State of the Journal Editorial

Sven Dickinson

I would like to take this opportunity to bring our readership up to date on the state of the journal. My yearly editorial usually appears in January of each year, but I've delayed it a month so that I can announce our largest cohort yet of new associate editors. By the time this editorial appears, I will have started my fourth year in the role and first of my two-year reappointment. It's once again been a very good year, with our impact factor leaping from 9.455 (2017) to 17.30 (2018), establishing *IEEE Transactions on Pattern Analysis and Machine Intelligence* as the top-ranked journal in all of computer science. Moreover, our submissions are up from last year; as of Nov 13, 2019, we had 1,101 submissions, compared to 992 as of Nov 13, 2018. For papers accepted in 2018, the average time from submission to first decision is 3.8 months, while the average time from submission to publication on Xplore is 10.8 months. While both metrics reflect an improvement over last year, I'm still short of my target. As I mentioned last year, with the increasing emphasis that our community places on conference papers, I'd ideally like to get the time from submission to online publication down to 6-7 months, which is comparable to the time from conference paper submission to conference paper presentation.

Over the past two years, a big part of my strategy to reduce time to acceptance and publication is to appoint more Associate Editors (AEs), which will reduce the workload per AE, hopefully allowing our AEs to focus their energy on fewer papers and shepherd them more efficiently. In 2017, we added 26 new AEs, and in 2018, we added another 22. I'm pleased to report that in the Fall of 2019, we've added 39 new AEs, who I'll introduce shortly. I plan to continue to add another 20-30 additional AEs over the coming year, as I continue to retire those that have exceeded their 2+2-year editorial board terms. Joyce Arnold and I have also converged on a new strategy for flagging papers that require urgent attention, and this should further reduce time to acceptance and publication.

As always, my favorite part of this exercise is introducing our new editorial board members and thanking our retiring board members. After 5 years of outstanding service to our journal, Kyoung Mu Lee has stepped down as Associate Editor-in-Chief (AEIC). I can't thank him enough for his exceptional diligence, responsiveness, judgement, and support. And I'm particularly grateful to Kyoung Mu for staying on in the role while we recruited his successor. His commitment to our journal has really been exemplary, and I will sincerely miss working closely with him. I'm pleased to announce that Kyoung Mu will be replaced by René Vidal, who brings leadership and broad expertise to the role. I'm really looking forward to working with René!

I'm pleased to announce a new cohort of 39 Associate Editors that have joined since January, 2019: Amr Ahmed, Xiang Bai, Dima Damen, Kosta Derpanis, Giovanni Farinella, Ryan Farrell, Yasu Furukawa, Jim Glass, Andras Gyorgy, Tim Hospedales, Brian Kingsbury, Ajay Kumar, Simon Lacoste-Julien, Lihong Li, Ce Liu, Tie-Yan Liu, Wei Liu, Chen Change Loy, Michael Maire, Deyu Meng, Vlad Morariu, Cheng Soon Ong, John Paisley, Thomas Pock, Liva Ralaivola, Xiaofen Ren, Irina Rish, Amit K. Roy-Chowdhury, Yaser Sheikh, Suvrit Sra, Ping Tan, Christian Theobalt, Radu Timofte, Lorenzo Torresani, Chong Wang, Jue Wang, Richard Wildes, Christian Wolf, and Lei Zhang. These individuals have been selected not only for their research excellence and leadership, but their good judgement and commitment to service. You'll find their pictures and brief bios at the end of this editorial. My sincerest thanks to all of these new AEs for their commitment to our journal! I'm really looking forward to working closely with them.

I'd like to take this opportunity to thank the following outgoing AEs for their service to our journal: Richard Bowden, Jason Corso, Gal Elidan, Ali Farhadi, Vittorio Ferrari, Jiaya Jia, Ivan Laptev, Svetlana Lazebnik, Erik Learned-Miller, Zhouchen Lin, Bryan Morse, Deva Ramanan, Clayton Scott, Noah Snavely, Massimo Tistarelli, Andrea Vedaldi, Ying Wu, Dong Xu, Huan Xu, Ruigang Yang, Jieping Ye, and Jingyi Yu. I'm really grateful to all of the above AEs for their service to our journal; as a past AE myself, I know how time-consuming the role can be at the most inconvenient times. My sincerest thanks to them all!

I'd like to once again express my thanks to the many individuals that have not only established *IEEE Transactions on Pattern Analysis and Machine Intelligence* as an elite journal, but are instrumental in its day-to-day operations. First, I could not do this job without the help of my stellar Associate Editors-in-Chief (AEICs): Christoph Lampert, Dale Schuurmans, and Jun Zhu on the machine learning side, and Kristen Grauman, Kyoung Mu Lee, Bernt Schiele, and René Vidal, on the computer vision side. I'm very grateful to them all, and feel really lucky to be able to work with such an exceptionally talented team of individuals. Second, while I thanked our outgoing and incoming AEs, I'd like to also thank the bulk of our active AE cohort – the foundation of our editorial board. The heavy lifting behind selecting papers for our journal is done by our

Digital Object Identifier no. 10.1109/TPAMI.2019.2958194

AEs, and I'm very grateful to them all for their outstanding service and commitment to our journal. Third, I'd like to thank the members of the *IEEE Transactions on Pattern Analysis and Machine Intelligence* Advisory Board. Their collective experience continues to be a valuable resource to me. Fourth, I'd like to thank our entire community of reviewers, without whom we would have no journal; I'm very grateful for their expertise, judgement, and commitment. Finally, I'd like to offer a special thanks to Joyce Arnold who continues to be of enormous help to me. I'd also like to thank the many other very helpful individuals at the IEEE Computer Society who have been assisting me this past year, including Hilda Carmen, Jennifer Carruth, Carrie Clark, Kathy Santa Maria, Pilar Etuk, Christine Shaughnessy, and Kimberly Sperka. My sincerest thanks to them all!

