

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 once again I've delayed it so that I can announce a large cohort of new Associate Editors (AE) – in fact, our largest yet! But before I begin, I would like to wish our entire readership, our editorial board, our staff, and our reviewers the very best in 2021. It's been such a difficult year this past year, and the toll the pandemic has taken on our families, our health, and our lives has been a significant one. For those who have lost friends and/or family to COVID-19, my sincerest condolences. I hope that the new vaccines and treatments will help protect us all, let us connect with each other again in person, and take us back to a normal world.

By the time this editorial appears, I will have started my fifth and final year in the role. It's once again been a very good year, with our impact factor nudging up slightly from 17.30 (2018) to 17.861 (2019), maintaining *TPAMI* as the top-ranked journal in all of computer science. Moreover, our submissions are up dramatically from last year; as of Nov. 13, 2020, we had 1571 submissions, compared to 1101 at the same time last year, an increase of over 40 percent in one year! Our key metrics are all trending in the right direction. For papers accepted in 2020, the average time from submission to first decision was 2.65 months, down from 2.72 months last year. For papers accepted in 2020, the average time from submission to acceptance was 9.49 months, down from 10.2 months last year. Finally, for papers accepted in 2020, the average time from submission to acceptance to publication on Xplore was 9.92 months, down from 10.6 months last year. My goal continues to be 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.

In 2020, we test drove a new model for a special issue of conference best papers that differs from how we handle our normal CVPR and ICCV best paper special issues. Teaming up with the organizers of the International Conference on Computational Photography (ICCP), all submissions to ICCP were in *TPAMI* format. The authors of the best papers chosen by the conference program chairs and vetted with two *TPAMI* AEs and myself were offered a choice: either fast-track your paper in a *TPAMI* ICCP Best Paper special issue available at the time of the conference and have your paper published *only* in the *TPAMI* special issue (and not in the conference proceedings) *or* choose to have your paper published only at the conference; in either case, best paper authors would be invited to present orally at the conference. For the *TPAMI* special issue, members of the ICCP Program Committee became *TPAMI* special issue guest editors, joined by two of our *TPAMI* AEs (my sincere thanks to Kristin Dana and Ko Nishino!). The experiment was considered so successful by both ICCP and *TPAMI* that we will continue the experiment in subsequent years through a MOU signed last Fall. I hope other elite conferences will explore this model with us.

Over the past two years, a big part of my strategy to reduce time to acceptance and publication was to appoint more AEs, which would 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, in 2018 we added another 22, and in 2019, we added another 39. I'm pleased to report that in the Fall of 2020, we added 49 new AEs, who I'll introduce shortly. I plan to continue to add more AEs over the coming year, as I continue to retire those that have exceeded their 2+2-year editorial board terms. This past year, Joyce Arnold and I evolved our strategy for flagging papers that require urgent attention, which should further reduce time to acceptance and publication.

As always, my favorite part of the editorial is introducing our new Editorial Board members and thanking our retiring board members. After 6 years of outstanding service to our journal, Dale Schuurmans has stepped down as one of our ML Associate Editors-in-Chief (AEIC). I can't thank him enough for his exceptional judgement, responsiveness, and exemplary commitment to our journal. And, I'm grateful to Dale for staying on in the role while we recruited his successor. I'm pleased to announce that Dale was replaced last Fall by Pradeep Ravikumar, who brings leadership and broad ML expertise to the role, including his prior service as a *TPAMI* AE. I'm really looking forward to working with Pradeep!

I'm pleased to announce a new cohort of 43 Associate Editors that joined in the Fall of 2020: Zeynep Akata, Pablo Arbelaez, Robert Babuska, Yuejie Chi, Rita Cucchiara, Florence d'Alché-Buc, Rina Dechter, Miroslav Dudik, Katerina Fragkiadaki, Pascale Fung, Bernard Ghanem, Tom Goldstein, Barbara Hammer, Zaid Harchaoui, James Hays, Wolfgang Heidrich, Judy Hoffman, Stefanie Jegelka, Shuiwang Ji, Evangelos Kalogerakis, Adriana Kovashka, Yong Jae Lee, Bastian Leibe, Victor Lempitsky, Xuelong Li, Yan Liu, Karen Livescu, Le Lu, Julien Mairal, Philippos Mordohai, Juan Carlos Niebles, Christopher Pal, Sinno Jialin Pan, Jonas Peters, Olga Russakovsky, Mathieu Salzmann, Dimitris Samaras, Cees Snoek, Justin Solomon, Kalyan Sunkavalli, Ivor Tsang, Anton van den Hengel, Martha White, Ole Winther, John Wright, Jianxin Wu, Tao Xiang, Sandra Zilles, and Wangmen Zuo. These individuals were 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 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: Tal Arbel, Tamara Berg, Gabriel Browstow, Gal Chechick, Pushmeet Kohli, R. Manmatha, Maja Pantic, Ilan Shimshoni, Josef Sivic, Gang Wang, and Kilian Weinberger. I'm really grateful to all 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 *TPAMI* 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, Jun Zhu, and Pradeep Ravikumar on the machine learning side, and Kristen Grauman, 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 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 *TPAMI* 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. She, too, has felt the burden of our dramatic growth in submissions, yet is always so positive and supportive. I'd also like to thank the many other very helpful individuals at the IEEE Computer Society and IEEE who have been assisting me this past year, including Robin Baldwin, Jennifer Carruth, Carrie Clark, Pilar Etuk, and Christine Shaughnessy. My sincerest thanks to them all!

Sven Dickinson
Editor-in-Chief



Zeynep Akata received the PhD degree in INRIA Rhone Alpes, France, in 2014. She is a professor of computer science with the Cluster of Excellence Machine Learning at the University of Tübingen in Germany. She worked as a post-doctoral researcher at the Max Planck Institute for Informatics (Germany) between 2014–2017, at UC Berkeley (USA) between 2016–2017, and as an assistant professor at the University of Amsterdam (The Netherlands) between 2017–2019. She received a Lise-Meitner Award for Excellent Women in Computer Science, in 2014 and an ERC Starting Grant, in 2019. Her research interests include multimodal learning in low-data regimes such as zero- and few-shot learning as well as explainable machine learning focusing on vision and language-based deep learning.



