THE INTERNATIONAL ASSOCIATION FOR PATTERN RECOGNITION



Click here to subscribe to the *IAPR Newsletter*

SSUE

From the Editor's Desk

CALLS for PAPERS

<u>Getting to Know...</u> <u>Tamás Szirányi, IAPR Fellow</u>

IAPR...The Next Generation: Yan Yan

From the ExCo

In Memoriam: Stefano Levialdi

<u>Meeting Reports</u>: <u>CCIW 2015, GbR 2015, ICB 2015</u> and <u>SCIA 2015</u>

Non-IAPR Report: <u>BMVA One Day Meeting on</u> <u>Real-time 3D scene understanding</u> <u>in 2020</u>

Free Book Offers

Bulletin Board and Pending Book Reviews

Meeting and Education Planner

From the Editor's Desk: Getting your paper rejected --Part 3 You convinced the reviewer to read your paper. What's next?

by Arjan Kuijper arjan.kuijper@igd.fraunhofer.de http://www.gris.tu-darmstadt.de/~akuijper/

In the two previous 'From the editor's desks' (*IAPR Newsletter* January 2015 and April 2015), I started counting down my top 10 reasons for rejecting a paper. They had to do quite a bit with visual characteristics, in the sense that an experienced reviewer can point out limitations of the paper within 20 seconds by noting:

- 10) Bad figures
- 9) No comparison with state-of-the-art
- 8) No experiments
- 7) Old references
- 6) No state-of-the-art or novelty

And, I added an extra reason, Reason 11: Wrong template.

Given any of the above observations, the reviewer doesn't feel obliged to even read the paper (seriously!). So avoiding these pitfalls is the number one task. Then, having successfully avoided the first five rejection reasons and convinced the reviewer that the paper is worth reading, you've got five more rejection reasons to avoid. These are discussed below. I should also menton, that before I (and I guess most people who will read your paper) start reading a paper, I usually scan it to get a first impression of the quality. So make sure this first impression is good.

The views expressed in this newsletter represent the personal views of the authors and not necessarily those of their host institutions or of the IAPR.



For the most up-to-date information on IAPR-supported conferences, workshops and summer schools, please visit the IAPR web site: <u>www.iapr.org/conferences/</u>

2015

ICVNZ 2015

Image and Vision Computing New Zealand Auckland, New Zealand Dates: Nov. 23-24, 2015 Deadline: Sep. 18, 2015

2016

ICB 2016 9th International Conference on Biometrics Halmstad, Sweden Dates: Jun. 14-17, 2016 Deadline: Jan. 11, 2016

ICPR 2016

23rd International Conference on Pattern Recognition Cancun, Mexico Dates: Dec. 4-8, 2016 Contest, tutorial and workshop proposal deadline: January 13, 2016 Paper submission deadline: April 3, 2016

Reason 5: Introduction and conclusions don't fit

Your paper describes a novel approach. So in the introduction you describe the problems you solve. The body of the paper gives the details and proofs. The conclusions – as the name says – show that the proofs indeed solve the problems tackled. The introduction and conclusions don't fit if, for example, you promise to solve three things in the introduction and conclude by mentioning only two of them, or by discussing four contributions; or you claim quite a lot and then weaken the claims later on. This leaves the reviewer puzzling what the true contribution is. You don't want to do that!

Reason 4: Conclusions vague

It should be clear whether or not you have solved the problem mentioned in the beginning of the paper. The conclusions must do more than rephrase the problem and restate the experimental results—they must also show how the experimental results support the claims you want to make. What can you conclude from and what are the implications of the experiments? That is the

ICISP 2016

7th International Conference on Image and Signal Processing Trois-Rivières, Québec, Canada Dates: May 30-June 1, 2016 Deadline: January 17, 2016

DAS 2016

12th IAPR International Workshop on Document Analysis Systems Santorini, Greece Dates: Apr. 11-14, 2016 Deadline: Oct. 2, 2016

> challenging task you have when writing the conclusions. As a reviewer, I don't want to guess what you have done and what the implications of your tests are. If I have to guess, I'll reject.

Reason 3: Contributions not stated in the introduction

This is connected to the previous points. In the introduction I want to get clear what the problem is and how you contribute to solving it. Mentioning that there is a problem is not enough. Then, what would the reader expect? Some ideas, hacks, and wild guesses to make some progress in solving the problem? Basically, I want to see here explicitly what is new and, also, what is not new. This answers two questions for me: What is this paper about? and What is this paper NOT about?" Getting a wrong impression of the contribution kills the paper.

Reason 2: Unclear abstract

Obviously, the abstract is the place where all these points come together. It summarizes the paper's contributions in very few sentences and gives the readers a clear overview of the problem – and that you solved it.

<u>Simon Peyton Jones</u>, in his excellent slides <u>How to write a</u> <u>great research paper</u>, gives a rule and example related to the abstract. The rule (referring to Kent Beck), is concisely stated in four sentences:

- 1. State the problem.
- 2. Say why it's an interesting problem.
- 3. Say what your solution achieves.
- 4. Say what follows from your solution.

An example of this rule:

- 1. Many papers are badly written and hard to understand.
- 2. This is a pity, because their good ideas may go unappreciated.
- 3. Following simple guidelines can dramatically improve the quality of your papers.
- Your work will be used more, and the feedback you get from others will, in turn, improve your research.

Restricting yourself to these four sentences is tough, and perhaps may seem even impossible. However, if you try it you'll see that you can improve your abstract significantly. Even if you end up with five sentences....! Surprisingly, a lot of abstracts I see mention only the first 3 points; some even only address point 3, like "In this paper we improve the SNR of noisy blurred images by 5%. By adding epsilon = 0.634 we even get to 5.34%".

Reason 1: No or insufficient contribution

I divide papers that get rejected for this reason into two classes. It is possible that many of the previous points will apply, and that there is no problem solved and, therefore, no contribution. However, sometimes people have learned how to write well (sometimes, such people are called politicians) and may be able to generate a paper without a contribution that lands on my desk.

After having gone through all the checkpoints and survived the hurdles, I finally decide that it is worth my time, and I get to read the paper. Now, I check if the claims are indeed true, if the contributions are significant enough to be accepted at the conference / in the journal. They

may have fooled me with the 'visuals': everything may look ok, but the body of the paper should provide the proofs and evidence. Here it may become clear that the theory and experiments do not make much sense or have already been done by others. It may also be that they are too limited, etc. The first two scenarios lead to a simple reject, but the latter case is trickier. The definition of "too limited" or "insufficient validation" is vague, differs per reviewer, and depends on the status of the journal / conference.

I hope that my list is useful to some of you. Just make sure that you avoid these pitfalls as much as possible. In my experience, most people only see these mistakes after their paper has been rejected. Before the submission deadline, they are so focused on getting the details (the body of the paper) right that they forget to check high level aspects (those mentioned here regarding introduction, conclusions, and abstract).

If you want to read more, my list is available in the pdf <u>10 ways to get</u> your paper rejected.



IAPR Newsletter, Vol. 38 No. 3, Jul. 2015

Page 3

Return to Page 1

Getting to know...Tamás Szirányi, IAPR Fellow

Prof. Tamás Szirányi received his Ph.D. and D. Sci. degrees in 1991 and 2001 from the Hungarian Academy of Sciences, Budapest. He was appointed to a full professor position in 2001 at the University of Pannonia, Veszprém, Hungary, and, in 2004, at the Péter Pázmány *Catholic University*, Budapest, Hungary. In 2014, he was appointed as a full professor at Budapest University of Technology and Economics. He leads the Distributed Events Analysis Research Laboratory at the

Institute for Computer Science and Control of the Hungarian Academy of Sciences. *His research activities include machine* perception, pattern recognition, *texture and motion segmentation, Markov random fields and stochastic* optimization, remote sensing, *surveillance systems for panoramic* and multiple camera systems, intelligent networked sensor systems, image quality, *digital film restoration, image rendering* and coding. Dr. Szirányi was the founder and past president (1997-2002) of the Hungarian Image Processing and Pattern *Recognition Society (the Hungarian member* society of the IAPR). He was an Associate Editor of the IEEE Transactions on Image Processing (2003-2009), and is currently an Associate Editor of Elsevier's Digital Signal Processing. He was honoured with the Master Professor Award in 2001 and the Széchenyi Professorship in 1999. He is a *Fellow of the IAPR and of the Hungarian Academy* of Engineering. He has more than 250 publications including 50 major scientific journals.

Tamás Szirányi, IAPR Fellow ICPR 2008, Tampa

For contributions to stochastic models and learning methods for video event analysis, and to energy optimization-based image and video segmentation.

by <u>Tamás Szirányi</u>, Institute for Computer Science and Control of the Hungarian Academy of Sciences and the Budapest University of Technology and Economics, Hungary

My involvement with image processing began in 1980 when I finished my MSc thesis about the mapping of thermal distribution on IC chips. Then, a method using linear programming and physical models helped us to estimate the image of thermal distribution with high precision. So, my first impression of image processing was the necessity of using tricky math to get useful results. Thanks to my supervisor, Prof. V. Székely, this impression determined my attitude towards image processing: math and pragmatism.

Later, I had special interests in recognizing shapes from the domain of not direct visibility by statistical methods: I introduced subpixel pattern recognition (Pattern Recognition, 1994), texture recognition with focal sensors (CVIU, 1998), depth maps from single images (IEEE Tr. PAMI, 2007), and finding hidden changes in remote sensing images (IEEE GRSL, 2014).