Sven Dickinson
Editor in Chief

Amr Ahmed is a senior staff research scientist at Google. He received the M.Sc and PhD degrees from the School of Computer Science, Carnegie Mellon University, in 2009 and 2011, respectively. He received the Best Paper Award at KDD 2014 , Best Paper Award at WSDM 2014, 2012 ACM SIGKDD Doctoral Dissertation Award, and Best Paper Award (runner-up) at WSDM 2012. He has co-chaired the WWW'18 track on Web Content Analysis and served as an area chair for IJCAI 2019, SIGIR 2019, SIGIR 2018, ICML 2018, ICML 2017, KDD 2016, WSDM 2015, ICML 2014, and ICDM 2014. His research interests include large-scale machine learning, deep learning, data/web mining, user modeling, personalization, social networks, and content analysis.

Xiang Bai received the BS, MS, and PhD degrees from the Huazhong University of Science and Technology (HUST), Wuhan, China, all in electronics and information engineering, in 2003, 2005, and 2009, respectively. He is currently a professor and vice dean with the School of Electronic Information and Communications, HUST. His research interests include object recognition, shape analysis, and document analysis. He has published more than 150 research papers. He was invited as a keynote speaker of ICDAR 2017. He is the recipient of 2019 IAPR/ICDAR Young Investigator Award for his outstanding contributions to scene text understanding. He serves as an area chair or senior PC member of CVPR20, AAAI20, ICDAR19, CVPR19, AAAI19, etc., and received AAAI-2019 Outstanding SPC Award. He is a senior member of the IEEE.

Amit Roy-Chowdhury received the PhD degree in electrical and computer engineering from the University of Maryland, College Park (UMCP), and joined the University of California, Riverside (UCR), in 2003, where he is a Bourns family faculty fellow and professor of electrical and computer engineering, director of the Center for Research in Intelligent Systems, and a cooperating faculty in the Department of Computer Science and Engineering. He leads the Video Computing Group at UCR, working on foundational principles of computer vision, image processing, and vision-based statistical learning, with applications in cyber-physical, autonomous, and intelligent systems. His current research projects include weakly-supervised and adaptive methods for video analysis, computer vision in resource-constrained environments, combined image and text analysis, image forensics, and bioimage analysis. his research has been supported by various US government agencies and private industries. He has published more than 200 papers in peer-reviewed journals and top conferences. He is the first author of the book "Camera Networks: The Acquisition and Analysis of Videos Over Wide Areas." His work on face recognition in art was featured widely in the news media, including a PBS/National Geographic documentary and in *The Economist*. He is a senior associate editor of the *IEEE Transactions on Image Processing* and on program committees of the main conferences in his area. His students have been first authors on multiple papers that received best paper awards at major international conferences, including ICASSP and ICMR. He is a fellow of the IEEE and IAPR, and received the Doctoral Dissertation Advising/Mentoring Award 2019 at UCR.

Dima Damen received the PhD degree from the University of Leeds, in 2009. She is a reader (associate professor) in computer vision with the University of Bristol, United Kingdom. Her research in computer vision and machine learning focuses on the automatic understanding of object interactions, actions and activities using static and wearable visual (and depth) sensors. She was selected as a Nokia research collaborator, in 2016, and as an outstanding reviewer in ICCV17, CVPR13, and CVPR12. She has more than 80 peer-reviewed publications, and frequently publishes in leading venues: CVPR, ICCV, and ECCV. Her recent works highlight neglected shortcomings of current approaches (e.g., reliability to action boundaries, time-reversal of actions and action completion), and introduce novel research questions (skill determination in video, fine-grained action retrieval and dual-domain LSTM). In 2018, her group released EPIC-Kitchens, the largest video dataset from wearable cameras. For more information, please visit http://dimadamen.github.io.

Kosta Derpanis received the honors BSc degree in computer science from the University of Toronto, Canada, in 2000, and the MSc (supervisors Prof. John Tsotsos and Prof. Richard Wildes) and PhD (supervisor Prof. Richard Wildes) degrees in computer science from York University, Canada, in 2003 and 2010, respectively. For his dissertation work, he received the Canadian Image Processing and Pattern Recognition Society (CIPPRS) Doctoral Dissertation Award 2010 Honourable Mention. Subsequently, he was a postdoctoral researcher in the GRASP Laboratory at the University of Pennsylvania under the supervision of Prof. Kostas Daniilidis. In 2012, he joined the Department of Computer Science at Ryerson University, Toronto, and is now an associate professor. He currently is also a research scientist with the Samsung AI Centre Toronto. His main research interests include computer vision with an emphasis on motion analysis and human motion understanding, and related aspects in image processing and machine learning.