Pablo Arbeláez received the PhD degree with honors in applied mathematics from the Université Paris Dauphine, in 2005. He was a senior research scientist with the Computer Vision Group at UC Berkeley from 2007 to 2014. He is currently the director of the Center for Research and Formation in Artificial Intelligence at the Universidad de los Andes in Colombia, where he also holds a faculty position at the Department of Biomedical Engineering. His research interests include computer vision and machine learning, in which he has worked on a number of problems, including perceptual grouping, object recognition, and the analysis of biomedical images.



Robert Babuska received the MSc (Hons.) degree in control engineering from the Czech Technical University, in Prague, in 1990, and the PhD (cum laude) degree from TU Delft, The Netherlands, in 1997. He has had faculty appointments with the Czech Technical University, in Prague, and with the Electrical Engineering Faculty, TU Delft. Currently, he is a full professor of intelligent control and robotics at TU Delft, Faculty 3mE, Department of Cognitive Robotics. He has extensive experience with leading research project teams and, in 2012, he established the TU Delft Robotics Institute (robotics.tudelft.nl) of which was the scientific director until 2016. His research interests include adaptive and learning control, reinforcement learning, non-linear system identification, state estimation, and dynamic multi-agent systems. He has been involved in the applications of these techniques in various fields, ranging from process control to robotics and aerospace. He has published three research monographs, three edited books, and more than 280 scientific papers. He is the recipient of the 2009 Andrew P. Sage Award for the best paper published annually in the *IEEE Transactions on Systems, Man, and Cybernetics*. He served as the chairman of the IFAC Technical Committee on Cognition and Control and as associate editor of several archived journals, including *Automatica*, the *IEEE Transactions on Fuzzy Systems*, and *Engineering Applications of Artificial Intelligence*.



Yuejie Chi received the BEng degree from Tsinghua University, and the PhD degree from Princeton University, both in electrical engineering. She is an associate professor with the Department of Electrical and Computer Engineering, and a faculty affiliate with the Machine Learning Department at Carnegie Mellon University, where she holds the Robert E. Doherty Early Career Development Professorship. Her research interests include the theoretical and algorithmic foundations of data science, signal processing, machine learning, and inverse problems, with applications in sensing systems, broadly defined. Among others, she received the Presidential Early Career Award for Scientists and Engineers (PECASE), and the inaugural IEEE Signal Processing Society Early Career Technical Achievement Award for contributions to high-dimensional structured signal processing.



Rita Cucchiara is currently a full professor of computer engineering and science (ING-INF/05 Sistemi di Elaborazione dell'Informazione) within the Dipartimento di Ingegneria "Enzo Ferrari" (DIEF) at the UNIMORE, Università di Modena e Reggio Emilia, Italy. She is the director of the Interdepartmental Center of Research in Ict for Enterprise Softech-ICT of the Modena Technopole co-funded by the Emilia Romagna High Technology Network, under EU FESR programs. She is the director of master in "Visual computing and Multimedia Technology" of UNIMORE. She holds the Research Lab Imagleab, active in computer vision, pattern recognition and multimedia within the Dipartimento di Ingegneria "Enzo Ferrari". Since 2016, she has been the president of the Italian Association of Pattern Recognition, Learning and Computer vision GIRPR (Gruppo di ricercatori italiani in pattern recognition), affiliated to IAPR, and is a member of the IAPR Governing Board member. Since 2015, she has also been an advisory board member of the Computer Vision Foundation, CVF.



Florence d'Alché-Buc is currently a professor with the Image, Data, and Signal Department of Télécom Paris (National School of Telecommunications), a founding member of the Institut Polytechnique de Paris. Her research interests include large spectrum of topics in machine learning, including structured prediction, kernel methods, representation learning, robust and explainable learning and applications to health and genomics. She has led the ANR Excellence Laboratory of Information, and Communication Tehcnologies in Paris-Salcay University (DigiCosme) until end of 2019 and was a program co-chair of NeurIPS 2019, one of the top international conference in machine learning. She is also the research chair on *Data Science* and *Artificial Intelligence* for industry and services. She has more than 110 publications, in international journals, conferences, and co-edited three books in the domain.



Rina Dechter received the BS degree in mathematics and statistics from the Hebrew University, in Jerusalem, the MS degree in applied mathematics from the Weizmann Institute, and the PhD degree from UCLA. Her research centers on computational aspects of automated reasoning and knowledge representation including search, constraint processing, and probabilistic reasoning. She is a Chancellor's Professor of computer science at the University of California, Irvine. She is the author of *Constraint Processing* published by Morgan Kaufmann (2003), and of *Reasoning with Probabilistic and Deterministic Graphical Models: Exact Algorithms* published by Morgan and Claypool Publishers (2013, second ed. 2019). She has coauthored close to 200 research papers and has served on the editorial boards of: *Artificial Intelligence*, the *Constraint Journal*, *Journal of Artificial Intelligence Research* (JAIR), and *Journal of Machine Learning Research* (JMLR). She is a fellow of the American Association of Artificial Intelligence, since 1994, was a Radcliffe fellow during 2005–2006, received the 2007 Association of Constraint Programming (ACP) Research Excellence Award, and became an ACM fellow, in 2013. She served as a co-editor-in-chief of *Artificial Intelligence* from 2011 to 2018, and is the conference chair-elect of IJCAI-2022.



Miroslav Dudík received the PhD degree from Princeton, in 2007. He is a senior principal researcher of machine learning at Microsoft Research, NYC, currently focusing on contextual bandits, reinforcement learning, and algorithmic fairness. He is currently a co-creator of the MaxEnt Package for modeling species distributions, which is used by biologists around the world to design national parks, model impacts of climate change, and discover new species.



