Introduction to the Special Section of CVPR 2017

Yanxi Liu, James M. Rehg, Camillo J. Taylor, and Ying Wu, *Fellow, IEEE*

1 Introduction

The past few years have been significant to our scientific community of Computer Vision and Pattern Recognition (CVPR) in many ways. We just (June 2022) held the first in-person conference in New Orleans since COVID-19. The number of total attendees at CVPR 2022 shot up to 10,000 while the number of physical attendees was 5K, a good reminder to us of the time when attendees hit the 5K mark for the first time at CVPR 2017 in Hawaii. Time indeed flies; our CVPR 2017 program chairs remain firm advocates for this set of five papers appearing here, before you, as a *TPAMI* special section. They are selected from CVPR 2017 and have gone through a rigorous process of paper review and revision, while enduring various time deferring factors that were out of everyone's control. Ultimately, they have stood the test of time.

2 Paper Review and Selection

All members of the program committee (PC) were involved in the selection of the papers in this special section from CVPR 2017 published papers. The selected papers then followed an anonymous double-blind protocol. Each paper was allocated at least four expert reviewers, most of whom had reviewed the same conference paper previously. After the reviews were released to the authors, the authors were given the opportunity to revise and resubmit. The final decisions were made after a second review by the same reviewers. The careful and thorough reviews from these domain experts are highly appreciated and acknowledged.

3 Papers in this Special Section

3.1 "Convolutional Networks With Dense Connectivity" by Huang, Liu, Pleiss, Maaten, and Weinberger

The authors introduce the Dense Convolutional Network (DenseNet), which connects each layer to every other layer in a feed-forward fashion. This paper expands upon the corresponding conference publication by adding ablation studies of the various architectural hyperparameter choices and providing details of a memory efficient implementation of the proposed architecture. The ablation studies are invaluable for understanding the contribution of various design choices to the strength of the described model. DenseNets obtain significant improvements over the state-of-the-art whilst requiring less parameters and computation to achieve high performance.

3.2 "Deep Audio-Visual Speech Recognition" by Afouras, Chung, Senior, Vinyals, and Zisserman

The goal of this work is to recognize phrases and sentences being spoken by a talking face, with or without the audio. This work presents one of the first audio-visual speech recognition systems, trained under an end-to-end architecture and including a relatively large scale dataset (publicly released). This journal version contains a CTC-based objective that accounts for alignment issues between the input signal and the output objective (text). This work demonstrates the efficacy of self-attention models on the audio-visual speech recognition task. The proposed models outperform previous approaches by a wide margin.

3.3 "Computational Imaging on the Electric Grid" by Sheinin, Schechner, and Kutulakos

Because "Night beats with alternating current (AC) illumination", this work leads to a new, revealing way to obtain scene information: the type of bulbs in the scene, the phases of the electric grid up to city scale, and the light transport matrix. The temporal response can then be measured using optical sensors, and the uniqueness of the transfer function provides an identifier that allows unmixing of different passive illuminants using temporal information alone. A number of impressive examples of computational illumination are shown, including decomposition of electric grid into 3-phase power lines, separation and computational relighting of passively illuminated scenes, and decomposition of transmission/reflections caused by independent sources.

3.4 "Estimation of Wetness and Color From A Single Multispectral Image" by Okawa, Shimano, Asano, Bise, Nishino, and Sato

This paper presents a new model for material surface wetness and color estimation from a single multispectral image. The model is theoretically derived with reasonable assumptions, and empirical validated. An alternating minimization method is proposed to recover the parameters, which leads

- *Yanxi Liu is with the School of EECS, Pennsylvania State University, State College, PA 16801 USA. E-mail: yul11@psu.edu.*
- *James M. Rehg is with the School of Interactive Computing, Georgia Institute of Technology, Atlanta, GA 30332 USA. E-mail: rehg@gatech.edu.*
- *Camillo J. Taylor is with Computer and Information Science Department, University of Pennsylvania, State College, PA 16801 USA. E-mail: cjtaylor@seas.upenn.edu.*
- *Ying Wu is with Electrical Engineering and Computer Science Department, Northwestern University, Evanston, IL 60208 USA. E-mail: yingwu@northwestern.edu.*

Manuscript received 1 January 1899; accepted 1 January 1899. Date of current version 3 November 2022.
(Corresponding author: Yanxi Liu.)
Digital Object Identifier no. 10.1109/TPAMI.2022.3201636

to an estimate of the spatially-varying wetness. An extension is also considered to multiple materials. Recognizing wet surfaces and their degrees of wetness may lead to many practical computer vision applications.