Giovanni Maria Farinella is an associate professor with the Department of Mathematics and Computer Science, University of Catania, Italy. His research interests include in the fields of computer vision and machine learning. His group's most recent effort is related to egocentric perception. He is author of more than 100 papers in international book chapters, journals, and conference proceedings, and inventor of five patents involving industrial partners. He serves as a reviewer and on the board program committee for major international journals and conferences. He has served as an area chair of ICCV 2017-19. He founded (in 2006) and currently directs the International Computer Vision Summer School (ICVSS). He also founded (in 2014) and currently directs the Medical Imaging Summer School (MISS). He was awarded the IEEE PAMI Mark Everingham Prize in October 2017. For more information, please visit http://www.dmi.unict.it/farinella.

Ryan Farrell received the BS degree from the University of California, Berkeley, in 2001, and the MS and PhD degrees from the University of Maryland, College Park, in 2006 and 2011, respectively. He is currently an associate professor of computer science with Brigham Young University (BYU) in Provo, Utah. His research interests include computer vision, on topics ranging from fine-grained recognition, few-shot learning, and pose estimation to wildlife monitoring, edge AI and AI for Conservation. He has served on numerous program committees and as an area chair for top vision conferences such as CVPR, ICCV, and ECCV. He is currently serving as a program chair for WACV 2020 and will serve as PC also for ICVGIP 2020. He has co-organized six workshops to date on fine-grained recognition (FGVC – FGVC6, and FGVC7 proposed for 2020).

Yasutaka Furukawa received the PhD degree in computer science from the Department of the University of Illinois, Urbana-Champaign, in 2008. He is an associate professor with the the School of Computing Science at Simon Fraser University (SFU), Canada. Prior to joining SFU, in 2017, he was an assistant professor at Washington University in St. Louis, USA, and a software engineer at Google. He received the Best Student Paper Award at the premier computer vision conference, European Conference on Computer Vision 2012, Best Paper Award at the International Conference on 3D Vision 2013, the NSF CAREER Award in 2015, and CS-CAN Outstanding Young CS Researcher Award 2018. He co-chaired the International Conference on 3D Vision in 2013 and 2017, respectively. He is an associate editor for the *International Journal of Computer Vision*, *Computer Vision and Image Understanding*, *Image and Vision Computing*, and the *IPSJ Transactions on Computer Vision and Applications*. He is a visiting researcher at the National Institute of Informatics (Japan) and National Institute of Advanced Industrial Science and Technology (Japan). He is a pioneer in 3D reconstruction research. His multi-view stereo (MVS) algorithm has been recognized as the best 3D reconstruction algorithm from images. His paper is the most cited MVS article in the last decade (i.e., 2,791 citations of 2019.8). His open-source software has been used in numerous academic and industrial settings, including Google, Industrial Light & Magic, and Weta-digital for real production purposes and even the conservation of wildlife.

James Glass received the SM and PhD degrees in electrical engineering and computer science from MIT. His research has focused on automatic speech recognition, unsupervised speech processing, and spoken language understanding. He is a senior research scientist with MIT, where he leads the Spoken Language Systems Group in the Computer Science and Artificial Intelligence Laboratory. He is also a member of the Harvard-MIT Health Sciences and Technology Faculty. He is a fellow of the IEEE and the International Speech Communication Association, and is currently an associate editor for *Computer, Speech, and Language*, and the *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

András György received the MSc (Eng.) degree (with distinction) in technical informatics from the Technical University of Budapest, in 1999, the MSc (Eng.) degree in mathematics and engineering from Queen's University, Kingston, ON, Canada, in 2001, and the PhD degree in technical informatics from the Budapest University of Technology and Economics, in 2003. He was a visiting research scholar in the Department of Electrical and Computer Engineering, University of California, San Diego, USA, in the spring of 1998. In 2002-2011, he was with the Computer and Automation Research Institute of the Hungarian Academy of Sciences, where, from 2006, he was a senior researcher and the head of the Machine Learning Research Group. In 2003-2004, he was also a NATO Science fellow in the Department of Mathematics and Statistics, Queen's University. He also held a part-time research position at GusGus Capital Llc., Budapest, Hungary, in 2006-2011. In 2012-2015, he was a researcher in the Department of Computing Science, University of Alberta, Edmonton, AB, Canada. In 2015-2019, he was a senior lecturer in the Department of Electrical and Electronic Engineering of Imperial College London, London, UK. Since 2018, he has been a research scientist at Deepmind, London, UK. His research interests include machine learning, statistical learning theory, online learning, adaptive systems, information theory, and optimization. He has published more than 80 papers in peer-reviewed journals and conferences on these topics, and regularly serves as a senior program committee member or area chair of leading conferences in machine learning and information theory (such as NeurIPS, COLT, ALT, IJCAI, and IEEE ISIT). He received the Gyula Farkas Prize of the János Bolyai Mathematical Society, in 2001, the Academic Golden Ring of the President of the Hungarian Republic, in 2003, and a Reviewer Award at the 2015 International Conference on Machine Learning.

Timothy Hospedales received the BA degree in computer science from the University of Cambridge, in 2002, and the PhD degree in neuro-informatics from the University of Edinburgh, in 2008. He is currently reader (UK associate professor) at the University of Edinburgh, principal researcher at Samsung AI Research Centre Cambridge, and Alan Turing Institute fellow. He has served as an area chair for several major events (ICCV, ECCV, AAAI, ACL) and program chair of BMVC 2018.