Katerina Fragkiadaki received the PhD degree from the University of Pennsylvania. She is an assistant professor with the Machine Learning Department, at Carnegie Mellon University, since 2016. She was a postdoctoral fellow at UC Berkeley and Google Research after that. Her work is on learning visual representations with little supervision and on combining spatial reasoning with deep visual learning. Her group develops algorithms for mobile computer vision, learning of physics, and common sense for agents that move around and interact with the world. Her work has been awarded with a Best PhD Thesis Award, in 2013, an NSF CAREER Award, in 2020, Google, Sony, and UPMC faculty research awards, since 2017. She has organized and co-organized more than six workshops in CVPR, on the topics of video understanding, perceptual grouping, and 3D vision. She has served as an area chair multiple times for CVPR, ICLR, ICML, and NeurIPS conferences.



Pascale Fung (Fellow, IEEE) is a professor with the Department of Electronic & Computer Engineering and Department of Computer Science & Engineering at the Hong Kong University of Science & Technology (HKUST), and a visiting professor at the Central Academy of Fine Arts in Beijing. She is an elected fellow of the Institute of Electrical and Electronic Engineers (IEEE) for her "contributions to human-machine interactions", and an elected fellow of the International Speech Communication Association for "fundamental contributions to the interdisciplinary area of spoken language human-machine interactions". She is the director of HKUST Centre for AI Research (CAiRE), an interdisciplinary research center on top of all four schools at HKUST. She co-founded the Human Language Technology Center (HLTC). She is an affiliated faculty with the Robotics Institute and the Big Data Institute at HKUST. She is the founding chair of the Women Faculty Association at HKUST. She is an expert on the Global Future Council, a think tank for the World Economic Forum. She represents HKUST on Partnership on AI to Benefit People and Society. She is on the Board of Governors of the IEEE Signal Processing Society. She is a member of the IEEE Working Group to develop an IEEE Standard - Recommended Practice for Organizational Governance of Artificial Intelligence. Her research team has won several best and outstanding paper awards at ACL, ACL, and NeurIPS workshops.



Bernard Ghanem received the bachelor's degree from the American University of Beirut (AUB), in 2005, and the MS and PhD degrees from the University of Illinois at Urbana-Champaign (UIUC), in 2010. He is currently an associate professor with the CEMSE Division and a research theme leader at the Visual Computing Center at the King Abdullah University of Science and Technology (KAUST) in Saudi Arabia. His research interests include computer vision and machine learning with emphasis on topics in video understanding, 3D recognition, and theoretical foundations of deep learning. His work has received several awards and honors, including the Henderson Graduate Award and CSE fellowship awards from UIUC, four best paper awards at workshops in CVPR2013, ECCV2018, CVPR2019, and ECCV2020, a two-year KAUST Seed Fund, and a Google Faculty Research Award, in 2015 (the first in the Middle East in Machine Perception). He has coauthored more than 100 peer reviewed conference and journal papers in his field as well as two issued patents. He has been invited to serve as area chair (AC) for CVPR 2018/2021, ICCV 2019, ICLR 2021, and AAAI 2021. For more information, please visit ivul.kaust.edu.sa and www.bernardghanem.com.



Tom Goldstein received the PhD degree in mathematics from UCLA. He is an associate professor of computer science with the University of Maryland. His research interests include the intersection of machine learning and optimization, and targets applications in computer vision, and signal processing. He works at the boundary between theory and practice, leveraging mathematical foundations, complex models, efficient hardware to build practical, and high-performance systems. He designs optimization methods for a wide range of platforms ranging from powerful cluster/cloud computing environments to resource limited integrated circuits and FPGAs. Before joining the faculty at Maryland, he was a research scientist at Rice University and Stanford University. He has been the recipient of several awards, including SIAM's DiPrima Prize, a DARPA Young Faculty Award, and a Sloan Fellowship.



Barbara Hammer received the PhD degree in computer science, in 1999, and the *her venia legendi* (permission to teach) degree, in 2003, both from the University of Osnabrueck, Germany. She is a professor for theoretical computer science with the CITEC Centre of Excellence at Bielefeld University, Germany. She was the head of an independent research group on the topic "Learning with Neural Methods on Structured Data" at the University of Osnabrueck, Germany. In 2004, she accepted an offer for a professorship at the Clausthal University of Technology, Germany, before moving to Bielefeld University in 2010. Her research interests include theory and algorithms in machine learning and neural networks and their application for technical systems, and the life sciences. She chaired the IEEE CIS Technical Committee on Data Mining and Big Data Analytics, in 2013/2014, and she is leading its task force on Data Analysis and Data Visualization. She is a chair of the special interest group on Neural Networks of the German Computer Science Society, and vice-chair of the German Neural Network Society. In 2016, she was elected as a member of the IEEE CIS Administrative Committee, she has been chairing the Distinguished Lecturer Program Committee of IEEE CIS, and she is currently chairing the IEEE CIS Neural Networks

Technical Committee. She has been an associate editor of the *IEEE Computational Intelligence Magazine*, she is an associate editor of the *IEEE Transactions on Neural Networks and Learning Systems*, the *Neural Processing Letters*, and the journal *Neurocomputing*.



Zaid Harchaoui received the doctoral degree from Telecom ParisTech. He is currently an associate professor with the Department of Statistics, an adjunct faculty at the Paul G. Allen School of Computer Science & Engineering, and a senior data science fellow at the eScience Institute at University of Washington. He is a member of the executive committee of IFDS and the NSF-TRIPODS Institute on the Foundations of Data Science. He was a CNRS fellow at the LTCI research unit of CNRS and Telecom ParisTech. He has held a visiting appointment at the Courant Institute of Mathematical Sciences at New York University in 2015 – 2016. He was a permanent researcher with the LEAR team at Inria from 2010 to 2015. He was also a member of the Microsoft Research - Inria Research Center. He was a postdoctoral fellow at Carnegie Mellon University, in 2009. He received the Inria Award for Scientific Excellence, the NIPS Reviewer Award, the Criteo Faculty Research Award, and the Google Faculty Research Award. He was appointed CIFAR associate fellow, member of the program "Learning in Machines and Brains" in 2015. He gave a tutorial on "Frank-Wolfe, greedy algorithms, and friends" at ICML 2014, on "Large-scale visual recognition" at CVPR 2013, and on "Machine learning for computer vision" at MLSS Kyoto 2015. He co-organized the "Future of AI" symposium

at New York University, the workshop on "Optimization for Machine Learning" at NIPS'14, the summer school "Foundations of Data Science" at University of Washington, in 2019, and the "Optimization and statistical learning" workshop (2013, 2015, 2017, and 2019) in Ecole de Physique des Houches (France). He served as area chair for ICML, ICLR, and NeurIPS.