I started my career in industry, namely in the R/D department of an electrical company. Until 1991, I was a Research Fellow and Project Leader of Optoelectronics Development at the VIDEOTON Development Institute, Budapest. At that time, my work involved designing optical data storage, laser printers and document scanning devices, and other applications of laser systems. I led a research department of more than 20 people and organized several state-of-the-art projects.

During this period, in 1987, I visited Italian universities, including the Research Center "E. Piaggio" (Pisa, Prof. Paolo Dario), as a postdoc researcher. At that time, I had my first publication in Pattern Recognition Letters, about the recognition of very low resolution images, applied for robot touch detection.

Since new technologies needed

new experts, we started a doctorate program at the Technical University of Budapest, where about a dozen young researchers were trained with the partial supervision of the Videoton company. At that time, I was supervising PhD students, and we also had fruitful collaborations with university labs in optics and optoelectronics. We also had good cooperation, including licensing and development activities, with leading international companies like Philips and Thomson. Then, I was interested in hardware/firmware based image processing methods, dealing with the theoretical limits of OCR considering optical resolution and sampling rates. After the economic changes in Hungary around 1990, the Videoton company went out of business, and the whole institute scattered. I continued my career at the Institute for Computer Science and Control (SZTAKI) of the Hungarian Academy of Sciences.

In parallel, I continued teaching at universities, and was invited for professorship at different universities. In 1997, we founded the Faculty of Information Technology at the University of Veszprém (presently University of Pannonia), where I developed and taught several new subjects: image processing, pattern recognition, biometrics in personal identification. I also organized an image processing research laboratory and supervised a group of excellent PhD students. My first PhD student later became the chair holder professor.

This research laboratory had some successful R&D projects:

 DIMORF: software for the restoration of archive films. Our first result, the restoration of the first color Hungarian movie "Ludas Matyi", was featured in cinemas, released on DVD, and is still occasionally being shown on TV.

 Image compression and surveillance systems with panoramic cameras: cooperative projects with Samsung Co., and Tateyama Co. (Japan) led to industrial applications and patents.

Later, I handed the leadership of the laboratory over to an earlier PhD graduate of mine, and I moved my professorship to the newly founded Faculty of Information Technology at the Péter Pázmány Catholic University, Budapest. This faculty had special areas connected to life sciences and new technologies in physics. I gave lectures, and I had a staff consisting of PhD students, partially cooperating with my new SZTAKI laboratory.

Also in 1997, I founded the Hungarian Image Processing and Pattern Recognition Society (KEPAF), the Hungarian IAPR member society, and I was its first president for five years. With the formation of KEPAF, a real working society was created from a group of interested image processing and pattern recognition researchers. The society continues to grow and holds a national conferences every second year.

Between 1990-2002, my research activities were focused on focal plane smart imaging sensors. I have designed sophisticated analogue and digital circuits and related software packages. I also supervised young researches in the institute (SZTAKI) and at the university labs.

My former PhD students have earned several international scholarships and outstanding job positions. Their scientific performance fulfilled the highest

international requirements. In 2006, I founded a new research laboratory in SZTAKI, with a starting staff consisting of my selected previous PhD students. Our goal was and is the interpretation and organization of information coming from distributed or embedded sensors. The sensors can be dynamic or static imaging devices, other multimedia sources or a network of different sensors. The challenge is in the evaluation, recognition and classification of events occurring at different locations or times. We place special emphasis on machine learning, data mining, human perception, geometrical optics, optimization methods and variational analysis. Our areas of interests include image and video analysis, biometrical identification, connections with sensor networks, virtual reality and computer graphics. Our first major project was the development of a surveillance event analysis system for the Hungarian police.

In 2014 I was invited to the Budapest University of Technology and Economics, Faculty of Transportation Engineering and Vehicle Engineering and was appointed full professorship.

I have contributions to image and video analysis in

- Retrieving geometrical information from videos without any a priori information about the image structure or possible shapes:
 - registration through co-motion statistics for matching image views, resulting in multimodal camera fusion or finding the vanishing point for planar reflected images or shadows,
 - o focus maps through

Bayesian iterations by using a theoretically new orthogonality criterion;

- Segmentation, change and motion detection by using novel effective Bayesian methods;
- Restoration of archive films by self-learning adaptive recognition methods.

I also have contributions in sensor systems and engineering solutions:

- Developing intelligent focal plane sensor toolboxes, designing sophisticated analogue and digital circuits and related software packages

 used by several R&D groups worldwide;
- Measuring the visual quality and human perception in human-computer interfaces.

We have several international cooperations, including French INRIA institutes, Italian, Turkish, Dutch and German universities and research groups. We also have active cooperation with researchers in machine learning, cognitive sciences and mathematics. Our group and my colleagues have been honoured by several prizes in recent years, and we achieved a high scientific level (number and quality of publications) in solving engineering tasks.

By now we have collected so many ideas about sensor systems, surveillance, and aerial imaging, that I can see several new areas of innovation, the more important ones being:

- Modeling high complexity systems containing an indefinite number of (freely positioned) networked sensors.
- Obtaining additional information from sensor networks, characterizing their geometrical or statistical relations and forming more

effective collaborations.

I started my career in industry and then continued into academia. Just as during my years in industry, I continue to pursue scientific achievement. Currently I strive to produce scientific results with industrial applicability. Such duality continues to drive my research work: my aim is to bring together new theoretical issues in creating a new thinking of networked equipment.

It is my conviction that a system containing a large number of sensors can provide much more information and knowledge than the simple summation of single sensors. During my earlier career in both industry and academia my favorite activities were around and about intelligent sensor systems. Today, it is clear that the time of smart, self-organized sensor systems has arrived, and I keep pursuing answers and solutions on how to face the related challenges.

IAPR Then and Now...IAPR Newsletter, Vol. 6 Nos. 2 and 3, October 1983

Excerpt from the Minutes of hte Ninth Meeting of the Executive Commitee of the IAPR

Hotel Eremitage, Lyngby (DK) July 14, 1983

Present were: Simon (President), Kohonen (Vice-President), Freeman (Treasurer), Devijver (Secretary), Levine (Chm. 7ICPR), Bajcsy (Newsletter editor), Danielsson (Chm. Membership committee), Rutovitx (Chm. Awards comm.).

Membership

• Danielsson reports that an application for membership was received from Hungary in October 1982, but the information provided at that time about the Hungarian national organization did not match the requirements of the IAPR Constitution. Additional information was received in July 1983 which presently allows the membership committee to propose that the Hungarian organization be accepted by the Governing Board as the twentieth national member of the IAPR.

The Hungarian national organization is the Section of Artificial Intelligence and Pattern Recognition of the John Von Neumann Society for Computer Science with a membership of 25 (category A). The proposed Hungarian national representative is Dr. G. Kozman, Center Research Institute for Physics, POB 42, Budapest H.1525.

• The IAPR secretary is directed to organize a mail ballot on the Hungarian application at his earlier convenience.

Follow-Up To the Minutes of the ninth meeting of the Executive Committee of IAPR

• A mail ballot dated September 5, 1983 has gone out to all members o fthe IAPR Governing board concerning the application for membership in the IAPR from the "Section of Artficial Intelligence and Pattern Recognition of the John Von Neuman Society for Computer Science, Hungary"

IAPR Then and Now...IAPR Newsletter, Vol. 20 No. 1, January 1998

New GB member for Hungary The new Hungarian representative on the GB of IAPR is Dr. Tamás Szirányi, the newly elected President of the Hungarian Association for Image Analysis and Pattern Recognition (KEPAF). Dr. Szirányi replaces Professor Dmitry Chetverikov who served on the GB since the early 80's. Dr. Szirányi is Senior Research Fellow at MTA SZTAKI, Budapest and Associate Professor at the University of Veszprem.

APR...The Next Generation In this series of Feature Articles, the IAPR Newsletter asks young researchers to respond to three questions. • Briefly: How did you get involved in pattern recognition and what technical work have you done? In more detail: What is/are your current research interest(s)? • How can the IAPR help young researchers? ~Arjan Kuijper, Editor -in-Chief

Yan Yan



by <u>Yan Yan</u>, Department of Information Engineering and Computer Science, University of Trento, Italy

Briefly: How did you get involved in pattern recognition and what technical work have you done?

I remember that when I was a third-year undergraduate student, several years before, and I took a course called 'Digital image Processing'. I found that it was really magic that we can use some methods to track people in the videos captured from cameras. That was the first time I touched pattern recognition. Inspired by this, I chose computer vision and pattern recognition as my major for my later Masters and Ph.D. studies.

My Ph.D. study mainly focused on multi-task learning and transfer learning with applications to

Yan Yan received his Ph.D. degree in *computer science from the University* of Trento, Italy. He is currently a Post-Doctoral Researcher with the MHUG Group, University of Trento. His research interests include machine learning and its application to computer vision and multimedia analysis, e.g. multi-view head pose estimation, multi-view action recognition, multimedia retrieval, image annotation, etc. He is the author of *refereed journal and conference papers* in IEEE TIP, IEEE TNNLS, CVIU, CVPR, ICCV, ACM Multimedia, AAAI, IJCAI, etc. Dr. Yan is the recipient of Best Student Paper Award in ICPR 2014. He was a Visiting Scholar with Carnegie Mellon University in 2013. He has served as a PC member for ACM Multimedia 2013-2015 and as a reviewer for refereed journals and conferences.

