAI like the need for the robot to understand what people actually mean (not say), viz., action representation. Szonya Durant (Royal Holloway, University of London) presented her and Professor Kostas Stathis' research on integrated cognitive user assistance for multiple-display systems. They focused on methods based

on eye-tracking, where location and duration of fixations are used to measure locus and timing of user attention. Sylvester Kacmarek (Imperial College) presented a series of works and studies carried out with several companies in the field of applied AI focusing on transparency, trust, reinforcement learning, and so on. Finally, Swen Gaudl (Metamakers Institute, Falmouth) brought to the workshop a 'games' perspective and talked about XAI, games, genetic programming and human mimicry. Swen presented an agent framework approach using genetic programming to derive new reasoning agents that can be interpreted by human users.

Panel discussion and final remarks

The meeting concluded with a panel discussion by the three keynote speakers. Several issues were discussed, such as "Why XAI?"; "What does the user want to know from the system? (e.g., does he want to be able to predict the system behaviour?); and lessons learned from other areas, such as decision support systems.

From our point of view, it was a very interesting meeting that fully engaged the participants in interactive discussion with the presenters, bringing together researchers and industry representatives from disparate areas such as robotics, AI, HCI, computer vision, cognitive science, etc. It is evident that there are many exciting challenges ahead for XAI, which need interdisciplinary solutions and the cooperation of several scientific areas.

Talk recordings

Selected recordings of the talks are available on the BMVA Youtube channel <https://tinyurl.com/ycbojtny>

Serge Thill
University of Plymouth
email: serge.thill@his.se

BMVA Workshop on Reinforcement Learning in Computer Vision



Location: British Computer Society, London, Wednesday 9 May 2018

Chairs: Dr Dmitry Kangin and Dr Nicolas Pugeault (Exeter)

www.bmva.org/meetings

Reinforcement learning offers the attractive feature of learning directly from a system's performance at a task. The fast development of deep learning has allowed applying reinforcement learning to new classes of problems: playing

video games, performing object tracking, predicting eye gaze direction, controlling robots and many others. However, many challenges remain. The high-dimensional input spaces require millions of data samples if they are to converge; this prohibits their use for many real-world tasks, in contrast to simulated tasks.


This workshop aims at bringing together experts in machine learning and computer vision to discuss the present and future of reinforcement learning in vision problems.

Programme

- 09.00 Registration
- 09.30 Welcome and introduction
- 09.40 Keynote: Dr Sen Wang
- 10.40 Sinisa Stekovic, TU Graz. "Experience replay for learning multiple environments"
- 11.10 Coffee Break
- 11.30 Keynote: Dr Subramanian Ramamoorthy
- 12.30 Lunch break
- 13.30 Keynote: Prof Shimon Whiteson
- 14.30 Dr Dmitry Kangin / Dr Nicolas Pugeault. "Reinforcement learning and Computer Vision"
- 15.00 Coffee break
- 15.30 Keynote: Dr Anil Bharath, Imperial College London
- 16.30 Discussion and conclusions
- 17.00 End.

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

BMVA Symposium: Robotics meets Semantics: Enabling Human-Level Understanding in Robots



This one-day BMVA Symposium will be held in London on Wednesday 18 July 2018 (www.bmva.org/meetings).

Chairs

Oscar Mendez University of Surrey and Mihai Bujanca University of Manchester

Keynote Speakers

Cyrril Stachniss, University of Bonn, Daniel Cremers, TU Munich and Andrew Davison, Imperial College London

Call for Papers

For decades, 'Bridging the Semantic Gap' has been a catchphrase in the fields of Computer Vision and Robotics. Whilst a nuanced semantic understanding of the world is still an open problem in both fields, there have been important advances in vision-based Machine Learning that enable semantic-based approaches in both Robotics and Vision. Specifically, the field of Semantic Segmentation has begun to produce results that are reliable enough to be used in robotic systems that perform navigation, grasping, localisation, reconstruction, SLAM and even human interaction. This workshop aims to bring together experts in Vision-based Machine Learning and Robotics to discuss the applicability of semantically-enabled vision to robotics.

Submission Deadline

We wish to encourage submissions from students as well as academic and industry researchers in the area. All those interested in presenting at this meeting are invited to submit a summary of their talk to o.mendez@surrey.ac.uk by 18 May 2018.

Registration

Book online at www.bmva.org/meetings: BMVA Members £16, non-Members £36 (lunch is included in both prices).