James Hays received the PhD degree from Carnegie Mellon University. He is an associate professor of computing at the Georgia Institute of Technology, since fall 2015. Since 2017, he has also worked with Argo AI to create self-driving cars. Previously, he was the Manning assistant professor of computer science at Brown University. He was a postdoc at the Massachusetts Institute of Technology. His research interests span computer vision, computer graphics, robotics, and machine learning. His research often involves exploiting non-traditional data sources (e.g., internet imagery, crowdsourced annotations, thermal imagery, human sketches, and autonomous vehicle sensor data) to explore new research problems (e.g., global geolocalization, sketch to real, and hand-object contact prediction). He is the recipient of the NSF CAREER Award and Sloan Fellowship.



Wolfgang Heidrich received the PhD degree from the University of Erlangen, in 1999. He is a professor of computer science and the director of the Visual Computing Center at the King Abdullah University of Science and Technology (KAUST), as well as the interim leader for KAUST's AI Initiative. He joined KAUST, in 2014, after 13 years as a faculty member at the University of British Columbia. He worked as a research associate at the Computer Graphics Group of the Max-Planck-Institute for Computer Science in Saarbrücken, Germany, before joining UBC in 2000. His research interests include the intersection of imaging, optics, computer vision, computer graphics, and inverse problems. His more recent interest are computational imaging, focusing on hardware-software co-design of the next generation of imaging systems, with applications such as high-dynamic range imaging, compact computational cameras, hyperspectral cameras, to name just a few. His work on High Dynamic Range Displays served as the basis for the technology behind Brightside Technologies, which was acquired by Dolby, in 2007. He has served on numerous program committees for top-tier conferences such as Siggraph, Siggraph Asia, Eurographics, EGSR, and he has chaired the papers program for both Siggraph Asia and the International Conference of Computational Photography (ICCP) among others. He is the recipient of a 2014 Humboldt Research Award.



Judy Hoffman received the PhD degree from UC Berkeley, in 2016 advised by Trevor Darrell and Kate Saenko. She is currently an assistant professor with the School of Interactive Computing at Georgia Tech. Her research includes the intersection of computer vision and machine learning and focuses on tackling real-world variation, and scale while minimizing human supervision. She develops algorithms which facilitate transfer of information through unsupervised and semi-supervised model adaptation and generalization. Prior to joining Georgia Tech, she was a research scientist at Facebook AI Research, a postdoctoral scholar at UC Berkeley with Alyosha Efros, and Trevor Darrell and a postdoctoral scholar at Stanford working with Fei-Fei Li. She was supported by an NSF Graduate Research Fellowship. In 2015, she cofounded Women in Computer Vision, a service organization which provides mentoring and travel support to encourage students and early-career women to excel in the Computer Vision Community. She routinely serves on program committees for conferences in Computer Vision (CVPR, ICCV, ECCV) and machine learning (NeurIPS, ICML, and ICLR), has served as an area chair for CVPR 2019–2021, ICLR 2019/2020, ICML 2020, ICCV 2019, and is an associate editor for the *International Journal of Computer Vision*. She is the recipient of the NVIDIA Female Leader in Computer Vision Award (2020) and was listed as one of the top 100 most influential scholars in machine learning in the last 10 years (aminer 2020).



Stefanie Jegelka received the PhD degree from ETH Zurich and the Max Planck Institute for Intelligent Systems. She is an X-Window Consortium Career Development associate professor with the Department of EECS at MIT, and a member of the Computer Science and AI Lab (CSAIL). Before joining MIT, she was a postdoctoral researcher at UC Berkeley. She has received a Sloan Research Fellowship, an NSF CAREER Award, a DARPA Young Faculty Award, the German Pattern Recognition Award, and a Best Paper Award at ICML. She has served multiple times as an area chair for NeurIPS and ICML, as a NeurIPS 2020 workshop co-chair, ICML 2020 press co-chair, KDD 2016 proceedings co-chair, and on ICML awards committees, and she is an action editor for the *Journal of Machine Learning Research*. She has also given multiple tutorials and co-organized workshops on topics in discrete (submodular) and continuous optimization in machine learning, negative dependence, and graph neural networks. Her research interests span the theory and practice of algorithmic machine learning, including discrete and continuous optimization, discrete probability, and learning with structured data.



Shuiwang Ji (Senior Member, IEEE) received the PhD degree in computer science from Arizona State University, Tempe, Arizona, in 2010. Currently, he is an associate professor with the Department of Computer Science and Engineering, Texas A&M University, College Station, Texas. His research interests include machine learning, deep learning, graph neural networks, and computational biology. He received the National Science Foundation CAREER Award, in 2014. He is currently an associate editor for the *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *ACM Computing Surveys*, *ACM Transactions on Knowledge Discovery from Data*, and an action editor for *Data Mining and Knowledge Discovery*.