Brian Kingsbury (M'97-S'09-F'18) received the BS degree in electrical engineering from Michigan State University, and the PhD degree in computer science from the University of California, Berkeley. He is a distinguished research staff member in IBM Research AI and manager of the Speech Technologies research group with the T. J. Watson Research Center in Yorktown Heights, NY. His research interests include deep learning, optimization, large-vocabulary speech transcription, and keyword search. From May 2012 until November 2016, he was co-PI and technical lead for LORELEI, an IBM-led consortium participating in the IARPA Babel program. He has contributed to IBM's entries in numerous competitive evaluations of speech technology, including Switchboard, SPINE, EARS, Spoken Term Detection, and GALE. He has served as a member of the Speech and Language Processing Technical Committee of the IEEE Signal Processing Society (2009-2011); as an ICASSP speech area chair (2010-2012); an associate editor for the *IEEE Transactions on Audio, Speech, and Language Processing* (2012-2016); and as a program chair for the International Conference on Representation Learning (2014-2016). He is an author or co-author on more than 100 publications on speech recognition, machine learning, and VLSI design. He is a fellow of the IEEE.

Ajay Kumar received the PhD degree from the University of Hong Kong, Hong Kong, in 2001. He was an assistant professor with the Department of Electrical Engineering, IIT Delhi, Delhi, India, from 2005 to 2007. He is currently working as a professor with the Department of Computing, The Hong Kong Polytechnic University. His research interests include biometrics and computer vision for industrial inspection. He holds seven US patents, and has authored a book entitled 'Contactless 3D Fingerprint Identification.' He served on the IEEE Biometrics Council as the vice president (Publications) in 2011-2015. He was an editorial board member of the *IEEE Transactions Information Forensics and Security* from 2010 to 2013. He has served as the area chair, program chair, and general chair for several international conferences and workshops in the field of his research interest. He is a fellow of the IEEE and a fellow of IAPR.

Simon Lacoste-Julien received the BSc degree in mathematics, physics, and computer science from McGill University, and the PhD degree in computer science from the University of California, Berkeley, in 2009. He is currently an associate professor with the Department of Computer Science and Operations Research at the Université de Montréal, a co-founding member of Mila, and heads part-time the SAIT AI Lab Montreal from Samsung. Before that, he completed a post-doctoral fellowship at the University of Cambridge as well as at Inria Paris, and was an Inria researcher in the Department of Computer Science at the Ecole Normale Supérieure (ENS) in Paris. His research interests include machine learning, optimization and statistics with applications to computer vision and natural language processing. He has published more than 40 scientific publications in machine learning and has served as an area chair for all the major machine learning or vision conferences. He received a Google Focused Research Award in 2016 and a CIFAR AI Chair in 2018.

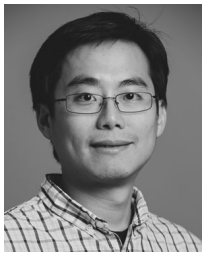

Dr. Lihong Li received the PhD degree in computer science from Rutgers University, in 2009, where he is currently a staff research scientist with Google Research. Before joining Google, he held research positions (researcher, senior researcher, and principal researcher) in Yahoo! Research and Microsoft Research. His main research interests include reinforcement learning, contextual bandits, online learning, active learning, and other related problems in AI. He has published more than 50 technical papers at top-tier AI journals and conferences, including several paper awards at ICML, AISTATS, and WSDM. His work has found applications in recommendation, advertising, Web search, and conversation systems. He has served as area chair or senior program committee member at major AI/ML conferences, such as AAAI, AISTATS, ICLR, ICML, IJCAI, and NeurIPS.

Ce Liu received the BE and ME degrees from Tsinghua University, in 1999 and 2002, respectively. He received the PhD degree from the MIT Department of Electrical Engineering and Computer Science, in 2009. He is a staff research scientist with Google Research, conducting research in the area of computer vision, computer graphics, and machine learning. He worked at Microsoft Research Asia from 2002 to 2003, and Microsoft Research New England from 2009 to 2014. He received the Best Student Paper Award at NIPS 2006 and CVPR 2009, and the Best Paper Award Honorable mention at CVPR 2019. He is a recipient of TPAMI Young Research Award in 2016. He has been serving as the area chair for CVPR/ICCV/ECCV/NeurIPS, and is serving as a co- program chair for CVPR 2020.

Tie-Yan Liu is an assistant managing director of Microsoft Research Asia, a fellow of the IEEE, and a distinguished member of the ACM. He is also an adjunct/honorary professor with Carnegie Mellon University (CMU), the University of Nottingham, and Tsinghua University. He is well known for his pioneer work on machine learning for information retrieval (learning to rank), and recently he has done impactful research on deep learning, reinforcement learning, and distributed learning. He published more than 200 papers in top conferences and journals, with tens of thousands of citations. He has been invited to serve as general chair, program committee chair, local chair, or area chair for a dozen of top conferences including WWW/WebConf, SIGIR, KDD, ICML, NIPS, IJCAI, AAAI, ACL, ICTIR, as well as associate editor of the *ACM Transactions on Information Systems*, *ACM Transactions on the Web*, and *Neurocomputing*. He won the Best (Student) Paper Award at SIGIR (2008) and ACML (2018), Most Cited Paper Award for the *Journal of Visual Communications and Image Representation* (2004-2006), the most cited Chinese researcher award by Elsevier (2017, 2018), and the Most Influential Scholar Award by AMiner (2007-2017). His team released LightGBM in 2017, which has become one of the most popularly used machine learning tools in Kaggle and KDD Cup; his team helped Microsoft achieve human parity in machine translation in 2018 and won eight champions in the WMT machine translation contest in 2019; his team also built the world-best Mahjong AI, named Suphx, which achieved 10 DAN on the Tenhou Mahjong platform in mid 2019.

