Journal of Signal Processing Systems (2018) 90:873–876
https://doi.org/10.1007/s11265-018-1365-8

Guest Editorial: Special Issue on Embedded Computer Vision

Stefano Mattoccia[1] · Branislav Kisačanin[2] · Margrit Gelautz[3] · Sek Chai[4] · Ahmed Nabil Belbachir[5] · Goksel Dedeoglu[6] · Fridtjof Stein[7]

Published online: 26 April 2018

It is with great pleasure that we present this Special Issue of the *Journal of Signal Processing Systems* (JSPS) dedicated to Embedded Computer Vision! We are pleased to include six state-of-the-art papers from the leaders in this field, both from industry and academia, who keep pushing the embedded computer vision technology forward.

While the idea for this special issue originated between the Guest Editors at one of the CVPR workshops on the same topic that we have organized, it is the work of the contributing authors that makes it a success. The papers were solicited from the workshop participants and through an open call for papers, so the initial submissions were in many ways already pre-filtered. Out of 24 submitted papers, the highly selective review process yielded the six papers included here. They cover a broad range of challenges that are encountered in practical deployment of embedded vision systems, especially when high computational performance needs meet limited resources. We present papers describing a range of novel solutions: a deep learning accelerator, a robust aerial tracking system, an FPGA-based aerial visual servoing task solution, an approach to use low-cost hardware for real-time vision, a real-time motion detector, and an image enhancement approach based on human vision.

In the following we summarize their contributions:

- In their paper "Efficient Object Detection Using Embedded Binarized Neural Networks, (https://doi.org/10.1007/s11265-017-1255-5)" the authors Jaeha Kung (Georgia Tech), David Zhang (SRI international), Gooitzen van der Wal (SRI international), Sek Chai (SRI international), and Saibal Mukhopadhyay (Georgia Tech) discuss their research on energy efficient deep learning accelerators and their application on vision-based object detection. The paper presents the training and implementation of a deep neural network with single bit weights, and it offers applications-based insights into the power-performance tradeoffs gained with a binarized approach.
- In their paper "Watch Out: Embedded Video Tracking with BST for Unmanned Aerial Vehicles, (https://doi.org/10.1007/s11265-017-1279-x)" the authors Francesco Battistone (MER MEC S.p.A.), Alfredo Petrosino (University Parthenope of Naples), and Vincenzo Santopietro (University Parthenope of Naples) present a real-time tracking system that is able to efficiently run on an Nvidia Jetson board mounted on a UAV. The approach to long term video tracking implemented in Watch Out is named Best Structured Tracker (BST) and its performance has been verified both on challenging datasets and in real

✉ Stefano Mattoccia
stefano.mattoccia@unibo.it

Branislav Kisačanin
b.kisacanin@ieee.org

Margrit Gelautz
margrit.gelautz@tuwien.ac.at

Sek Chai
sek.chai@gmail.com

Ahmed Nabil Belbachir
nabil@belbachir.info

Goksel Dedeoglu
goksel@perceptonic.com

Fridtjof Stein
fridtjof.stein@gmail.com

[1] University of Bologna, Viale Risorgimento, 2, 40136 Bologna, Italy

[2] Nvidia Corporation, 3113 Ipswich Drive, Plano, TX 75025, USA

[3] Vienna University of Technology, Favoritenstrasse 9-11, 1040 Vienna, Austria

[4] SRI International, 53 York Dr, Princeton, CA 08540, USA

[5] Teknova, Tordenskjoldsgate 9, 5 etg, NO-4612 Kristiansand, Norway

[6] PercepTonic, 2208 Valcourt Dr, Plano, TX 75023, USA

[7] Daimler AG, Helene-Lange-Strasse 52, 73760 Ostfildern, Germany

situations using an Nvidia Jetson board mounted on a drone. Results show that a robust system, like Watch Out, can track almost every possible target in real time.