Evangelos Kalogerakis received the PhD degree from the University of Toronto, in 2010. He is currently an associate professor with the College of Information and Computer Sciences at the University of Massachusetts Amherst (UMass Amherst), where he leads a group of students working on graphics+vision. He joined UMass Amherst, in 2012. He was a postdoctoral researcher at Stanford University from 2010 to 2012. He has served as an area chair in CVPR, and on technical paper committees for ACM SIGGRAPH, ACM SIGGRAPH ASIA, Eurographics, and Symposium on Geometry Processing. He is currently an associate editor for the Editorial Boards of *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, *IEEE Transactions on Visualization & Computer Graphics (TVCG)*, and *Computers & Graphics* journal - Elsevier. He co-chaired the Shape Modeling International (SMI) Conference, in 2018. He was listed as one of the 100 most cited computer graphics scholars in the world between 2009 and 2019, and received the “Most Influential Scholar Award Honorable Mention for outstanding and vibrant contributions to the field Computer Graphics” awarded by the AMiner Academic Network. His research is supported by NSF awards and donations from Adobe. He received the NSERC Alexander Graham Fellowship and the International EPFL Dimitris N. Chorafas Award for his PhD work.



Adriana Kovashka is currently an assistant professor of computer science at the University of Pittsburgh. Her research interests include computer vision and machine learning. Specifically, her contributions lie in vision and language tasks, object recognition and detection, multimodal learning, image retrieval, modeling visual persuasion in the media, weakly supervised and transfer learning, and domain adaptation. She has co-authored 25 conference publications (CVPR, ICCV, ECCV, NeurIPS, AAAI, ACL, BMVC, ACCV, and WACV), three journal publications (*IEEE Transactions on Pattern Analysis and Machine Intelligence* and *International Journal of Computer Vision*), and ten workshop publications. She will serve as program co-chair at ICCV 2025. She has served as an area chair for CVPR 2018–2021, NeurIPS 2020, ICLR 2021, AAAI 2021 (Senior AC), and WACV 2016–2017. She has served as tutorials chair for CVPR 2020 and WACV 2018, and doctoral consortium chair / co-chair at CVPR 2015–2017. She was the guest editor for Distributed and Parallel Databases (DAPD) in 2019–2020, and Panelist for NSF (2015, 2017, 2019, and 2020). She has co-organized seven workshops at top-tier conferences. Her research is funded by the National Science Foundation, Google, Amazon, and Adobe.



Yong Jae Lee received the PhD degree from the University of Texas at Austin, in 2012, advised by Kristen Grauman. He is currently an associate professor with the Computer Science Department at the University of California, Davis, and will join the Computer Sciences Department at the University of Wisconsin-Madison, in Fall 2021. His research interests include computer vision, machine learning, and computer graphics, with a focus on creating robust visual recognition systems that can learn to understand the visual world with minimal human supervision. He was a post-doc at Carnegie Mellon University (2012–2013) and UC Berkeley (2013–2014) advised by Alyosha Efros. He is a recipient of several awards including an Army Research Office Young Investigator Program Award, NSF CAREER Award, UC Davis College of Engineering Outstanding Junior Faculty Award, Most Innovative Award at the COCO Object Detection Challenge 2019, and the Best Paper Award at BMVC 2020.



Bastian Leibe received the MS degree from the Georgia Institute of Technology, in 1999, the diploma degree from the University of Stuttgart, in 2001, and the PhD degree from ETH Zurich, in 2004, all in computer science. He is currently a full professor of computer science with RWTH Aachen University, Germany, where he leads the Computer Vision Group. His main research interests include computer vision and machine learning for dynamic visual scene understanding, encompassing object recognition, tracking, segmentation, and 3D reconstruction. He has published more than 120 articles in peer-reviewed journals and conferences, and his research work has accumulated more than 20000 citations (h-index 64, both according to Google Scholar). Over the years, he has received several awards for his research work, including the ETH Medal and the DAGM Main Prize, in 2004, CVPR Best Paper Award, in 2007, DAGM Olympus Prize, in 2008, IEEE ICRA Best Vision Paper Award, in 2009 and 2014, ISPRS Journal of Photogrammetry and Remote Sensing Best Paper of the Year Award, in 2010, and U.V. Helava Award for the Best Paper of the 4-Year Period 2008–2011 in the *ISPRS Journal of Photogrammetry and Remote Sensing*. In 2012, he was awarded a European Research Council (ERC) Starting Grant, and in 2017 an ERC Consolidator Grant. He has been program chair for ECCV 2016 and area chair and program committee member for all major computer vision conferences.



Victor Lempitsky received the PhD (“kandidat nauk”) degree from Moscow State University, 2007. He leads the Samsung AI Center in Moscow as well as the Vision, Learning, Telepresence (VIOLET) Lab at this center. He is also an associate professor with the Skolkovo Institute of Science and Technology (Skoltech). In the past, he was a researcher at Yandex, at the Visual Geometry Group (VGG) of Oxford University, and at the Computer Vision Group of Microsoft Research Cambridge. His research interests are in various aspects of computer vision and deep learning, in particular, generative deep learning, and telepresence applications. He has served as an area chair for top computer vision and machine learning conferences (CVPR, ICCV, ECCV, ICLR, NeurIPS) on multiple occasions. His recent work on neural head avatars was recognized as the most-discussed research publication of 2019 by Altmetric Top 100 rating.




Xuelong Li (Fellow, IEEE) received the BEng and PhD degrees from the University of Science and Technology of China (USTC). He is a professor at the Northwestern Polytechnical University, Xi’an, P. R. China, where he founded the Center for OPTical IMagery Analysis and Learning (OPTIMAL). He is a Clarivate analytics highly cited researcher in both Computer Science and Optical Engineering. He is also a member of the Academia Europaea.



Yan Liu received the PhD degree from Carnegie Mellon University. She is an associate professor with the Computer Science Department and the director of Machine Learning Center at the University of Southern California. She was a research staff member at IBM Research, in 2006–2010, and chief scientist in Didi Chuxing, in 2018. Her research interest includes machine learning and its applications to climate science, health care, and sustainability. She has received several awards, including NSF CAREER Award, Okawa Foundation Research Award, New Voices of Academies of Science, Engineering, and Medicine, Biocom Catalyst Award Winner, ACM Dissertation Award Honorable Mention, and the Best Paper Award in SIAM Data Mining Conference



Karen Livescu received the bachelor’s degree in physics from Princeton University, and the PhD degree and post-doc in electrical engineering and computer science at MIT. She is an associate professor at TTI-Chicago. Her main research interests include speech and language processing, and machine learning. She is an IEEE Signal Processing Society distinguished lecturer for 2021–2022. She is an associate editor for the *IEEE Open Journal of Signal Processing*, and has also served as an associate editor for the *IEEE Transactions on Audio, Speech, and Language Processing* (2014–2017) and a subject editor for *Speech Communication* (2012–2013). Other professional activities include serving as a member in the IEEE Spoken Language Technical Committee; as a technical co-chair of IEEE ASRU 2015, 2017, and 2019; and as a program co-chair of ICLR 2019.