Wei Liu (M'14-SM'19) received the PhD degree in electrical engineering and computer science from Columbia University, in 2012. He is currently a Distinguished Scientist of Tencent, China, and the director of Computer Vision Center at Tencent AI Lab. Prior to that, he has been a research staff member of the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA, from 2012 to 2015. He has long been devoted to research and development in the fields of machine learning, computer vision, pattern recognition, information retrieval, big data, etc. He has published extensively in these fields with more than 200 peer-reviewed technical papers. His research works win a number of awards and honors, such as the 2011 Facebook Fellowship, the 2013 Jury Award for Best Thesis of Columbia University, 2016 and 2017 SIGIR Best Paper Award Honorable Mentions, and 2018 "AI's 10 To Watch" honor. He currently serves on the editorial boards of the *IEEE Transactions on Neural Networks and Learning Systems*, *IEEE Transactions on Circuits and Systems for Video Technology*, *IEEE Access*, *Pattern Recognition*, etc. He is an elected member of the International Statistical Institute (ISI) and a fellow of the British Computer Society (BCS).

Chen Change Loy (S'06-M'10-SM'17) received the PhD degree in computer science from the Queen Mary University of London, in 2010. He is an associate professor with the School of Computer Science and Engineering, Nanyang Technological University (NTU). Prior to joining NTU, he served as a research assistant professor with the Department of Information Engineering, The Chinese University of Hong Kong, from 2013 to 2018. His research interests include computer vision and deep learning. He received the Nanyang Associate Professorship (NTU Career Award), in 2019. He serves as an associate editor of the *International Journal of Computer Vision* and *IET Computer Vision*. He also serves/served as an area chair of CVPR 2019, BMVC 2019, ECCV 2018, and BMVC 2018. He is a senior member of the IEEE.

Michael Maire received the PhD degree in computer science from the University of California, Berkeley, in 2009. He is an assistant professor of computer science with the University of Chicago. Previously, he was a research assistant professor at the Toyota Technological Institute, Chicago, following a time as a postdoctoral scholar in the Department of Electrical Engineering at Caltech. His current research interests include computer vision and deep learning. His publications cover multiple areas, including image segmentation, spectral clustering, object recognition, large-scale dataset construction, self-supervised learning, and neural network architectures. According to Google Scholar, he has multiple highly cited papers, including work on image segmentation and the Common Objects in Context (COCO) dataset. This dataset has served as a primary benchmark of modern object detection systems, with yearly workshop competitions featured at major computer vision conferences.

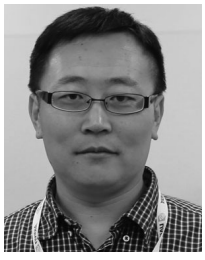

Deyu Meng received the BSc, MSc, and PhD degrees from Xi'an Jiaotong University, Xi'an, China, in 2001, 2004, and 2008, respectively. He is currently a professor with the Institute for Information and System Sciences, School of Mathematics and Statistics, Xi'an Jiaotong University. From 2012 to 2014, he took his two-year sabbatical leave at Carnegie Mellon University. He has published more than 100 peer-reviewed papers on international journals like *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *International Journal of Computer Vision*, *Transactions on Image Processing*, *Transactions on Knowledge and Data Engineering*, *Transactions on Neural Networks and Learning Systems* and top conferences, such as CVPR, ICCV, ECCV, AAAI, IJCAI, ICML, NeurIPS, and ACM MM. He has served as SPC or PC members of multiple top conferences, such as CVPR, ICCV, ECCV, AAAI, IJCAI, ICML, NeurIPS, and ACM MM. His current research interests include self-paced learning, noise/loss modeling, and tensor analysis.

Vlad Morariu received the BS and MS degrees from the Pennsylvania State University, in 2005, with Professor Octavia I. Camps as his thesis advisor. He received the PhD degree from the University of Maryland, in 2010, with Professor Larry S. Davis as his advisor. He is a senior research scientist with Adobe Research. His research interests include combining computer vision, natural language, machine learning, and artificial intelligence techniques to develop rich visual and linguistic models that enable intelligent reasoning about images, videos, and related linguistic descriptions. After completing his doctoral studies, he continued as a postdoctoral researcher and then as a research scientist at the University of Maryland until 2018, when he joined Adobe Research. He has co-authored more than 40 peer-reviewed publications including more than 20 in top-tier computer vision conferences (ICCV, CVPR, ECCV, and NeurIPS). He served as area chair for WACV 2017 and 2020, and he served as program committee member for many computer vision and AI conferences, including *Computer Vision and Pattern Recognition*, ICCV, ECCV, AAAI, and IJCAI.