- In their paper "FPGA-Based Fast Response Image Analysis for Orientational Control in Aerial Manipulation Tasks, (https://doi.org/10.1007/s11265-017-1286-y)" the authors Robert Ladig, Suphachart Leewiwatwong, and Kazuhiro Shimonomura (all from Ritsumeikan University) explore the feasibility of using an on-board field-programmable gate array (FPGA) for aerial visual servoing tasks. They successfully designed a novel derivative of Fast Increment Hough Transform 2 (FIHT2) fit for implementation on an FPGA. The algorithm is able to identify the orientation and distance of bar-like objects even in cluttered environments, solely by using the FPGA and a monocular camera image. The algorithm is then tested against a practical real-life application, where a gripper mounted on an aerial robot is able to autonomously align towards a bar-like object using this method.
- In their paper "Image Processing Units on Ultra-low-cost Embedded Hardware: Algorithmic Optimizations for Real-time Performance, (https://doi.org/10.1007/s11265-017-1267-1)" the authors Suraj Nair (TUM CREATE, Singapore), Nikhil Somani (TUM CREATE, Singapore), Artur Grunau (Technische Universitat Munchen), Emmanuel Dean-Leon (Technische Universität München), and Alois Knoll (Technische Universität München) discuss the growing popularity of low cost single board computers (SBC) and how they can be used to build real-time computer vision (CV) applications. One of the key challenges presented and addressed in the paper is how popular computer vision algorithms can be mapped to the graphics processor of such SBCs given no support for high-level GPU APIs such as CUDA or OpenCL. Taking RaspberryPi, one of the most commonly used SBCs, as a demonstration platform the authors present the re-engineering of CV algorithms to overcome the hardware limitations of the SBC's GPU while still achieving real-time performance.
- In their paper "Real-Time Embedded Motion Detection via Neural Response Mixture Modeling, (https://doi.org/10.1007/s11265-017-1265-3)" the authors Mohammad Javad Shafiee (University of Waterloo), Parthipan Siva (Aimetis Corporation), Paul Fieguth (University of Waterloo), and Alexander Wong (University of Waterloo) discuss their new research on utilizing deep neural networks for real-time applications on embedded systems. They propose a new framework to do real-time motion detection via the rich deep feature extracted from the neural response of an efficient, stochastically-formed deep neural network, the so-called StochasticNet. The neural response features are utilized to construct Gaussian mixture models to detect motion in a scene. The Neural response mixture (NeRM) model is examined and performed on an Axis surveillance camera in a real-time manner. Results show that the proposed NeRM approach can improve the GMM performance with fewer false detections and illustrates the possibility of utilizing deep neural networks on embedded devices in real-time applications.
- In their paper "Digital Image Fusion Using HVS in Block Based Transforms, (https://doi.org/10.1007/s11265-017-1252-8)" the authors Vadhi Radhika (Jawaharlal Nehru Technological University), Kilari Veeraswamy (QIS College of Engineering and Technology), and Samayamantula Srinivas Kumar (Jawaharlal Nehru Technological University) present an image enhancement method based on human visual system (HVS) and block transforms. Their method shows improvements in achieved quality while having reduced complexity compared to earlier methods.

Embedded computer vision as a field has been at the forefront of solving practical issues of computer vision and we hope this special issue will help shed light on the many current and future applications that it has made possible.

Guest Editors of the JSPS Special Issue on Embedded Computer Vision.
Stefano Mattoccia, Branislav Kisačanin, Margrit Gelautz, Sek Chai, Nabil Belbachir, Goksel Dedeoglu, Fridtjof Stein

Stefano Mattoccia received a Master degree in Electronic Engineering and a Ph.D. in Computer Science and Engineering from University of Bologna, respectively, in 1997 and 2002. Currently he is Associate Professor at the School of Engineering and Architecture of the University of Bologna, Department of Computer Science and Engineering. His research interests include computer vision, 3D vision, machine learning, and embedded vision and he regularly serves as reviewer for major journals and conferences in these fields. He was General Chair of the 2016 CVPR Embedded Vision Workshop, Program Chair of the same workshop in 2015 and 2014, Area Chair of the IEEE International Conference on Multimedia & Expo ICME in 2014 and 2013 and a Guest Editor in 2013 for IEEE Journal of Selected Topics in Signal Processing for a special issue on Emerging Techniques in 3D.