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

Workshop on Robot Perception of Humans

Call for Papers

This workshop will take place in Baden-Baden, Germany on 11 June 2018. It is being organised in conjunction with the 15th International Conference on Intelligent Autonomous Systems (IAS-15):

<http://www.aisl.cs.tut.ac.jp/ws-rph/>

There is an increasing demand for robots and intelligent systems that co-exist with and support humans in their everyday lives. A key to realize such robots and systems is reliable perception of humans. They can behave intelligently and effectively by reliably detecting humans and understanding their behaviours, identities, and physical/mental states.

Visual and sensor-based perception of humans has been an important research area in robotics, and its application area is expanding, accelerated by recent progress with machine learning algorithms and deep neural networks. This full-day workshop brings together researchers and industry

engineers working on novel approaches and applications and practical applications of perception of humans.

The goal of the workshop is to recognize the state-of-the-art of human perception research and to discuss its future directions and potential application scenarios. The workshop consists of several invited talks by leading researchers and oral and poster presentation of contributed papers.

Research and application papers are solicited in, but will not be limited to, the following area topics:

- human detection and tracking
- person re-identification
- human pose estimation
- face detection and recognition
- gesture recognition
- human anomaly detection for care and service robots
- human behavior analysis and surveillance
- human-robot interaction and collaboration
- perception of humans for mobile robots and self-driving cars
- deep learning-based approaches to perception of humans
- systems and software tools for perception of humans
- human detection and recognition in/for industrial scenarios.

Selected contributed papers are to be invited to a special issue on “robot perception of humans” of the Robotics Journal (ISSN 2218-6581).

Important dates

Submission deadline: 15 April
 Notification of acceptance: 8 May
 Camera ready copy deadline: 22 May
 Workshop at IAS-15: 11 June.

Organizers

- Jun Miura, Chair (Toyohashi University of Technology)
- Emanuele Menegatti (University of Padua)
- Nicola Bellotto (University of Lincoln)
- Stefano Ghidoni (University of Padua).

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

Interested in running BMVC in 2020?

As you will know, the first British Machine Vision Conference was held in 1990, so the conference in 2020 marks thirty years since its institution. It has grown dramatically over recent years, becoming one of the most important annual vision conferences globally – so much so that we expect over 500 delegates in 2020. We would like to make the 2020 event not just a conference but a celebration of computer vision: *Vision 2020* has a something of a ring to it.

The BMVA Executive Committee is looking for institutions interested in hosting BMVC in this milestone year. At this stage only an expression of interest is required: following discussion by the Executive Committee, these expressions of interest may be followed up by requests to supply an official bid to hold the conference. For the expression of interest please supply the following details:

- Main contact for the conference: name, email, phone.
- Prospective members of the conference organising committee: general chair, programme chairs, local arrangements chairs, and so on. There is no need for all these members to be from the same institution; indeed, we encourage you to spread the load across several institutions as BMVC is now a big conference.
- Provisional dates for the conference, including details of accommodation, rooms and arrangements for conference venue and meals.
- Ideas you have for making the event a celebration of computer vision as well as a conference.

BMVC is traditionally held in one of the first two weeks of September and runs as a single-track conference with podium and poster sessions from Tuesday through to Thursday lunchtime. The pre-conference Monday afternoon is normally scheduled for a tutorial session, and separate workshops may be held on the Thursday afternoon or on Friday.


The main conference auditorium needs to be large enough to accommodate over 500 delegates and should preferably be a tiered auditorium with adequate A/V facilities. The poster space should be capable of holding at least 150 1m-wide poster boards, ideally in the same contiguous space (e.g., a large hall). For further information on preparing a proposal to run BMVC please refer to:

http://www.bmva.org/bmvc_proposals

Please send expressions of interest to the BMVA Chair, Dr Adrian Clark, by 1 July 2018.

Adrian Clark
 BMVA Chair
 email: chair@bmva.org

Book Review: E. R. Davies *Computer Vision*



Computer Vision: Principles, Algorithms, Applications, Learning, E. R. Davies, Academic Press (5th edition, 2017)

What can one say about a Fifth Edition of a book that has not been said before? Well, in this case, quite a bit! For those who have not heard of this atlas of computer vision before, you should know that that the book is possibly THE best companion for immediate access to a myriad of techniques and methodologies in your beloved field.

With each new edition since the first, the author has significantly updated the book to dispose of less relevant or practiced methods and to incorporate the latest techniques, trends, and developments.