Editor's note:

Yan Yan was the recipient of the IBM Best Student Paper Award at ICPR 2014 in Track 4: Document Analysis, Biometrics and Pattern Recognition Applications for the paper "Clustered Multi-Task Linear Discriminant Analysis for View Invariant Color-Depth Action Recognition" authors: Yan Yan, Elisa Ricci, Gaowen Liu, Ramanathan Subramanian and Nicu Sebe.

~ Arjan Kuijper, Editor-in-Chief

different pattern recognition problems, in particular, head pose estimation from multiple, large field-of-view surveillance cameras and multi-view human action recognition.

For the head pose estimation problem, a novel graph-guided multi-task learning framework was proposed for classifying head pose of moving targets from multiple camera views. Starting from a dense 2D spatial grid, two graphs which respectively model appearance similarity among grid partitions and head pose classes guide the learner to output regionspecific pose classifiers and the optimal space partitioning.

For the action recognition problem, Self-Similarity Matrices (SSMs) have been found to be effective view-invariant action descriptors. To enhance the performance of SSM-based methods, we proposed Multi-task LDA, a novel multi-task learning framework for multi-view action recognition that allows for the sharing of discriminative SSM features among different views (i.e. tasks). Inspired by the mathematical connection between multi-variate linear regression and Linear Discriminant Analysis (LDA), we modeled multi-task multi-class LDA as a single optimization problem by choosing an appropriate class indicator matrix. This work was presented at ICPR2014 in the paper "Clustered Multi-Task Linear Discriminant Analysis for View Invariant Color-Depth Action Recognition". The figure shows output from the method against images from the ACT4² dataset.

In more detail: What is/are your current research interest(s)?

My current research interests mainly focus on ego-centric daily activity recognition and Multimedia Event Detection.



Recognizing human activities from videos is a fundamental research problem in computer vision. Recently, there has been a growing interest in analyzing human behavior from data collected with wearable cameras. First-person cameras continuously record several hours of their wearers' lives. To cope with this vast amount of unlabeled and heterogeneous data, novel algorithmic solutions are required. In my current research, we propose a multi-task clustering framework for activity of daily living analysis from visual data gathered from wearable cameras. Our intuition is that, even if the data are not annotated, it is possible to exploit the fact that the tasks of recognizing everyday activities of multiple individuals are related, since typically people perform the same actions in similar environments (e.g. people working in an office often read and write documents). In our framework, rather than clustering data from different users separately, we propose to look for clustering partitions which are coherent among related tasks. Specifically, two novel multi-task clustering algorithms, derived from a common optimization problem, are

introduced.

Multimedia Event Detection is a retrieval task with the goal of finding videos of a particular event in a large-scale unconstrained internet video archive, given example videos and text descriptions. Nowadays, different multimodal fusion schemes of low-level and high-level features are extensively investigated and evaluated for the complex event detection task. However, how to effectively select the high-level semantic meaningful concepts from a large pool to assist complex event detection is rarely studied in the literature. In my current research, we propose a novel strategy to automatically select semantic meaningful concepts for the event detection task based on both the events-kit text descriptions and the concepts high-level feature descriptions. Moreover, we introduce a novel event oriented dictionary representation based on the selected semantic concepts. Toward this goal, we leverage training images (frames) of selected concepts from the semantic indexing dataset with a pool of 346 concepts, into a novel supervised multitask Lp-norm dictionary learning framework.

How can the IAPR help young researchers?

The IAPR is a famous organization in the area of pattern recognition. It is important for young researchers to have their own funding to support and continue their research. Therefore, it would be really useful for young researchers if the IAPR can provide some information about funding opportunities in different countries. Moreover, it would be good if the IAPR can provide more travel funding for students and young researchers to attend IAPR international conferences.

Page 8

From the



News from the Executive Committee of the IAPR

by <u>Alexandra Branzan-Albu</u> IAPR Secretary





ICDAR Workshops (* denotes a non-IAPR event)

- 6th International Workshop on Camera Based Document Analysis and Recognition (CBDAR 2015)
- * 5th International Workshop on Multilingual OCR (MOCR 2015)
- * 4th International Workshop on Automated Forensic Handwriting Analysis (AFHA 2015)
- 3rd International Workshop on Historical Document Imaging and Processing (HIP 2015)
- 11th International Workshop on Graphics Recognition (GREC 2015)

Victoria, July 12 2015

Summer is usually great for conference travel, mainly because not many academics teach at this time of the year, so they have more time for their research and for connecting with their peers at conferences. There is a lot of effort that goes into the organization of an IAPR conference, and most of the work is shouldered by the local organizers and the advisory committee. There are times when these committees have to make very difficult decisions under considerable time and financial constraints, such as in the case of the 13th International Conference on Document Analysis and Recognition, ICDAR 2015 and its five related workshops. The recent terrorist attack in Tunisia rendered the original venue of the conference potentially unsafe, so the conference venue was moved to Nancy (France), while the dates remain the same (23-26 August 2015). The IAPR Executive Committee fully supports the Tunisian IAPR Society, which is in charge of organizing ICDAR 2015. We applaud their success in finding an excellent solution to a complex problem under considerable time constraints. We also extend our thanks to the University of Lorraine and to the LORIA research laboratory for their local support.

We received sad news from the Italian Association for Pattern Recognition (GIRPR) about the loss of one of their most prominent members, Professor and IAPR Fellow Stefano Levialdi. This issue of our Newsletter contains a heartfelt obituary written by his colleague and friend Virginio Cantoni.

The Interim ExCo Meeting will be hosted by the IAPR Past President Prof. Kim Boyer at the new <u>College of Engineering and</u> <u>Applied Sciences at the University of Albany, NY</u>, on August 18 and 19 2015. This year, Ms Linda O'Gorman (IAPR Secretariat and IAPR Newsletter Layout Editor) will attend part of the meeting. We are convinced that her presence will be very helpful, and that her participation will make our discussions more informed and productive.

I hope that you will find the July 2015 issue of the *IAPR Newsletter* an inspiring read.

Page 9

In Memoriam...

Stefano Levialdi IAPR Fellow and IEEE Life Fellow (1936-2012)



by <u>Virginio Cantoni</u>, IAPR and IEEE Fellow, University of Pavia, Italy

Stefano Levialdi died on July 1, 2015, in Rome, after a final, long and uneven battle, and after a series of incredible victories over diseases, often called incurable. With Stefano's passing, the IAPR and the communities of Visual Languages (VL) and Human Computer Interaction (HCI) have lost an inspiring and valuable protagonist who leaves a lasting impression on the national and international university context.

Stefano graduated at the end of the 1950's in Telecommunication Engineering in Buenos Aires. In 1960, he came to Italy, joining a group at the University of Genoa that was carrying out research into Pattern Recognition, developing a machine, which after a short training period, could learn a given class of patterns. In 1961, he was appointed Lecturer in Electronics in the same university and in 1965, joined the Italian National Research Council (CNR) in Naples where, for 13 years, he headed a research group working on Parallel Image Processing. In 1981, he became Full Professor of Computer Science at the University of Bari, moving to La Sapienza in Rome in 1983. Here, he drew up a new curriculum in Computer Science, which saw the light of day in 1989 after six hard years of discussion. He finally retired in November 2010, after working for 49 years in different universities.

Stefano's university teaching duties included a wide range of subjects: Electronics Laboratory, Computer Architectures, Image Processing and Human-Computer Interaction, all aimed at Computer Science students. As regards publications, Stefano wrote over 250 papers. His activity for the Journal of Visual Languages and Computing (Elsevier Press), which began in 1990, deserves a special mention. He began this in 1990 (together with Prof. S.K. Chang). Moreover, Stefano travelled widely, he visited over 45 countries where he gave talks and started collaborations with local researchers. He spoke English, French, Italian and Spanish fluently.

For his work on Parallel Image Processing, he was awarded an IEEE Fellowship in 1988 (and IEEE Life Fellowship in 2001); for his overall work in Image Processing, he became an IAPR Fellow in 1994; in 2008 for his activity in Human/Computer Interaction and Visual Languages, he became a VL/HCI Fellow.

Our friendship was soon full and firm, probably because we were, without knowing it, complementary, but certainly, because we shared the same ideals about university and university research. From the beginning, Stefano gave me a lot of his time and energy, so generously. We started working together in the late 1970s, more than 35 years ago, during the preparation for the first International Conference on Image Analysis and Processing (ICIAP, 1980); he served as Chairman of the Scientific Committee, while I was local chair, looking after the proceedings. Over the years, we have organized many initiatives together: in particular, and with others, the first five editions of ICIAP as well as a NATO Advanced Research Workshop on Pyramidal Systems for Computer Vision. In 1988, we started the first school on 'La Visione delle Macchine' which, with contributions from all units of GIRPR, led to the publication of a volume which

was adopted by the majority of Italian university courses on image processing and pattern recognition. I want to mention two other joint projects given their intensive nature over a long period. The first was PAPIA, an Italian national research project that involved seven units (universities and the Italian National Research Council). Throughout the 1980s, it investigated the simulation and construction of a multiprocessor system with a pyramidal architecture. The second lasted for even longer. 20 years in fact. This was the writing of the book with our co-author Betrand Zavidovique 3C Vision: cues, contexts and channels. It began in the 1990s and ended in 2012.