Le Lu received the PhD degree in computer science from the Johns Hopkins University, in 2007. He is the executive director of Bethesda Research Lab of PAII Inc., in Bethesda, Maryland, USA. He founded the Deep Learning for Medical Imaging and Clinical Informatics Group at NVIDIA, in 2017, and was a senior research manager until June 2018. He was a staff scientist at the National Institutes of Health Clinical Center, Bethesda, Maryland, during 2013–2017. From 2006 to 2013, he worked at Siemens Corporate Research and Siemens Medical Solutions on various research roles. He has authored 170 highly selective journal and conference papers, 40 US patents granted or pending, and 32 first-tier clinical abstracts. Together with his trainees, has won two research trainee awards at RSNA 2016 and 2018, the Young Scientist Runner-Up Award at MICCAI 2017, and the “Test of Time” Publication Award at MICCAI 2018. He has made instrumental contributions on the public releases of several large-scale radiology imaging databases, including NIH-ChestXray14 and NIH-DeepLesion. He edited two books on the topics of deep learning for medical image computing by Springer, in 2017 and 2019. He won the NIH “Mentor of the Year” Award, in 2015 and NIH Clinical Center Director’s Award for “research excellence and significant patient care

impacts” in 2017. He serves as an area chair of major computer vision conferences for ten times; and won two outstanding reviewer awards at CVPR 2018 and BMVC 2017.



Julien Mairal received the PhD degree from Ecole Normale Supérieure, Cachan. He is a research scientist with Inria Grenoble, where he leads the Thoth research team. He joined Inria Grenoble in 2012, after a post-doc with the Statistics Department of UC Berkeley. His research interests include machine learning, computer vision, mathematical optimization, and statistical image and signal processing. In 2016, he received a Starting Grant from the European Research Council. He was awarded the Cor Baayen Prize, in 2013, the IEEE PAMI Young Researcher Award, in 2017, and the Test-of-Time Award at ICML 2019.



Philippos Mordohai received the PhD degree from the University of Southern California. He is a professor of computer science with the Stevens Institute of Technology, which he joined, in 2008. He held postdoctoral appointments at the University of North Carolina and the University of Pennsylvania. His research interests span 3D reconstruction from images and video, 3D segmentation and recognition, perception for robotics, and active vision. In the past few years, his group has focused on synthesizing geometric and machine learning techniques to address problems in geometric computer vision. He is an associate editor for the *Image and Vision Computing Journal* and *Computer Vision and Image Understanding*. He was selected for the 2020 Image and Vision Computing Associate Editor Award. He has served as area chair for CVPR, ICCV, and ECCV, and program co-chair of the International Conference on 3D Vision (3DV), 2019 and will be one of the program chairs of ICCV 2025.



Juan Carlos Niebles received an engineering degree in electronics from the Universidad del Norte, Colombia, in 2002, the MSc degree in electrical and computer engineering from the University of Illinois at Urbana-Champaign, in 2007, and the PhD degree in electrical engineering from Princeton University, in 2011. He is currently a co-director of the Stanford Vision and Learning Lab, associate director of research at the Stanford-Toyota Center for AI Research, and a senior research scientist at the Stanford AI Lab, since 2015. He was also an associate professor of electrical and electronic engineering at the Universidad del Norte (Colombia) between 2011 and 2019. His research interests include computer vision and machine learning, with a focus on visual recognition and understanding of human actions and activities, objects, scenes, and events. He has served as area chair for the top computer vision conferences CVPR, and ICCV. He is also a member of the AI Index Steering Committee and is the curriculum director for Stanford-AI4ALL. He is a recipient of a Google Faculty Research Award (2015), the Microsoft Research Faculty Fellowship (2012), a Google Research Award (2011), and a Fulbright Fellowship (2005).



Christopher Pal received the PhD degree from the University of Waterloo in 2005 and collaborated extensively with the University of Toronto's Machine Learning Group and Microsoft Research during his PhD. He is currently a full professor at Polytechnique Montréal and one of the original core faculty members at Mila, the Québec Artificial Intelligence (AI) Institute. He holds a Canada CIFAR AI Chair from the Canadian Institute for Advanced Research. He is also an adjunct professor in the Department of Computer Science and Operations Research at the University of Montréal and he is a Distinguished Scientist at Element AI (a ServiceNow company). He is one of the program chairs of the International Conference on Computer Vision (ICCV) 2021.



Sinno Jialin Pan received the PhD degree in computer science from the Hong Kong University of Science and Technology (HKUST), in 2011. He is a Provost's chair associate professor with the School of Computer Science and Engineering at Nanyang Technological University (NTU), Singapore. Prior to joining NTU, he was a scientist and lab head of text analytics with the Data Analytics Department, Institute for Infocomm Research, Singapore, from 2010 to 2014. He joined NTU as a Nanyang assistant professor (university named assistant professor), in November 2014. He was named to "AI 10 to Watch" by the *IEEE Intelligent Systems* magazine, in 2018. His research interests include transfer learning and its real-world applications.