3.5 "Multi-View Supervision for Single-View Reconstruction via Differentiable Ray Consistency" by Tulsiani, Zhou, Efros, and Malik

This work investigates the notion of consistency between a 3D shape and a 2D observation, and proposes a differentiable formulation that allows computing gradients of the 3D shape given an observation from an arbitrary view. A differentiable ray consistency (DRC) constraint is incorporated in a learning framework to leverage different types of multi-view observations and with input modalities such as RGB, foreground/background segmentation and semantic segmentation. The proposed approach improves over existing techniques for single-view reconstruction of objects from the PASCAL VOC dataset.

Yanxi Liu is a professor with the School of EECS, Penn State University. She is director of the Human Motion Lab for Smart Health and co-directs the Lab for Perception, Action and Cognition (LPAC). She is trained in physics/EE (BS), computer science (MS) and theoretical robotics/AI (PhD). A central theme of her research is "computational regularity" with diverse applications in robotics, human/machine perception, and human health. She was a visiting professor with ETH (2016-2017), Tsinghua University (2015) and Stanford University (2013-2014). Her industrial experience includes a summer visit to Microsoft Silicon Valley, and a year-long visit to Google, Mountain View, which resulted in two granted US patents on applying computational symmetry to urban scenes. She led three international competitions on Symmetry Detection (ICCV 2017, CVPR 2013, CVPR 2011). She has served as a program co-chair for CVPR 2017 and WACV 2019, area chair for multiple CVPR/ECCV/ICCV/MICCAI conferences, and as an associate editor for *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

James M. Rehg received the PhD degree from CMU in 1995 and worked with the Cambridge Research Lab of DEC (and then Compaq) from 1995-2001, where he managed the computer vision research group. He is a professor with the School of Interactive Computing with the Georgia Institute of Technology, where he co-directs the Center for Health Analytics and Informatics. He received an NSF CAREER Award in 2001 and a Raytheon faculty fellowship from Georgia Tech in 2005. He and his students have received a number of best paper awards, including best student paper awards at ICML 2005, BMVC 2010, Mobihealth 2014, Face and Gesture 2015, and a 2018 Distinguished Paper Award from ACM IMWUT and a Method of the Year Award from Nature Methods. He served as the general co-chair for CVPR 2009 and the program co-chair for CVPR 2017. He has authored more than 200 peer-reviewed scientific papers and holds 26 issued US patents.

Camillo J. Taylor received the AB degree in electrical computer and systems engineering from Harvard College in 1988 and the MS and PhD degrees from Yale University in 1990 and 1994 respectively. He is the Raymond S. Markowitz President's distinguished professor with the Computer and Information Science Department and GRASP Laboratory, University of Pennsylvania. His research interests lie primarily in the fields of computer vision and robotics and include reconstruction of 3D models from images, vision-guided robot navigation and scene understanding. He has served as an associate editor for *IEEE Transactions of Pattern Analysis and Machine Intelligence*. He has also served on numerous conference organizing committees. He is a general chair of the International Conference on Computer Vision 2021 and was a program chair of the 2006 and 2017 editions for IEEE Conference on Computer Vision and Pattern Recognition and of the 2013 edition of 3DV. In 2012 he was awarded the Christian R. and Mary F. Lindback Foundation Award for distinguished teaching with the University of Pennsylvania.

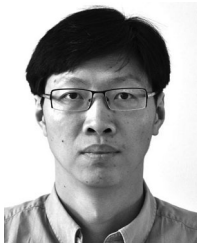

Ying Wu (Fellow, IEEE) received the BS degree from the Huazhong University of Science and Technology, Wuhan, China, in 1994, the MS. degree from Tsinghua University, Beijing, China, in 1997, and the PhD degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, Urbana, Illinois, in 2001. He is currently full professor of Electrical Engineering and Computer Science with Northwestern University. His current research interests include computer vision, robotics, image and video analysis, pattern recognition, machine learning, multimedia data mining, and human-computer interaction. He serves as the editor-in-chief for IAPR *APR Journal of Machine Vision and Applications*, and associate editor for *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *IEEE Transactions on Image Processing, IEEE Transactions on Circuits and Systems for Video Technology*, *SPIE Journal of Electronic Imaging*. He served as a program chair and area chairs for CVPR, ICCV and ECCV. He received the Robert T. Chien Award at UIUC in 2001, and the NSF CAREER Award in 2003. He is a Fellow of the IEEE and IAPR.