His enthusiasm for science and technology research, an expression of his innate curiosity, but richly nourished by culture, ensured he overcame the strenuous initial difficulties of academia. With his total commitment, he managed to fulfil his life in a way that was remarkably consistent and valid. His passion for research was so addictive that he never allowed himself to be satisfied with what he had already achieved, even though this would have been the easiest thing to do.

Whoever approached him had the immediate perception of a person who was intensely free and fair. Stefano avoided makeshift solutions, however convenient. when the values in which he believed were likely to suffer, and coherent to the end, he would hold the line at all costs, fighting alongside friends in the defense of a different, and almost always new and effective, line of thinking. His research led to publications of high scientific value, written in a brilliant style, which, at the very same time, were concrete and precise.

Countless times we have seen Stefano at a conference standing up to give his opinion. He always knew how to adjust perfectly to the audience, choosing words and a tone of voice that were never inappropriate: always calm, even when his ideas were strong and when the difference in opinion was a complete rift; he was always trying to persuade, even when it seemed most improbable that the other party would concede.

Anybody who came close to him in recent years could only wonder at his physical, as well as mental, resistance. He fought with so much courage to lead a life, which, at end, was very hard. While his dramatic physical condition was very evident, his extraordinary spirit amazed us all. We all wondered how it was that Stefano had lost none of his traditional lustre, in particular, when we realized that, even though so bruised, he was planning new activities and developing new proposals. Beyond the substantial scientific and academic merit. this is the Stefano Levialdi we knew and that we have lost forever.

IAPR Then and Now...*IAPR Newsletter* 1980-1993

When the IAPR Newsletter editorial staff learned this sad news, we began to search through the Newsletter archives for references to Prof. Levialdi. We found too many to reprint here, so, instead, we share our overall impression. Prof. Levialdi was very active within the IAPR from 1980 to 1993. Through the Newsletter, we see him as someone who was always willing to help, share knowledge and take on extra responsibility. He was steadily involved in the IAPR community through conference organization and through participation in the IAPR Governing Board, its Standing Committees and its Executive Committee. We invite you to reread "Getting to Know...Stefano Levialdi, IAPR Fellow" in the April 2012 issue of the IAPR Newsletter.

~Arjan Kuijper, Editor-in-Chief, and Linda O'Gorman, Layout Editor

Meeting Reports Conferences, Workshops & Summer Schools



CCIW 2015

5th IAPR Computational Color Imaging Workshop March 24-26, 2015 - Saint-Etienne, France

www.cciw2015.org/

Workshop Chairs:

<u>Alain Trémeau</u> (University Jean Monnet, France) <u>Raimondo Schettini</u> (University of Milano Bicocca, Italy) <u>Shoji Tominaga</u> (Chiba University, Japan)

by the Workshop Organizers/ Chairs

The 5th Computational Color Imaging Workshop (CCIW'2015) had six invited talks and 17 oral presentations (a 56% acceptance rate). Progress and challenges in various areas of computational color imaging were discussed over the course of the three-day workshop.

The workshop began on March 24, 2015, with a keynote presentation by Mathieu Hébert (University Jean Monnet, France) called "Color and spectral mixings in printed surfaces". Next, in the **Color reproduction** session, Dmitry Kuzovkin presented an automatic method for descriptor-

based image colorization and a robust color artifact regularization method. This paper proposed new improvements for color mapping which were not covered by the recent survey on color mapping written by H. S. Faridul et al. in 2014. In the next presentation, G. M. Atigur Rahaman proposed to extend the Murray–Davies reflectance model used for modelling spectral halftone images and to improve the efficiency prediction of this model to changes in reflectance by a power function. Lastly, Ryosuke Nakahata (Chiba University, Japan) presented a dynamic relighting method for moving planar objects with unknown reflectance. By acquiring the surface spectral reflectance



Click above to go the publisher's web site for this volume.

of moving objects, this method is able to reproduce accurate colors on a display device. This research topic is still largely unexplored. Most of the solutions using a projector-camera system published recently in the state of the art address issues related to photometric compensation but not to color compensation. Moreover, most of these papers deal with static surfaces and not with moving surfaces. However, a wide range of applications such as augmented reality, education, cultural heritage and interactive art installations could benefit from progress in this field.

Another challenging issue in computational color imaging is related to the color rendering of color reproduction. Some researchers tried to improve this issue by improving first the color acquisition. In his invited talk, titled "The good, the bad and the ugly: the color we would like, the color we have, its appearance and dynamic range", Alessandro Rizzi (University of Milano, Italy) discussed some "hidden" issues, often not taken into account, related to color acquisition that can introduce severe errors in the color information. Among these issues he focused on the limits of accurate camera acquisition, the usable range of light of our vision system, and the role of accurate vs. non-accurate luminance recording for the final appearance of a scene.

In the **Color sensation and perception** session, G. M. Atiqur Rahaman presented issues related to the acquisition and analysis of memory colors of objects found in natural scenes. In the next presentation, João M. M. Linhares (University of Minho, Portugal) discussed the effect on a display gamut of varying the optical density and the position of the maximum sensitivity of the cones spectra of anomalous trichromatic observers. Lastly, Jorge L. A. Santos (University of Beira Interior, Portugal) discussed reaction times for normal color and dichromatic observers in a visual search experiment. These three papers addressed a wide range of problems related to vision science (i.e. to sensation and perception).

The second day of the workshop began with an invited presentation by Joost Van de Weijer (University of Barcelona, Spain) on color features in the era of big data. The process of unsupervised feature learning (e.g. deep learning) has recently received a lot of attention in the field of computer vision. Meanwhile many studies were carried out only on color data, since recently few studies have been carried out on multimodal data (e.g. RGB and depth data or RGD and temporal data). However, several papers have demonstrated that machine learning plays an important role in bridging the gap between feature representations and decision making (e.g., for object/scene recognition, human pose estimation, and gesture/activity recognition) by learning useful information from a large set of RGB-D data. In his presentation Joost Van de Weijer provided an overview of color names applications in computer vision, including image classification, object recognition, texture classification, visual tracking and action recognition, and demonstrated that, in general, color names outperform photometric invariants.

In the **Color image processing** session, Pablo Martínez-Cañada (University of Granada, Spain) presented a configurable simulation platform that reproduces the analog neural behaviour of different models of the human visual system at the early stages. Next, Yann Gavet (EMSTE, France) proposed to use color logarithmic image processing (CoLIP) for white balance correction and color transfer. Gianluigi Ciocca (University of Milano, Italy) discussed the influence of color on the perception of image complexity. These three papers addressed a wide range of color image processing solutions related to vision science. Lastly, Andreas Kleefeld (BTU Cottbus-Senftenberg, Germany) presented a new framework for color-valued median filters.

Another challenging issue in computational color imaging is related to the processing of spectral imaging data. In his invited talk, titled "Optics and computational methods for hybrid resolution spectral imaging", Masahiro Yamaguchi (Tokyo Institute of Technology, Japan) introduced the concept of "Hybrid **Resolution Spectral Imaging**" (HRSI) and presented algorithms for reconstructing spectral images. The goal is to combine a highresolution RGB image and a low-resolution spectral image in order to capture high resolution spectral video with a compact and handy camera system. Thanks to hybrid spectral imaging systems, it becomes possible to think of new applications and new developments of spectral imaging.

In the **Spectral imaging** session, Simone Bianco (University of Milano, Italy) discussed if the performance of hyperspectral face recognition algorithms can be improved by considering 1D projections of the whole spectral data along the spectral dimension. Feature band selection, dimensionality reduction and feature extraction are challenging issues for the face recognition task and also for other computer

vision tasks. Even if, in the last ten years, many studies tried to solve these issues in the field of face recognition, several improvements remain to be overcome. Similar issues/improvements remain to be overcome in other application fields. In the next presentation, Hilda Deborah presented a spectral noise model using a spectral database of uniform color/ pigment patches, which answers the challenge of identifying a spectral noise model. Lastly, Xingbo Wang discussed the colorimetric performance of CFA/ MSFA based image acquisition systems. Despite the number of studies dealing with spectral imaging, little attention has been given to the evaluation of the quality of spectral images and of spectral imaging systems. However, we can note that there is a general tendency to address this issue from several different perspectives (physics and psychophysics).

On the last day of the workshop, in cooperation with the European COST action TD 1201 (http:// cosch.info/), a special session on color in digital cultural heritage was scheduled. The process of color digitization of 3D objects in cultural heritage has recently received much attention due to continuing improvements in the quality and resolution of digital objects. One tendency to improve color accuracy is to use a multispectral system. Thus, Ari Ide-Ektessabi (University of Kyoto, Japan) proposed to use high-resolution multispectral scanning for mesoscopic investigation of discoloration of traditional Japanese pigments. Another tendency is to improve the performance of colordifference formulas or to evaluate with these formulas if a color digitization system is accurate

enough. Thus, Tatiana Vitorino (Consiglio Nazionale delle Ricerche, Florence, Italy) proposed using the ColorChecker chart to assess the usefulness and comparability of data acquired with two hyperspectral systems. The process of 3D object visualization in cultural heritage has also received much attention due to the development of color rendering and color correction algorithms. Some authors studied these issues from the observer's perspective (i.e. visual observation). Sergio Nascimento (University of Minho, Portugal) discussed which color compositions observers prefer when they look at some paintings. Other researchers proposed to address these issues using photometry/spectrophotometry models. For example, Lindsay MacDonald (University College of London, UK) in his invited talk titled "Representation of cultural objects by image sets with directional illumination" discussed problems related to the modelling of the diffuse and

specular reflectance of 3D objects and to 3D surface reconstruction from photometric stereo. Another approach is to address these issues using computer vision models. Zoltan Kato (University of Szeged, Hungary) in his invited talk titled "Relative pose estimation and fusion of 2D spectral and 3D lidar images" discussed problems related to pose estimation without the use of any special calibration pattern or explicit point correspondence. This paper addressed one of the most challenging issues in digital cultural heritage which is the fusion of 2D RGB/spectral imagery with other 3D range sensing modalities (e.g. Lidar). On the other hand, Citlalli Gamez Serna presented a semi-automatic 2D-3D registration framework to produce accurate realistic results from a set of 2D uncalibrated images and a sparse 3D point cloud representation of an object digitized with a laser scanning.