Cheng Soon Ong received the BE (information systems) and BSc (computer science) degrees from the University of Sydney, Australia, and the PhD degree in computer science from the Australian National University, in 2005. He is a principal research scientist with the Machine Learning Research Group, Data61, CSIRO. He is also an adjunct associate professor with the Australian National University. He was a postdoc at the Max Planck Institute of Biological Cybernetics and the Fredrich Miescher Laboratory in Tübingen, Germany. From 2008 to 2011, he was a lecturer in the Department of Computer Science at ETH Zurich, and in 2012 and 2013 he worked in the Diagnostic Genomics Team at NICTA in Melbourne. Since 2014, he has been researching ways to enable scientific discovery by extending statistical machine learning methods with the Machine Learning Research Group in Data61, CSIRO, in Canberra. Prior to his PhD degree, he researched and built a search engine and Bahasa Malaysia technologies at Mimos Berhad, Malaysia.

John Paisley received the BS, MS, and PhD degrees in electrical and computer engineering from Duke University. He was then a postdoctoral researcher in the Computer Science Departments at Princeton University and the University of California at Berkeley. He is currently an associate professor with the Department of Electrical Engineering at Columbia University, where he is also a member of the Data Science Institute. His research interests include the area of machine learning, focusing on Bayesian models and inference techniques, as well as applied deep learning, with applications to data science problems involving text, audio, and images.

Thomas Pock received the both MSc degrees during 1998-2004, and the PhD degree, in 2008 in computer engineering (Telematik) from the Graz University of Technology. After a post-doctoral position at the University of Bonn, he returned to the Technical University of Graz, where he worked as an assistant professor in the Institute of Computer Graphics and Vision. In 2013, he was awarded the START Prize of the Austrian Science Fundation (FWF) and the German Pattern Recognition Award of the German Society for Pattern Recognition (DAGM) and in 2014 a start grant from the European Research Council (ERC). He has been a professor of computer science at the Graz University of Technology since June 2014 and a board member of the Austrian Science Fund (FWF) since April 2019. His research interests include the development of mathematical models and machine learning techniques for imaging and vision, as well as the development of efficient convex and non-smooth optimization algorithms.

Liva Ralaivola has been full professor of computer science at Aix-Marseille University (AMU) since 2011 (on leave since September 2019), member of Institut Universitaire de France (2016 – 2019) and research director in AI at Criteo since September 2018. He received the Ph.D. degree in computer science from Université Paris 6, in 2003, and the Habilitation Diriger des Recherches (HDR) in computer science from AMU, in 2010. His research focuses on statistical and algorithmic aspects of machine learning with a focus on theoretical issues and algorithmic processes (risk of predictors, concentration inequalities in non-IID settings, bandit problems, greedy methods).

Xiaofeng Ren received the PhD degree from U.C. Berkeley. He is currently the chief scientist of Amap (AutoNavi), mapping and navigation subsidiary of the Alibaba Group. He holds affiliate faculty positions at the University of Washington and Zhejiang University. He is a computer vision researcher whose work spans both academic and applied research and product development. His research topics include boundary detection and image segmentation, object detection and recognition, scene understanding, motion and video analysis, and activity recognition. He served as area chairs for CVPR, ICCV, and AAAI. He co-organized series of workshops on RGB-D perception and its applications, and on Egocentric (First-Person) Vision. Previously, he held various computer vision R&D positions: chief scientist at the Institute of Data Science and Technologies at Alibaba, senior principal scientist at Amazon (Amazon Go), senior research scientist at Intel Labs (Seattle), and research assistant professor at TTI-Chicago.

Irina Rish is an associate professor in the Computer Science and Operations Research Department at the Université de Montréal (UdeM) and a core member of MILA—Quebec AI Institute. Before joining UdeM and MILA in 2019, she was a research scientist at the IBM T.J. Watson Research Center, where she worked on various projects at the intersection of neuroscience and AI. Specializing in machine learning, neural data analysis, and bio-inspired computing, she focuses her current research on continual learning and optimization in deep neural networks, sparse modeling and probabilistic inference, dynamical systems, and information theory.

Yaser Sheikh directs the Facebook Reality Lab in Pittsburgh, which is devoted to achieving photorealistic social interactions in AR and VR. He is also an adjunct professor with Carnegie Mellon University, where he was part of the Robotics Institute faculty for over a decade. His research broadly focuses on machine perception and rendering of social behavior, spanning sub-disciplines in computer vision, computer graphics, and machine learning. With colleagues and students, he has won the Honda Initiation Award (2010), Popular Science's "Best of What's New" Award, Best Student Paper Award at CVPR (2018), best paper finalists at (CVPR 2019), best paper awards at WACV (2012), SAP (2012), SCA (2010), ICCV THEMIS (2009), Best Demo Award at ECCV (2016), and he received the Hillman Fellowship for Excellence in Computer Science Research (2004). He has served as a senior committee member at leading conferences in computer vision, computer graphics, and robotics including SIGGRAPH (2013, 2014), CVPR (2014, 2015, 2018), ICRA (2014, 2016), ICCP (2011), and served as an associate editor of *Computer Vision and Image Understanding*. His research has been featured by various media outlets including The New York Times, BBC, MSNBC, Popular Science, and in technology media, such as WIRED, The Verge, and New Scientist.