Branislav Kisačanin received M.S. and Ph.D. degrees in EE/CS from the University of Illinois at Chicago in 1998. While at Delphi (1998–2007) he pioneered the embedded vision approach in development of automotive vision systems and was inducted into the Delphi Innovation Hall of Fame. While at Texas Instruments (2007–2013) he co-founded a skunkworks project that delivered EVE, a popular automotive vision accelerator. At Interphase (2013–2015) he contributed embedded vision algorithms and production code for PenVeu, an interactive electronic pen. Since 2015 he has been with Nvidia, working on development of PVA, a vision accelerator for Xavier, the first SoC for autonomous vehicles. Branislav was a General Chair of the CVPR Embedded Vision Workshop several times, most recently in 2013, and a co-editor of the books Embedded Computer Vision (Springer, 2009) and Advances in Embedded Computer Vision (Springer, 2014).

Margrit Gelautz is an associate professor at the Department of Informatics at Vienna University of Technology, Austria. She obtained her Ph.D. degree in Telematics from Graz University of Technology, Austria, in 1997. Afterwards, she spent two years as a visiting scholar at the EE Department of Stanford University, USA, working in the field of radar remote sensing. At Vienna University of Technology, she built up a research group on Image and Video Analysis & Synthesis, with a focus on 3D reconstruction and algorithms for stereo analysis, image matting and related applications. She was director of the Doctoral College on Computational Perception at Vienna University of Technology (2010–2013), and vice chair of the IEEE Austria Section (2012–2014). She acts as a reviewer for major journals, including IEEE PAMI, IEEE TCSVT and IEEE TIP, and serves as a project evaluator for the European Union and several national research funding agencies. Margrit was a General Chair of the 2013 CVPR Embedded Vision workshop and a co-editor of the book Advances in Embedded Computer Vision (Springer, 2014).

Sek Chai is a technical director at SRI International. He manages research and development programs in power efficient systems, high performance computing, computer vision, and machine learning. He currently leads projects in lifelong learning by enabling AI systems to know what to learn and when. Prior to joining SRI, he developed imaging and video solutions for next-generation mobile devices and home broadband products at Motorola Labs. He is a senior member of IEEE and ACM, and holds a Ph.D. in Electrical Engineering from Georgia Tech. Sek was a co-editor of the book Embedded Computer Vision (Springer, 2009). He is also a key member of the CVPR Embedded Vision Workshop (2008–2018) and CVPR Vision Industry and Entrepreneur Workshop (2013–2017).

Ahmed Nabil Belbachir was awarded a Ph.D. in Computer Science from Vienna University of Technology (Austria) in 2005 and the engineering degree in electronics from the University of Science and Technology of Oran (Algeria) in 1996. Currently, he is the head of the department "Smart Instrumentation" at Teknova (Norway) with the strategical objective to bring leading edge technologies and outstanding research ideas to industrial prototypes within the key markets of its partners and customers that are the Offshore Oil and Gas, metal-production process industries and energy sectors. He is also chief scientist with a research focus on smart cameras and machine intelligence. His research interests include pattern analysis, computer vision, computational intelligence and (embedded) real-time systems. He has published more than 80 peer-reviewed scientific papers, including two books: "Smart Cameras" as editor and "On-board Processing for Infrared Observatories" as author. He is the holder of four patents. Nabil was a General Chair of the CVPR Embedded Vision Workshop several times, most recently in 2012.

Goksel Dedeoglu is the Lab Director of PercepTonic, his Computer Vision consulting practice based in Dallas. After receiving a Ph.D. degree from the Robotics Institute at Carnegie Mellon, Goksel joined Texas Instruments' R&D Labs where he made product contributions and received patents in automotive safety, video security, retail analytics, wearables and augmented reality. Goksel is a frequent contributor to the embedded vision industry, helping organize the annual Embedded Vision Workshop, serving on various advisory boards and review panels, and presenting at conferences. Goksel is currently pioneering Computer Vision on UAVs for infrastructure asset management as Co-Founder and CTO of Brains4Drones. Goksel was a General Chair of the CVPR Embedded Vision Workshop a couple of times, most recently in 2015.

Fridtjof Stein received the Dipl.-Inf. degree in computer science from University of Karlsruhe, Karlsruhe, Germany. Subsequently, he continued to study computer science and received the Ph.D. degree from University of Southern California, Los Angeles, in 1992. Since then he has been with Daimler AG, where his main research interests have been intelligent vehicles and computer vision, with focus on camera based environment perception in the fields of driver assistance systems and autonomous driving. Fridtjof was a General Chair of the 2014 CVPR Embedded Vision Workshop.