Jonas Peters received the PhD degree jointly from MPI and ETH. He is currently a professor of statistics with the Department of Mathematical Sciences at the University of Copenhagen. Previously, he has been a group leader at the MPI for Intelligent Systems in Tuebingen and a Marie Curie fellow at the Seminar for Statistics, ETH Zurich. He studied mathematics at the University of Heidelberg and the University of Cambridge. He is interested in inferring causal relationships from different types of data and in building statistical methods that are robust with respect to distributional shifts. In his research, he seeks to combine theory, methodology, and applications. His work relates to areas such as computational statistics, causal inference, graphical models, independence testing, or high-dimensional statistics. Together with D. Janzing and B. Schölkopf he co-authored the book *Elements of Causal Inference: Foundations and Learning Algorithms* (MIT Press) that is available as an open access. Since 2016, he has been a member of the Junge Akademie. He won several awards, such as the Guy Medal in Bronze, awarded by the Royal Statistical Society (2019), the ASA Causality in Statistics Education Award (2018), Teacher of the Year at SCIENCE, University of Copenhagen (2018), and the ETH Medal for an Outstanding PhD Thesis.



Olga Russakovsky received the PhD degree from Stanford University, in 2015. She completed her postdoctoral fellowship at Carnegie Mellon University, in 2017. She is currently an assistant professor with the Computer Science Department, Princeton University. Her research interests include computer vision, closely integrated with the fields of machine learning, human-computer interaction and fairness, accountability, and transparency. She has been awarded the AnitaB.org's Emerging Leader Abie Award in honor of Denice Denton, in 2020, the CRA-WP Anita Borg Early Career Award, in 2020, the MIT Technology Review's 35-under-35 Innovator Award, in 2017, the PAMI Everingham Prize, in 2016 and Foreign Policy Magazine's 100 Leading Global Thinkers Award, in 2015. In addition to her research, she co-founded and continues to serve on the board of directors of the AI4ALL Foundation dedicated to increasing diversity and inclusion in Artificial Intelligence (AI).



Mathieu Salzmann received the PhD degree from EPFL, in January 2009, under the supervision of Prof. Pascal Fua. He is a senior researcher at EPFL-CVLab, and since May 2020, a part-time artificial intelligence engineer at ClearSpace. Previously, he was a senior researcher and research leader in NICTA's Computer Vision Research Group. Prior to this, from September 2010 to January 2012, he was a research assistant professor at TTI-Chicago, and, from February 2009 to August 2010, a postdoctoral fellow at ICSI and EECS at UC Berkeley under the supervision of Prof. Trevor Darrell. His research interests include domain adaptation and deep network compression/architecture search for visual recognition, as well as many areas related to image-based 3D reconstruction, such as 6D pose estimation, 3D human pose estimation, and surface shape reconstruction. He has recently become highly interested in the application of computer vision in space engineering and research.



Dimitris Samaras received the diploma degree in computer science and engineering from the University of Patras, in 1992, the MSc degree from Northeastern University, in 1994, and the PhD degree from the University of Pennsylvania, in 2001. He is currently a SUNY Empire Innovation Professor of computer science with Stony Brook University, where he directs the Computer Vision Lab. His research interests include human behavior analysis, generative image models, illumination modeling and estimation for recognition and graphics, and biomedical image analysis, especially digital histopathology. He is program chair of CVPR 2022.



Cees G.M. Snoek received the MSc degree in business information systems, in 2000, and the PhD degree in computer science, in 2005, both from the University of Amsterdam, The Netherlands. He is currently a full professor of computer science with the University of Amsterdam, where he heads the Video & Image Sense Lab. He is also a director of three public-private AI Research Labs: QUVA Lab with Qualcomm, Atlas Lab with TomTom, and AIM Lab with the Inception Institute of Artificial Intelligence. At University spin-off Kepler Vision Technologies, he acts as the chief scientific officer. He was previously a visiting scientist at Carnegie Mellon University and UC Berkeley, head of R&D at Euvision Technologies, and managing principal engineer at Qualcomm Research Europe. His research interests focus on making sense of videos and images by pattern analysis and machine intelligence. He is the lead researcher of the MediaMill Semantic Video Search Engine, which was the most consistent top performer in the yearly NIST TRECVID evaluations for over a decade. He was previously the general chair of ACM Multimedia 2016 in Amsterdam, and an associate editor for *IEEE MultiMedia*, *IEEE Transactions on Multimedia*, *ACM Transactions on Multimedia*, and currently for *Computer Vision and Image Understanding*. He is the recipient of NWO

Veni and Vidi career awards, a Fulbright Junior Scholarship, and the Netherlands Prize for ICT Research. Together with his PhD students he has won several awards, including the Best Paper of ACM Multimedia and the IEEE Transactions on Multimedia Prize.



Justin Solomon received the BS degree in computer science & mathematics, and the MS and PhD degrees in computer science, all from Stanford University. He is currently an associate professor of electrical engineering and computer science with the Massachusetts Institute of Technology (MIT). He leads the Geometric Data Processing Group in the MIT Computer Science, and Artificial Intelligence Laboratory (CSAIL), which studies problems at the intersection of geometry, optimization, and applications. He also completed postdoctoral research in the Princeton Program in Applied and Computational Mathematics.



Kalyan Sunkavalli received the MS degree in computer science from Columbia University, in 2006, and the PhD degree in computer science from Harvard University, in 2012. He is a senior research scientist at Adobe Research. His research interests lie at the intersection of computational photography, computer vision, graphics, and machine learning and specifically focus on reconstructing and editing visual appearance—shape, reflectance, illumination, and motion—in images and videos. He has published more than 60 papers on these topics in the leading computer vision and graphics conferences, and journals. His research has led to features that have shipped in commercial products. He was the program chair for IEEE ICCP 2020.



Ivor W. Tsang is currently a professor of artificial intelligence with the University of Technology Sydney. He is also the research director of the Australian Artificial Intelligence Institute (AAIL). His research focuses on transfer learning, learning from noisy supervision, deep generative models, and big data analytics. In 2013, he received his prestigious Australian Research Council Future Fellowship for his research regarding machine learning on Big Data. In 2019, his JMLR paper titled “Towards ultrahigh dimensional feature selection for big data” received the International Consortium of Chinese Mathematicians Best Paper Award. In 2020, he was recognized as the AI 2000 AAAI/IJCAI Most Influential Scholar in Australia for his outstanding contributions to the field of AAAI/IJCAI between 2009 and 2019. His research on transfer learning granted him the 2014 IEEE TMM Prize Paper Award and the Best Student Paper Award at CVPR 2010. In addition, he received the IEEE TNN Outstanding 2004 Paper Award, in 2007, and the 2012 ECCV Outstanding Reviewer Award. He serves as a senior area chair/area chair for NeurIPS, ICML, AISTATS, AAAI, and IJCAI. He also serves as on the editorial board for the *Journal of Machine Learning Research*, *Machine Learning Journal*, and *IEEE Transactions on Pattern Analysis and Machine Intelligence*.