Suvrit Sra received the PhD degree in computer science from the University of Texas at Austin. He is an associate professor in the Department of Electrical Engineering and Computer Science and IDSS, which he joined in January 2018. Prior to this, he was a principal research scientist in the MIT Laboratory for Information and Decision Systems (LIDS). Before coming to LIDS, he was a senior research scientist at the Max Planck Institute for Intelligent Systems, in Tübingen, Germany. During this time, he was also a visiting faculty member at the University of California, Berkeley (EECS Department) and Carnegie Mellon University (Machine Learning Department). His research bridges a variety of mathematical topics including optimization, matrix theory, differential geometry, and probability with machine learning. His recent work focuses on the foundations of geometric optimization, an emerging subarea of nonconvex optimization where geometry (often non-Euclidean) enables efficient computation of global optimality. More broadly, his work encompasses a wide range of topics in optimization, especially in machine learning, statistics, signal processing, and related areas. His work has won several awards at machine learning conferences, the 2011 "SIAM Outstanding Paper" Award, faculty research awards from industry, and an NSF CAREER Award. In addition, he founded (and regularly co-chairs) the popular OPT "Optimization for Machine Learning" series of Workshops at the Conference on Neural Information Processing Systems. He has also edited a well-received book with the same title (MIT Press, 2011).

Ping Tan received the PhD degree from the Hong Kong University of Science and Technology (HKUST), in 2007. He is an associate professor with Simon Fraser University (SFU), Canada. Before that, he was an associate professor in the National University of Singapore (NUS). His research interests include computer vision and computer graphics. He serves an associate editor for the *International Journal of Computer Vision (IJCV)*, *Computer Graphics Forum (CGF)*, and *Machine Vision and Applications (MVA)*, and serves as an area chair for IEEE International Conference on Computer Vision and Pattern Recognition (CVPR), ACM SIGGRAPH, ACM SIGGRAPH Asia, and IEEE International Conference on Intelligent Robots and Systems. He received an NSERC DAS award in 2014.

Christian Theobalt received the MSc degree in artificial intelligence from the University of Edinburgh, the Diploma MS degree in computer science from Saarland University, and the PhD (Dr.-Ing.) degree from the Max-Planck-Institute for Informatics. He is a professor of computer science and the head of the research group "Graphics, Vision, & Video" at the Max-Planck-Institute (MPI) for Informatics, Saarbrücken, Germany. He is also a professor of computer science with Saarland University, Germany. From 2007 until 2009, he was a visiting assistant professor in the Department of Computer Science at Stanford University. In his research, he looks at algorithmic problems that lie at the intersection of computer graphics, computer vision, and machine learning, such as static and dynamic 3D scene reconstruction, marker-less motion and performance capture, virtual and augmented reality, computer animation, appearance and reflectance modelling, intrinsic video and inverse rendering, machine learning for graphics and vision, new sensors for 3D acquisition, advanced video processing, as well as image and physically-based rendering. He is also interested in using reconstruction techniques for human computer interaction. For his work, he received several prestigious awards, including the Otto Hahn Medal of the Max-Planck Society in 2007, EUROGRAPHICS Young Researcher Award in 2009, German Pattern Recognition Award 2012, and Karl Heinz Beckurts Award in 2017 (one of Germany's highest science awards). He received two ERC grants, one of the most prestigious and competitive individual research grants in Europe: An ERC Starting Grant in 2013 and an ERC Consolidator Grant in 2017. He is also a fellow of ELLIS—the European Laboratory for Learning and Intelligent Systems. In 2015, he was elected as one of the top 40 innovation leaders under 40 in Germany by the business magazine *Capital*. He is also a co-founder of an award-winning spin-off company from his group—www.thecaptury.com—that is commercializing one of the most advanced solutions for marker-less motion and performance capture.

Radu Timofte received the Dipl. Eng. from the Technical University of Iasi, Romania, in 2006, the MSc degree the University of Eastern Finland, in 2007, and the PhD degree in electrical engineering from the KU Leuven, Belgium, in 2013. He is a lecturer and research group leader with the Computer Vision Laboratory, at ETH Zurich, Switzerland. He serves as a reviewer for top journals (such as *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *IEEE Transactions on Image Processing*, *International Journal of Computer Vision*, *IEEE Transactions on Neural Networks and Learning Systems*, *Transactions on Circuits and Systems for Video Technology*, *Computer Vision and Image Understanding*, and *Pattern Recognition*) and conferences (ICCV, CVPR, ECCV, NeurIPS, and ICLR), as area editor for Elsevier's *CVIU* (from 2017) journal, and as associate editor for SIAM's *SIIMS* (from 2020). He served as area chair for ACCV 2018, ICCV 2019, and as SPC for IJCAI 2019 and 2020. He received a NIPS 2017 Best Reviewer Award. His work received several awards, including a Best Scientific Paper Award at ICPR 2012, the Honorable Mention Award at FG 2017, the Best Student Paper Award at BMVC 2019, and his team won a number of challenges including traffic sign detection (IJCNN 2013), apparent age estimation (ICCV 2015), and real world super-resolution (ICCV 2019). He is co-founder of Merantix and co-organizer of NTIRE, CLIC, PIRM, and AIM events. His current research interests include deep learning, augmented perception, domain translation, image/video compression, manipulation, restoration, and enhancement.