CCIW 2017 will take place in Milano, Italy, in Spring 2017.

The image below is from Citlalli Gamez's talk on "Data fusion of objects using techniques such as Laser Scanning, Structured Light and Photogrammetry for Cultural Heritage Applications"



General Chairs: <u>Cheng-Lin Liu</u> (Institute of Automation of CAS, China) Bin Luo (Anhui University, China) <u>Walter G. Kropatsch</u> (Vienna University of Technology, Austria)

by the General Chairs

The 10th IAPR Workshop on Graph-based Representations in Pattern Recognition (GbR) was sponsored by the International Association for Pattern Recognition (IAPR) and hosted by the National Laboratory of Pattern Recognition of CASIA.

GbR is a workshop organized by the IAPR's Technical Committee 15 on Graph-Based Representation

(IAPR TC-15), aimed at encouraging research works in pattern recognition and image analysis within the graph theory framework. This workshop series, held every two years, traditionally provides a forum for presenting and discussing research results and applications in the intersection of pattern recognition, image analysis and graph theory. Recently, graph-based structural models and probabilistic graphical models have been widely applied to pattern recognition, computer vision, machine learning and data mining.

The GbR2015 is the tenth in the series and is the first to be held outside of Europe. It received 53 submissions, of which 36 were accepted for presentation after peer review. 33 program committee members and 12 additional reviewers were invited to review submissions. Each submission was assigned to three reviewers. The accepted papers were presented in nine oral sessions, and each paper was given 25 minutes for presentation and discussion. There were also two IAPR invited talks. Professor Marcello Pelillo, University of Venice University, Italy, addressed a speech titled "Revealing Structure in Large Graphs: Szemerédi's Regularity Lemma and Its Use in Pattern Recognition". Dr. Xing Xie, Microsoft Research Asia, Beijing, China, gave a talk titled "Understanding Users by Connecting Large Scale Social Graphs".

More than 50 participants attended the workshop, including 20 students. The organizing committee offered stipends to six students from Europe, including four Class A stipends that waived the registration fee and offered accommodations, and two Class B stipends that offered accommodation only.

The proceedings of GbR2015 were published by Springer in Lecture Notes in Computer Sciences, Vol. 9069, <u>http://link.springer.com/book/</u> <u>10.1007%2F978-3-319-18224-7</u>.

Click above to go the publisher's web site for this volume.

Locations of previous editions of GbR:

Gbr1997, France GbR1999, Austria GbR2001, Italy GbR2003, UK GbR2005, France GbR2007, Spain GbR 2009, Italy GbR2011, Germany GbR2013, Austria GbR2015, China

ICB-2015 The 8th IAPR International Conference on Biometrics May 19-22, 2015. Phuket, Thailand http://icb2015.org/

General Co-Chairs: <u>Massimo Tistarelli</u>, University of Sassari, Italy <u>Wei-Yun Yau</u>, Institute for Infocomm Research (I²R), Singapore <u>Vutipong Areekul</u>, Kasetsart University, Thailand

by the General Chairs

ICB 2015 was held at the Merlin Beach Resort which is located at a secluded and serene section of Phuket that is famed for its world class beaches and beautiful scenery. This biometric flagship conference was co-sponsored by the IAPR Technical Committee 4 on Biometrics (IAPR TC-4) and the IEEE Biometrics Council and hosted by Kasetsart University. The conference was attended by 165 delegates from 26 different countries.

To maintain the high standards of the conference, the Program Co-Chairs, Prof. Vijayakumar Bhagavatula (Carnegie Mellon University, USA), Prof. Mark Nixon (University of Southampton, UK), and Prof. Stan Z. Li (Chinese Academy of Sciences, China) led 25 Area Chairs and 176 reviewers to review all the papers received. This year the conference received 169 paper submissions from 35 different countries. After a rigorous review process that involved 560 reviews, the program committee selected 76 papers from 24 countries: 31 papers were presented in 9 oral sessions while the remaining 45 papers were presented in 3 poster sessions. The poster presenters had the opportunity to provide a

two-minute oral overview of their papers to the audience prior to the poster session. For this year, the poster session lasted the whole day allowing for more interaction among the researchers. The oral sessions were classified into the following areas, representing a wide range of biometric topics: iris and ocular, face, fingerprint, hand and vein, other biometrics, mobile and remote, and security and spoofing. Moreover, two oral sessions were allocated for the best reviewed papers.

ICB-2015 had three out-standing keynote speakers: (1) Dr. Deepak Chandra (Google's Advanced Technology and Projects (ATAP), USA) who gave a keynote entitled "Multi Modal Smartphone Authentication"; (2) Dr. Sébastien Marcel (IDIAP Research Institute, Switzerland) who introduced a new tool for the biometric research community in his talk entitled "Reproducible research in Biometrics: Moving to the BEAT"; and Prof. Shuicheng Yan (National University of Singapore) who centered his presentation on the topic: "Deep Learning and Biometrics".

ICB-2015 featured <u>four tutorials</u> given by outstanding lecturers in the field: (1) Prof. Adams Wai-Kin Kong (Nanyang Technological University, Singapore) on the topic "An Overview of the Recent **Development of Computational** Forensics for Criminal and Victim Identification"; (2) Prof. Brian Lovell (University of Queensland, Australia) on the topic "Remote, Mobile, and Wearable Video Face Recognition for Surveillance"; (3) Dr. Sébastien Marcel (IDIAP research institute, Switzerland) on the topic "Spoofing and Anti-Spoofing in Biometrics: Lessons Learned from the TABULA RASA Project"; and (4) Prof. Wolf Lior (Tel Aviv University, Israel) on the topic "Deep Learning Methods for Face Recognition". Each of the tutorials was attended by several participants, and there were lively discussions between the lecturers and the participants.

All ICB-2015 conference details, images, videos of keynote speaker talks and the keynote slides can be accessed at <u>http://www.icb2015.org</u>

> ICB 2015 Proceedings are available through <u>IEEEXplore</u>

Two <u>panel sessions</u> were organized.

The first panel session on "Biometric Technology Adoption and Challenges for Seamless Border Control in ASEAN", was moderated by Dr. Wei-Yun Yau and Prof. Pratit Santiprabhob (Assumption University, Thailand) and featured Prof. Vutipong Areekul, Mr. Azman azra bin Abdul (Ministry of Home Affairs, Malaysia), Mr. Chittipat Tongprasroeth (Ministry of Foreign Affairs, Thailand), and Mr. Peerapol Promwong (Chanwanich Security Printing Co.Itd., Thailand). The ASEAN Economy Community (AEC2015) will be started at the end of this year. Therefore, the aim of this panel was to highlight some of the research challenges faced for border control in ASEAN that could motivate the R&D community.

The second panel session, entitled "Promoting High Quality Biometrics Research", was moderated by Prof. Anil Jain (Michigan State University) and featured Dr. Deepak Chandra, Prof. Josef Kittler (University of Surrey, UK), Dr. Brendan Klare (Noblis, USA), and Prof. Stan Z. Li (Chinese Academy of Sciences, China).

The <u>Doctoral Consortium</u> provided a forum for young PhD researchers to interact with senior researchers in the field by presenting their research work in a poster session and discussing during the lunch time. Prof. Arun Ross (Michigan State University, USA) led a team of mentors comprising Prof. Mark Nixon (University of Southampton, UK), Prof. Josef Kittler (University of Surrey, UK), Dr. Brendan Klare, and Prof. Julian Fierrez (Universidad Autonoma de Madrid, Spain).

A number of <u>awards</u> were presented at ICB-2015.

The IAPR Young Biometrics Investigator Award (YBIA) was given to Prof. Gang Hua from Stevens Institute of Technology, USA, who presented a plenary talk on "Unconstrained Face Recognition in Images and Videos".

Gang Hua (r) receives his award from Arun Ross

- The IAPR Best Biometrics Student Paper Award (BBSPA) was given to Jianqing Zhu, Shengcai Liao, Dong Yi, Zhen Lei, and Stan Z. Li from the Chinese Academy of Sciences, China, for the paper entitled "Multi-label CNN-Based Pedestrian Attribute Learning for Soft Biometrics".
- The Siew-Sngiem Best Paper Award was given to

Marta Gomez-Barrero, Javier Galbally, Julian Fierrez, Javier Ortega-Garcia, and Réjean Plamond from the Universidad Autonoma de Madrid, Spain and the École Polytechnique de Montreal, Canada, for the paper entitled "Enhanced On-line Signature Verification Based on Skilled Forgery Detection Using Sigma-log Normal Features".