This edition is no exception. Indeed, this edition, compared to the most previous one, has 5 new chapters, drops 5 other chapters, and has material from 6 chapters rolled into 3. Despite this equilibrium in terms of quantity of content and the number of pages, the use of better quality paper and production wizardry has resulted in a remarkably thinner volume and a much, much better look-and-feel outcome. Further, this edition has also a dash of colour added to around 90 of the diagrams that helps with the quality and ease of understanding.

The reader will not be surprised to find coverage of deep learning, and CNNs (concept, architectures, and more) as some of the latest additions in this 5th Edition of the book. This is simply the logical next step in the book's evolution. What is pleasantly surprising is the seamless manner in

which the author has managed to integrate these topics into the flow of the book. The “Deep Learning Networks” chapter (Chapter 15) smoothly picks up where the earlier Chapter 13 left off with the introduction of artificial neural networks and the back-propagation algorithm. It covers the concept of CNNs in general and the parameters required for defining CNN architectures. Chapter 15 continues by reviewing several CNN architectures in some detail, such as AlexNet, VGGNet, and SegNet. The chapter concludes with a review of Recurrent NNs. Sandwiched between Chapters 13 and 15 is the chapter on “Machine learning: probabilistic methods”, which ties those two together.

The book content is far-reaching and of encyclopaedic extent. For reasons of brevity however, I can only attempt to provide an overview of the content using the neat, 5-part division of the book. The first part is Low-level Vision techniques, i.e., image filtering for feature extraction, e.g., edges, corners, invariants, and texture descriptors. The second part is Intermediate-level Vision and is centred around various aspects of shape analysis. Part 3 is on classification, machine learning, and deep learning, as detailed in the previous paragraph. The five chapters of the fourth part are all about the 3D world and motion. Finally, in Part 5, practical examples are provided that put what we have learnt to work, including chapters on in-vehicle vision systems, surveillance and monitoring of humans and animals, and the use of deep learning for face recognition.

This book will serve a spectrum of readers, from undergraduates studying on image processing and computer vision modules, to researchers who need a library of theories, techniques and methodologies at their fingertips.

All in one place! The author provides guidance that is worth briefly mentioning here: Chapters 1–13 and Chapter 16 would be mostly sufficient for an undergraduate course on computer vision for Electronic Engineering and Computer Science students, while MSc and PhD students might benefit from Chapters 13–20 in depth, and a selection of the final chapters, 21–24. In summary, I cannot think of a comparable book that provides the same “pragmatic blend of theory, implementation, and algorithms” – this time to quote an endorsement from the Foreword to this book by an authority no less than Professor Mark Nixon.

Professor Majid Mirmehdi
University of Bristol
email: m.mirmehdi@bristol.ac.uk

This page is deliberately left blank.

Nominations for the BMVA Executive Committee

Nominations are requested for the forthcoming election of Executive Committee members of the BMVA. Nominees must be paid-up members of the Association and agree to serve for a period of two years. A member of the Committee is expected to participate in roughly five committee meetings per year, taking place typically in London and Birmingham. Reasonable travelling expenses to attend meetings are paid by the BMVA.

Completed nomination forms should be sent to the BMVA Chair and must be received by 1 July. The nomination form must be signed by the individual standing and by one other member, and should also include a brief biographical statement for distribution to BMVA members.

The Executive Committee normally consists of ten elected members, five of whom are elected each year. Details of the current members may be found at:

http://www.bmva.org/executive_committee

If more than five nominations are received for the elected places then a postal ballot will be held. Voting papers will be sent out in early July and will need to be returned by the end of August. Each member will be able to vote for up to five candidates. The results will be announced at BMVC and in *BMVA News*.

Adrian Clark
BMVA Chair
email: chair@bmva.org

Nomination Form for the BMVA Executive Committee

To be completed by the Nominator

As a fully paid up member of the BMVA, I,

Name:

Address:

.....

.....

wish to nominate:

Name:

Institution:

for the Executive Committee of the BMVA.

Signed:

To be completed by the Nominee

I am a fully paid up member of the BMVA and am willing to serve for a minimum period of one year on the BMVA committee.

Name:

Address:

.....

.....

email:

I attach below a brief biography for distribution to BMVA members.

Signed:

Brief Biography of Nominee

Please send completed nomination form by post or email to:

Dr Adrian F Clark
Computer Science and Electronic Engineering
University of Essex
Wivenhoe Park
Colchester
Essex, CO4 3SQ
email: chair@bmva.org