Lorenzo Torresani received the Laurea degree in computer science with summa cum laude honors from the University of Milan (Italy), in 1996, and the MS and PhD degrees in computer science from Stanford University, in 2001 and 2005, respectively. He is an associate professor with tenure in the Computer Science Department at Dartmouth College and a research scientist at Facebook AI Research (FAIR) and Applied Research (FAIAR). In the past, he has worked at several industrial research labs including Microsoft Research, Like.com, and Digital Persona. His research interests include computer vision and deep learning. He is the recipient of several awards, including a CVPR Best Student Paper Prize, a National Science Foundation CAREER Award, a Google Faculty Research Award, three Facebook Faculty Awards, and a Fulbright US Scholar Award.

Chong Wang received the PhD degree from Princeton University. He is currently the head of the applied machine learning (AML) research team at Seattle, ByteDance Inc. he is interested in machine learning and their applications to various problems. Previously, he worked as a research scientist at Google and Microsoft Research. For more information, please visit https://chongw.github.io

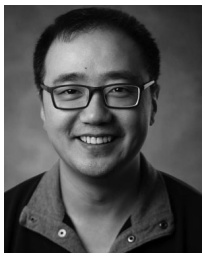

Jue Wang received the BE and MSc degrees from the Department of Automation, Tsinghua University, Beijing, China, in 2000 and 2003, respectively, and the PhD degree in electrical engineering from the University of Washington in Seattle WA, USA in 2007. He is currently a senior director with Megvii Research. He was a principle research scientist at Adobe Research from 2007 to 2017. He has published more than 100 peer-reviewed research articles in the areas of computer vision, computer graphics, and HCI, and holds more than 60 international patents. He was a co-organizer for the IEEE International Conference on Multimedia and Expo 2016 and IEEE International Conference on Computational Photography 2012. He was an associate editor for *The Visual Computer and International Journal of Computer Games Technology*. He was the recipient of the Microsoft Research Fellowship and the Yang Research Award of the University of Washington in 2006. He is a board member of X-Lab at Tsinghua University in Beijing and Haitian scholar of the Dalian University of Technology. He is a senior member of the IEEE and ACM.

Richard P. Wildes received the PhD degree from the Massachusetts Institute of Technology, in 1989. Subsequently, he joined Sarnoff Corporation (now SRI) as a member of the technical staff in the Vision Technologies Lab. In 2001, he joined the Department of Electrical Engineering and Computer Science (EECS) at York University, Toronto, as an associate professor. Currently, he is associate director of York's Canada First Research Excellence Fund Supported Vision: Science to Applications (VISTA) program. He also has served as associate director of York's Centre for Vision Research (CVR) and chair of York's EECS Department. His research interests are in computational vision, especially video understanding and machine vision applications, as well as artificial intelligence. His work received award and honours including the IEEE D.G. Fink Prize Paper Award, Sarnoff Corporation Technical Achievement Award, twice giving invited presentations at the USA National Academy of Sciences, the Lassonde School of Engineering Graduate Mentorship Award and currently holding a York Research Chair, Tier I.

Christian Wolf received the MSc degree in computer science from TU Vienna, Austria, in 2000, and the PhD degree in computer science from INSA de Lyon, France, in 2003. In 2012, he obtained the habilitation diploma, also from INSA de Lyon. He has been an associate professor (Maître de Conférences, HDR) at INSA de Lyon and LIRIS, a CNRS laboratory, since September 2005. His research interests include machine learning and computer vision, especially the visual analysis of complex scenes in motion. His work puts an emphasis on modelling complex interactions of a large amount of variables: deep learning, structured models, and graphical models, and more recently the connections between machine learning and control. Between September 2017 and August 2019, he was on leave at INRIA, at the Chroma Work Group at the CITI laboratory. He is currently the national project coordinator of ANR/NSERC project "Deepvision" "Learning highly complex visual problems with deep structured models" (9/2016-3/2020) and the national coordinator of project ANR Delicio "Data and Prior, Machine Learning and Control" (2019-2023); he has coordinated the INSA-Lyon partner of ANR Canada (2007-2010) and PIA Interabot (2012-2016). He is member of the scientific committee of GDR IA and member of the directing committee of GDR ISIS and co-leader of the topic "*Machine Learning*." He has supervised 10 defended PhD theses.

Lei Zhang (F'18) received the BSc, MSc, and PhD degrees in 1995, 1998, and 2001, respectively. He joined the Department of Computing, the Hong Kong Polytechnic University, in 2006, and since 2017, he has been a chair professor in the department. He is a fellow of the IEEE, for his contributions in sparsity based image modeling and image quality assessment. His research interests include computer vision, image and video analysis, pattern recognition, biometrics, etc. He has published more than 200 papers in those areas. His publications have been cited more than 44,000 times in literature. He was listed as a "Clarivate Analytics Highly Cited Researcher" consecutively from 2015 to 2019. He is a senior associate editor of *IEEE Transactions on Image Processing*, and is/was an associate editor of the *SIAM Journal of Imaging Sciences*, *IEEE Transactions on CSVT*, *IEEE Transactions on SMC*, *Image Vision Computing*, etc.