The Siew-Sngiem Best Poster Award was given to Christian Rathgeb, Harald Baier, Christoph Busch, and Frank Breitinger, from the Hochschule Darmstadt, Germany and the University of New Haven, USA for the poster entitled "Towards Bloom Filter-based Indexing of Iris Biometric Data".

The Siew-Sngiem Award was the Platinum sponsor of ICB-2015. The Gold sponsors were Kasetsart University and Chanwanich co.ltd. The Silver sponsors were Safran-Morpho and ST Electronics. To promote awareness of the stateof-the-art biometric works, several Thailand government officers were invited to attend and supported by the Ministry of Information and Communication Technology, Ministry of Justice, Ministry of Foreign Affair, Ministry of Labor, and Immigration Bureau, Thailand.

<u>The 9th International Conference</u> on Biometrics (ICB-2016) will be held in Halmstad, Sweden, on June 14-17, 2016.

Page 17

General Chairmanship: <u>Rasmus Larsen</u> (DTU Compute, Denmark) <u>Mads Nielsen</u> (University of Copenhagen, Denmark) <u>Dan Witzner Hansen</u> (IT University of Copenhagen, Denmark)

Program Chairs: <u>Rasmus R. Paulsen</u> (DTU Compute, Denmark) <u>Kim Steenstrup Petersen</u> (University of Copenhagen, Denmark)

by Rasmus Larsen, Chairman of the Danish Pattern Recognition Society

SCIA is one the oldest conference series on image analysis in the world.

SCIA 2015 was organized by the IT University of Copenhagen, the University of Copenhagen, and the Technical University of Denmark on behalf of the Danish Pattern Recognition Society. This biennial conference series is organized in turns by the Norwegian, Swedish, Finnish, and Danish pattern recognition societies under the auspices of the International Association of Pattern Recognition. The current President of the IAPR **Executive Committee Ingela** Nyström is from the Swedish Pattern Recognition Society, and she was with us at SCIA 2015. The IAPR sponsored SCIA 2015 and made it possible to invite two distinguished IAPR speakers.

At SCIA, we bring together scientists from all the Scandinavian

university image analysis and computer vision research groups as well as scientists outside Scandinavia. Although we are all internationally oriented and meet each other at conferences and summer school all over the world, it is also of great value for us to have a closer contact. Geography and proximity still matter. We are four small countries more or less sharing language and culture. And, even though we are small countries, it is fair to say that in the field of science we are concerned with here we are punching above our weight.

The field of study – image analysis and computer vision – is not esoteric or dwindling. It is very much in focus across a long series of application areas including smart manufacturing, health care and wellbeing, and secure societies. Although image analysis is not a dominant discussion in popular debate, it pops up in many contexts: in discussions of privacy in societies where we are monitored constantly by surveillance cameras; in discussions of smart manufacturing and automated production; in the global battle for production jobs; and as show cases from the universities. In Sweden, a recent donation of 1.8 BSEK (250 MEuro) by the Wallenberg foundation to the Wallenberg Autonomous Systems Program will involve much computer vision. In Lund,

web site for this volume.

Sweden the European Spallation Source with major investments from Sweden and Denmark and 15 other partner countries is currently being constructed. This will be the most power full neutron source in the world. At the same time the MAX IV synchrotron is being built also in Lund. This will be the most brilliant x-ray source in the world. Among other experiments, these large scale facilities will produce images of the insides of materials. Our competencies are in demand for analyzing these.

In Denmark, a recent investment was made by the A. P. Moller Foundation for the establishment of the Danish Hydrocarbon Research Center. The DHCRC will draw upon spatial data analysis and image analysis over many scales. They range from the entire North Sea to the micro structure of the chalk in the oil reservoirs below the sea bed.

At the conference two IAPR sponsored speakers Professor Tim Cootes (Manchester University, UK) and Research Director Dorin Comaniciu (Siemens Corporate Research, New Jersey, USA) gave keynote addresses. Dr. Comaniciu spoke on *"Shaping the Future through* Innovations: Towards Imagebased Personalized Medicine" and professor Cootes on "Image Segmentation using Statistical Shape Models". Moreover, a series of five invited speakers for the Scandinavian countries gave additional keynote addresses: Professor Carolina Wählby (Sweden), Associate Professor Robert Jenssen (Norway), Thomas B. Moeslund (Denmark), Kalle Åström (Sweden), Christoph Busch (Norway).

Furthermore, the programme consisted of 45 contributed papers. 26 were presented as oral presentations and 19 as poster presentations. The conference had 116 participants. The programme also included an industrial exhibition in which seven Danish computer vision companies - that also were sponsors of the conference - exhibited their products. These included Fingerprints – a biometric solution provider, Videometer – an hyper spectral imaging instrument and solution provider, Innospexion – an X-ray inspection system provider, Biomedig - a contract research company for the pharmaceutical industry, Trackman – a provider of radar and vision systems for sports measurement equipment, <u>IHFood</u> – provider of vision technology for food production, and <u>Chemometec</u> – a provider of cell imaging and analysis. Finally the Otto Mønsted Foundation provided the financial security of the event.

The papers awards for the best paper and best student paper were sponsored by Siemens Health Care, Denmark, and were awarded to Jennifer Alvén (Sweden) for the paper "Überatlas: robust speedup of feature-based registration and multi-atlas segmentation" and Hong Pan (Denmark) for the paper "Joint spatial-depth feature pooling for RGB-D object classification". Finally, at the SCIA conference dinner held at the Trekroner Fortress in the Copenhagen harbour the award for the best Scandinavian PhD thesis for the vears 2013 and 2014 was awarded to Petter Strandmark (Sweden) for the thesis entitled "Discrete Optimization in Early Vision".

At the conference dinner representatives from the Norwegian Pattern Recognition Society announced that the 20th Scandinavian Conference on Image Analysis will be held in Tromsø, Norway, on June 12th-14th, 2017.

Non-IAPR Meeting Report

Real-time 3D Scene Understanding in year 2020

One Day British Machine Vision Association and Society for Pattern Recognition Technical Meeting

London, UK June 17, 2015 <u>http://www.bmva.org/meetings</u>

Chairs: <u>Zeeshan Zia</u> (Imperial College London, UK) <u>Andrew Davison</u> (Imperial College London, UK)

by Zeeshan Zia, General Co-chair

I recently co-chaired a British Machine Vision Association (BMVA) meeting with Prof. Andrew Davison in London, carrying the above title. We planned on having a 100% invited program and managed to get excellent speakers. The objective of the meeting was to look at softwarehardware issues in real-time 3D scene understanding, in the context of future heterogeneous manycore computing architectures.

A lot of progress in computer vision over the past two decades has been driven by exponential increase in computational resources available to researchers. Up until 2005, programmers could rely on the next technology node to deliver twice as many transistors in the same chip area (Moore's law and Dennard scaling) which were arranged into increasingly complex superscalar processors, exploited automatically by modern compilers. As Dennard scaling broke down, the next generation of computing architectures are rapidly evolving towards heterogeneous multi-cores - including GPUs and application-specific hardware accelerators. Unfortunately, with multi-core architectures, the burden is back on the shoulders of programmers to exploit the

available computational resources - state-of-the-art compilers are not able to automatically extract the parallelism inherent in our programs and offload tasks automatically to the various types of computational cores that may be available on a system. In addition, platform portability is becoming a bigger challenge - a piece of code written for my big Nvidia GPU needs many changes before it can run on my smartphone which has a very different configuration of heterogeneous embedded cores and GPU. Yet another important issue, as computer vision moves towards practical consumer applications is that of power consumption. Even with the most sophisticated customdesigned hardware for handheld mobile vision applications, for instance on the Google Tango phones, the battery discharges completely within fifteen minutes. Since battery technology is not expected to improve by leaps and bounds anytime soon, there is an urgent need to think about power consumption when designing computer vision algorithms - if we want computer vision to become commercially relevant.

The speakers correspondingly included experts in computing architectures, computer vision, and application domains including robotics and augmented reality as well as people who are working on software-hardware co-design for vision.

Andrew Davison opened the meeting by highlighting the exponential increase in computational capacity enabled by GPGPU in the last 7-8 years as opposed to the far slower improvement in CPUs. He admitted that early adoption of GPGPU technology by his group was responsible for its leading position in dense SLAM in the recent years. I particularly liked when he contrasted that a 1000 pound laptop from 2003 could run MonoSLAM, the first realtime monocular SLAM system, representing maps as 10s to 100s of sparse interest points; versus a 1000 pound laptop of today which can run ElasticFusion, the highest quality dense SLAM system that represents maps as 10s of millions of points (surfels). Its really exciting to think of this 6-7 orders of magnitude improvement in the fidelity of real-time 3D scene representation enabled almost entirely by GPUs! But does everybody in the vision community possess mastery of such unique programming models as CUDA, to keep pushing the boundaries?