Anton van den Hengel received the bachelor’s degree in mathematical science, in 1991, the bachelor of Laws degree in 1993, the master’s degree in computer science, in 1994, and the PhD degree in computer vision, in 2000, from The University of Adelaide. Currently, he is a professor in the School of Computer Science of the University of Adelaide, a director of applied science at Amazon, and the director of The Australian Institute for Machine Learning (AIML), one of Australia’s most successful Computer Vision and Machine Learning Research Groups. He has won a number of awards, including the CVPR Best Paper Prize, in 2010. His current research interests include vision-and-language problems, visual reasoning, and weakly supervised learning.



Martha White is currently an associate professor of computing science at the University of Alberta. Before joining the University of Alberta, in 2017, she was an assistant professor of computer science at Indiana University. She is a PI of AMII—the Alberta Machine Intelligence Institute—which is one of the top machine learning centres in the world, and a director of RLAI—the Reinforcement Learning and Artificial Intelligence Lab at the University of Alberta. She holds a Canada CIFAR AI chair and has authored more than 40 papers in top journals and conferences. She has served as an area chair or meta-reviewer for the top conferences in AI and ML, including ICML, NeurIPS, AAAI and IJCAI, as well as co-program chair for ICLR, and is an associate editor for the *IEEE Transactions on Pattern Analysis and Machine Intelligence*. Her research focus is on developing algorithms for agents continually learning on streams of data, with an emphasis on representation learning, and reinforcement learning.




Ole Winther received the PhD degree. He is professor of data science and complexity at DTU Compute, and professor MSO and group leader in genomic bioinformatics at the Department of Biology, KU/Rigshospitalet. He is currently a leading expert on statistical artificial intelligence with research focus on basic machine learning method development, and applications to the medical and natural language domains. He has been a part time group leader at the Department of Biology/BRIC, KU for some years and acquired expertise and built a close collaboration with experimental and clinical groups on biological high throughput data analysis. He works with text data in commercial setting within the medical diagnostic setting (findzebra.com) and the corporate (raffle.ai). He has been engaged in education (started DTU deep learning course with presently 300+ students), continued education, and industrial consultancies and public dissemination of AI.



John Wright received the PhD degree in electrical engineering from the University of Illinois at Urbana Champaign, in 2009. He is an associate professor with the Electrical Engineering Department at Columbia University. He is also affiliated with the Department of Applied Physics and Applied Mathematics and Columbia’s Data Science Institute. Before joining Columbia, he was with Microsoft Research Asia from 2009–2011. His research interests include sparse and low-dimensional models for high-dimensional data, optimization (convex and otherwise), and applications in imaging and vision. His work has received a number of awards and honors, including the 2012 COLT Best Paper Award and the 2015 PAMI TC Young Researcher Award.




Jianxin Wu received the BS and MS degrees in computer science from Nanjing University, and the PhD degree in computer science from the Georgia Institute of Technology. He is currently a professor with the Department of Computer Science and Technology at Nanjing University, China, and is associated with the National Key Laboratory for Novel Software Technology, China. He has served as an (senior) area chair for CVPR, ICCV, ECCV, AAAI, and IJCAI, and is an associate editor for the *Pattern Recognition* journal. His research interests include computer vision and machine learning, in particular, computer vision and deep learning tasks under the constraint that computing and data resources are limited or even highly insufficient.



Tao Xiang received the PhD degree from the National University of Singapore, in 2002. He is a professor of computer vision and machine learning and University Distinguished Chair at the University of Surrey. Before joining Samsung and the University of Surrey, he was a professor at the Queen Mary University of London. His research in computer vision has been focused on video surveillance, daily activity analysis, and sketch analysis. He also has interests in large-scale machine learning problems including zero/few-shot learning and domain adaptation. He has published more than 180 papers with 18K citations (h-index 71).



Sandra Zilles is currently a professor of computer science at the University of Regina, where she holds a Canada Research Chair in Computational Learning Theory. Her research on machine learning and artificial intelligence is funded by government agencies and industry partners and has led to more than 100 research publications. More than 45 highly qualified personnel have been trained under her (co-)supervision. Methods developed in her lab have found applications in research on autonomous vehicles, in research on genetics, and in cancer research. She has won several awards for her research and is also a member of the College of New Scholars, Artists, and Scientists of the Royal Society of Canada. Her expertise is often called upon on program committees and organizational committees of leading conferences in her field, and she is an associate editor for the reputable *Journal of Computer and System Sciences*. She also serves on the board of directors for Innovation Saskatchewan and on the Board of Directors for the Pacific Institute for the Mathematical Sciences (PIMS).



Wangmeng Zuo is currently a professor with the School of Computer Science and Technology, Harbin Institute of Technology. His research interests include low level vision, image/video generation, visual tracking, and image classification. He has published more than 100 papers in top-tier academic journals and conferences, including 10 IEEE *TPAMI* papers, and more than 50 papers on CVPR/ICCV/ECCV. His google scholar citations have reached 17000+. He also served as area chairs of ICCV 2019, CVPR 2020, a tutorial organizer in ECCV 2016, and a program co-chair of ECCV workshop on AIM 2020. He also serves as an associate editor of the *Visual Computer* journal and a senior associate editor of the *Journal of Electronic Imaging*. He was also a guest editor of *Neurocomputing*, *Pattern Recognition*, *IEEE Transactions on Circuits and Systems for Video Technology*, *IEEE Transactions on Neural Networks and Learning Systems*, and *Computer Vision and Image Understanding*.