Doug Watt, who is Multimedia Strategy Manager at Imagination Technologies talked about their PowerVR Imaging Framework. Imagination Technologies identifies applications of Computer Vision and Computational Photography as opportunities to differentiate from their competition - in the domain of low-power embedded graphics cores. They have been exploiting GPUCompute and implementing a number of vision pipelines for image stabilization, face detection, background removal, face beautification etc on their heterogeneous CPU and GPU fabric. After the talk, I was wondering why such big SoCs do not incorporate specific hardware accelerators for lowlevel vision kernels - the sort that are available on the Myriad Vision Processing Units (VPUs) from Movidius or the EyeQ line of chips from MobilEye. This thought was automatically addressed by Gerhard Reitmayr, Principal Engineer at Qualcomm Research, when he gave his perspective on the issue. Explaining the business of big System-on-Chip (SoC) producers, he mentioned that his company sells 10s of millions of SoCs, and everything added to the SoC must be of value to the end OEMs. These SoCs have a multitude of heterogeneous computational cores, which can be reused at the OEM level if needed. On-chip real estate is very expensive and whenever you wish to add something new, you have to justify why it couldn't be done with existing compute capability on the SoC, for example, why couldn't the modem DSP be exploited for a new computational module instead of integrating a new function-specific hardware accelerator? In general, when deciding on compute options, their first choice is to do it on the CPU (exploiting Neon SIMD engines), if

their application has already fully exploited the CPU, they go for the DSP, whereas the embedded GPU is usually busy rendering for their Augmented Reality applications. Gerhard also showed videos for a number of 3D scene understanding applications (sparse and dense SLAM, segmentation, object recognition) running on their lowpower mobile platforms.

While these speakers supported the pragmatic approach given available resources, the fact that GPUs are not well-suited to "branchy" code and the von Neumann bottleneck were not tackled until Simon Knowles. CTO of XMOS talked about a completely new computing device that he calls a "graph processor". He claimed that computational vision is nothing like graphics, for which GPUs were originally designed. In particular, highlighting machine learning workloads, he said that the associated data structures are often sparse, and arbitrary sparse structures are naturally represented as graphs. He further talked about the dynamic nature of machine learning research, citing "Dropout" in deep neural networks for regularization and Rectified Linear Units replacing Sigmoid units which were the standard for a long time. Since the field is evolving rapidly, he does not believe in committing to fine-grained details of present algorithmic pipelines. Thus, his team is working on a traditionally neglected computation model called the "Bulk Synchronous Model" employing large quantities of on-chip SRAM which will work as a "fat" communication model. The system will have multiple asynchronous cores - such that the overall system is well-suited for sparse computation. The system will need a unique "graph-parallel programming abstraction" with a sequential outer control program,

controlling "codelet" vertex functions. As far as I understood, they are evaluating this abstraction represented as an extension to the Python programming language. However, they do not propose this machine as a replacement for CPUs and GPUs, rather as a third component in a troika, with each component specializing in different kinds of workloads. I and my lab mates are really looking forward to playing with these devices soon!

Trying to act as a bridge between the architecture experts and the vision community, and as a representative of the PAMELA project, I talked about two of the projects that we on the PAMELA team are involved in. PAMELA stands for a Panoramic Approach to the Manycore Landscape, and is a project funded by a 5-year EPSRC programme grant. We have three partner universities involved in the project: Universities of Edinburgh and Manchester, and Imperial College London, with groups specialising in computing architectures, compiler and runtime design, domain specific optimisation and languages, and robot vision. The objective of the project is the same as the aims of this meeting, "to exploit future heterogeneous manycore architectures" with 3D scene understanding as the unifying challenge application. Specifically, I talked about our SLAMBench framework, which is a publiclyavailable software framework for quantitative, comparable, and validatable experimental research to investigate trade-offs in performance, accuracy, and energy consumption of various SLAM systems (KinectFusion already available, LSD-SLAM and ORB-SLAM integration under way). SLAMBench provides implementations in C++, OpenMP, OpenCL, CUDA, ARM NEON and in our paper we perform

experimentation on a variety of hardware platforms. I am particularly proud of the software level energy measurement instrumentation we have provided in the framework, which is able to read off fine-grained energy consumption at the level of individual kernels and hardware components. We have recently been performing design space exploration on top of SLAMBench, using active learning to jointly model algorithmic, compiler, and hardware parameter spaces. This machine learning enabled autotuning has been very beneficial in terms of providing us configuration points which allow significant speed-ups at much lesser energy consumption. Overall the work can allow compiler and architecture experts to optimise their designs for vision workloads - as well as computer vision researchers to compare and contrast various vision kernels in full pipeline context.

Simon Lynen, researcher at ETH-Zurich and Google Project Tango, presented a number of experiments done with Tango devices. To give some background: Google introduced its Project Tango last year, which comes in the form of two mobile platforms (a smartphone and a tablet) incorporating a depth camera as well as a high-performance computing device for vision processing. The project allows highly accurate real-time SLAM by tightly coupling inertial sensing with visual tracking. While the visual tracking appears to be based on sparse keypoints in selected keyframes, the map representation comprises of an occupancy grid populated from the depth sensor. This allows for a wide range of use cases specially in the area of augmented reality (some amazing videos are available on Youtube). The phone

performs its vision processing on a Myriad vision processing unit from Movidius - having generalpurpose vector processing capabilities together with a number of low-level hardware accelerator for vision tasks such as edge detection - all fitting within a tiny power envelope. The tablet on the other hand, comprises of an Nvidia Jetson TK1 SoC, which is a low-power version of Nvidia's standard GPU technology, which we thoroughly evaluate in our SLAMBench work also. The existence of these devices further emphasize the importance of domain-specific optimization and architectures in enabling low-power high-performance vision. Simon mentioned how the platform has been maturing rapidly, with the odometry drift going from 1% (of the trajectory) just two years ago to 0.1% now. He has been testing the device under challenging conditions from fourteen consecutive roller coaster rides (resulting in only 3.5% drift) to Zero gravity flights to real-time obstacle avoidance on a quadrotor. The device has even been sent to the International Space Station. One interesting challenge he highlighted was that the integrated depth camera, due to eye-safety concerns, delivers less than 2.5% of the data as a standard Kinect camera.

Mike Aldred, Electronics lead at Dyson, gave the perspective of an integrator of these technologies, as he talked about their recent Dyson 360 Eye vacuum cleaner robot. As opposed to its predecessor, the famous Roomba robot, which randomly roams around the environment changing its direction when it hits an obstacle, the 360 Eye is a serious product with a powerful 120W vacuum cleaner. To avoid wasting energy, the device needs to do methodical motion and thus incorporates an omnidirectional camera (128 x 1024 resolution) used for sparse feature tracking and mapping (together with a few other sensors). He mentioned how even successful academic algorithms such as MonoSLAM or PTAM, still desired more than a decade of real-world testing and tuning, before they could reach a mature prototype. They have tested the robot in 100,000 homes already, and still keep finding new challenges - he described some funny and unexpected ways in which humans interact with these devices sometimes resulting in failure. He stressed that the vision community should explicitly report fallibilities together with capabilities of our algorithms - so the integrators can know what the challenges are. He shared his wisdom on other issues including pitfalls of premature optimisation and importance of keeping the architecture generic as far as possible.

Thomas Whelan, Dyson Research Fellow at Imperial College London talked about "the rise of dense methods" is SLAM. He started with a quick survey of dense SLAM algorithms and how these have been enabled primarily by recent advances in cheap consumergrade hardware: GPUs and RGB-D cameras. He sees desktop/ laptop GPUs reaching tens of teraflops and sensors moving towards higher resolution and frame rates by 2020 - and stressed that these improvements will in turn make things much easier for the algorithms. He also described his state-of-the-art large scale dense SLAM algorithms Kintinuous and the latest ElasticFusion (being presented at RSS 2015). The audience loved Tom's live demo of ElasticFusion. He commented that in the domain of dense methods, we are moving towards 4D reconstruction methods

Page 22

IAPR Newsletter, Vol. 38 No. 3, Jul. 2015

(for dynamic scenes), where approaches will be going beyond the lambertian assumption and incorporating increasingly more sophisticated semantic concepts in real-time pipelines. Finally, he gave a sneak peak into his current research project, where he is able to recover the direction and position of lighting sources from his dense reconstructions!

Apart from speakers directly focusing on the intersection of computing and vision, we also had a few amazing speakers sharing their experiences in pure vision and robotics, while still citing the importance of computation.

Maurice Fallon, Lecturer at Edinburgh University and Perception lead on MIT's entry to the recent DARPA Robotics Challenge, talked about his experiences designing and operating tele-operated humanoid robots working in disaster scenarios. Further, he showed some recent work involving autonomous walking with a passive camera, merging Kintinuous with a footstep planning algorithm. I wondered, seeing the recent deep learning revolution, if it might be possible to learn the manipulation and footstep planning from a combination of real and synthetic simulated data. Unfortunately, in Maurice's experience, physical contacts and motion are still difficult to model faithfully and thus he doesn't believe in

simulated training for mid-level control problems. In fact, he is soon going to get a NASA Valkyrie robot in Edinburgh to advance his research with humanoids - where he plans to work on autonomous operation. Really exciting!!

While the vision component of the workshop was SLAM-heavy, we were lucky to also have serious machine learning muscle amongst ourselves. Jamie Shotton, Principal Researcher at Microsoft is well known for his contributions to random forests applied to computer vision problems. He showcased their latest research in real-time hand tracking from depth cameras, and how the technology is rapidly evolving into a robust and useable input device. Responding to a question about why there are never multiple hands tracked in his demo videos, he responded that the required computation to estimate a single hand pose takes up the whole GPU - again emphasizing the need for holistic thinking when it comes to highperformance vision!

Renato Salas-Moreno, co-founder of Surreal Vision, recently acquired by Oculus Research/Facebook talked about his PhD research at Imperial College London. He likes to think of computer vision as "inverse video game design", and is the developer of SLAM++ and Dense Planar SLAM (check out YouTube for some impressive videos, which have received 50,000+ views in total). These realtime SLAM pipelines incorporate semantic concepts into localisation and mapping - enabled by Renato's thorough understanding of Nvidia's GPUs and OpenGL. Renato demonstrated how performing 3D object detection and plane fitting inside the SLAM pipeline, allow for unique augmented reality applications.

The last half-hour of the programme was dedicated to a panel discussion with all of our speakers. Without giving away the "answers" we got, I will let you know the questions I scripted to ask our panelists! We will be happy to include your thoughts on these in future issues of this Newsletter... :-)

1. What is the most impressive real-time 3D scene understanding demo that you expect to see at CVPR or ISMAR 2020?

2. Seeing that computer vision and machine learning are evolving rapidly, do you think it is prudent to already start optimizing them for software/hardware implementation?

3. What is the right level of abstraction to target for vision DSLs or SoCs? low-level kernels, generic parallel patterns, full pipelines, or something else.

A version of this report will also be published in the BMVA Newsletter.

Thoughts on this or other articles you've read in the IAPR Newsletter or elsewhere?

Send your comments to:

Arjan Kuijper, Editor-in-Chief arjan.kuijper@igd.fraunhofer.de

BOOKSBOOKSBOOKS

FREE BOOKS

The *IAPR Newsletter* is looking for reviewers for the books listed below. If you have interest and some knowledge in the topic, email us with your mailing address. We will send you a copy of the book—which you may keep—and will expect in return a review for the Newsletter. ~<u>Zeeshan Zia</u>, Associate Editor for Book Reviews

We are offering the following latest titles for review. These have been published (or will very soon be available) in the "Advances in Computer Vision and Pattern Recognition" series from Springer.

* *Computer Vision in Sports* by Thomas B. Moeslund et al. (ISBN 978-3-319-09395-6):<u>http://www.springer.com/computer/</u> <u>image+processing/book/978-3-319-09395-6</u>

* *Smart Information Systems* by Frank Hopfgartner (ISBN 978-3-319-14177-0): <u>http://www.springer.com/computer/inform</u> <u>ation+systems+and+applications/book/978-3-319-14177-0</u>

* *Advances in Embedded Computer Vision* by Branislav Kisačanin and Margrit Gelautz (ISBN 978-3-319-09386-4): <u>http://www.springer.com/computer/image+processing/book/978-3-319-09386-4</u>

* *Video Text Detection* by Tong Lu et al. (ISBN 978-1-4471-6514-9):<u>http://www.springer.com/computer/image+processing/</u> <u>book/978-1-4471-6514-9</u>

* Computer Vision and Machine Learning with RGB-D Sensors by Ling Shao et al. (ISBN 978-3-319-08650-7):<u>http://www.springer.com/computer/image+processing/book/978-3-319-08650-7</u>

Other Springer titles of interest recently published include:

* *Advanced Video Coding Systems* by Wen Gao and Siwei Ma (ISBN 978-3-319-14242-5):<u>http://www.springer.com/com-puter/theoretical+computer+science/book/978-3-319-14242-5</u>

* *Visual Cryptography for Image Processing and Security* by Feng Liu and Wei Qi Yan (ISBN 978-3-319-09643-8):<u>http://www.springer.com/computer/image+processing/book/978-3-319-09643-8</u>

* *Natural User Interfaces in Medical Image Analysis* by Marek R. Ogiela and Tomasz Hachaj (ISBN 978-3-319-07799-4): <u>http://www.springer.com/computer/image+processing/book/978-3-319-07799-4</u>

* *Fusion in Computer Vision* by Bogdan Ionescu et al. (ISBN 978-3-319-05695-1): <u>http://www.springer.com/computer/</u> <u>image+processing/book/978-3-319-05695-1</u>

* *Biometrics in the New World* by Julian Ashbourn (ISBN 978-3-319-04158-2): <u>http://www.springer.com/computer/</u> <u>image+processing/book/978-3-319-04158-2</u>

* *Image Blending Techniques and their Application in Underwater Mosaicing* by Ricard Prados, Rafael Garcia, and László Neumann (ISBN 978-3-319-05558-9): <u>http://www.springer.com/computer/image+processing/book/978-3-319-05557-</u>2

* *Computational Intelligence in Multi-Feature Visual Pattern Recognition* by Pramod Pisharady, Prahlad Vadakkepat and Loh Ai Poh (ISBN 978-981-287-055-1): <u>http://www.springer.com/us/book/9789812870551</u>

And from MIT Press:

* *Developmental Robotics: From Babies to Robots* by Angelo Cangelosi and Matthew Schlesinger (ISBN 978026202801): <u>http://mitpress.mit.edu/books/developmental-robotics</u>

This bulletin board contains items of interest to the **IAPR Community**

Call for Nominations

IAPR Fellow Award

Deadline: January 31, 2016

IAPR Fellows requested to submit, in-depth survey papers in Pattern Recognition Letters

Click here to go to the IAPR web site for In addition to capturing ground breaking research work in a timely manner through short letters, Pattern Recognition Letters (PRL) has started a new category of paper submission of survey papers authored by IAPR Fellows.

These papers are meant to capture general trends of key research areas of broad interest and to offer guidance about future research directions. They go substantially beyond just a literature review. They are meant to be mileposts along a research theme, marking the end of a period and the beginning of a new one.

These papers do not have to comply with the length restriction. The review process will determine the appropriateness of the length.

Pattern Recognition Letters is an official publication of the IAPR.

~G. Borgefors, G. Sanniti di Baja, S. Sarkar, Editors-in-Chief, PRL

Pending Book Reviews

We are looking forward to reports on the following books under review:

1. Airborne and Terrestrial Laser Scanning by George Vosselman, Hans-Gerd Maas (Whittles Publishing, 2010). Reviewer: Giuseppe Maino. Under review since February 2010.

2. <u>Support Vector Machines for Pattern Classification, 2nd</u> Ed., by Shigeo Abe (Springer, 2010). Reviewer: Huthaifa Abderahman. Under review since July 2013.

3. Concise Computer Vision by Reinhard Klette (Springer 2014). Reviewer: Tayyab Naseer. Under review since February 2014.

4. Decision Forests for Computer Vision and Medical Image Analysis by Antonio Criminisi and Jamie Shotton, Eds. (Springer 2013). Reviewer: Zeeshan Zia. Under review since April 2014.

5. *Imaging Spectroscopy for Scene Analysis* by Antonio Robles-Kelly and Cong Phuoc Huynh (Springer 2013). Reviewer: Tanish H. Zaveri. Under review since August 2014.

6. Scalable Pattern Recognition Algorithms by Pradipta Maji and Sushmita Paul. Reviewer: Munish Kumar. Under review since February 2015.

7. Handbook of Biometric Anti-Spoofing edited by Sebastien Marcel et al. Reviewer: George A. Papakostas. Under review since February 2015.

Meeting and Education Planner

The IAPR web site has the most up-to-date information on IAPR events. Click <u>here</u>. NOTE: Highlighting indicates that the paper submission deadline has not yet passed. * Asterisks denote non-IAPR events *

		Meeting	Report on previous edition	Venue
		GREC 2015: Eleventh International Workshop on Graphics Recognition	<u>GREC 2013</u>	France
2015	SEP AUG	<u>CBDAR 2015</u> : 6th International Workshop on Camera Based Document Analysis and Recognition	<u>CBDAR</u> 2013	France
		HIP 2015: 3rd International Workshop on Historical Document Imaging and Processing	<u>HIP 2013</u>	France
		ICDAR 2015: 13th International Conference on Document Analysis and Recognition	<u>ICDAR</u> 2013	France
		ICIAP 2015: 18th International Conference on Image Analysis and Processing	ICIAP 2013	Italy
		* <u>ICIP 2015</u> : IEEE International Conference on Image Processing *		Canada
	NOV	ACPR 2015: 3rd Asian Conference on Pattern Recognition	ACPR 2013	Malaysia
		CIARP 2015: XX Iberoamerican Congress on Pattern Recognition		Uruguay
		DICTA 2015: International Conference on Digital Image Computing: Techniques and Applications	DICTA 2012	Australia
		ICVNZ 2015: Image and Vision Computing New Zealand		New Zealand
		PSIVT15: 7th Pacific Rim Symposium on Image and Video Technology	PSIVT13	New Zealand
$\overline{\bigcirc}$				
2016	APR	DAS 2016: 12th IAPR International Workshop on Document Analysis Systems	<u>DAS 2014</u>	Greece
	MAY	ICISP 2016: 7th International Conference on Image and Signal Processing	<u>ICISP 2014</u>	Canada
	NU	ICB 2016: 9th International Conference on Biometrics	<u>ICB 2015</u>	Sweden
	DEC	ICPR 2016: 23rd International Conference on Pattern Recognition	<u>ICPR 2014</u>	Mexico
• The IA DR Newslatter is published in association with the IADD website numuriate are				
•	The IAPR Newsletter is published four times per year January April July and October			
•	Click here to subscribe			
Chick here to subscribe.				
• Deadline for the next issue: September 25, 2015				
10 contact us:				

Arjan Kuijper, Editor-in-Chief, arjan.kuijper@igd.fraunhofer.de Zeeshan Zia, Associate Editor for Book Reviews, <u>zeeshan.zia@imperial.ac.uk</u> Linda J. O'Gorman, Layout Editor, <u>logorman@alumni.duke.edu</u>

.

.

.