FEATURE ARTICLE: ARTIFICIAL INTELLIGENCE; AI'S 10 TO WATCH AWARD

AI's 10 to Watch for 2024

Jurgen Dix, *Clausthal University of Technology, 38678, Clausthal, Germany*

Zhongfei (Mark) Zhang, *Binghamton University, State University of New York, Binghamton, NY, 13902, USA*

This article features the official announcement of the 2024 edition of the IEEE Computer Society biennial competition of AI's 10 to Watch focusing on identifying the rising stars in various areas of the broad AI communities. The 10 awardees are introduced in this article.

IEEE Intelligent Systems is promoting young and aspiring artificial intelligence (AI) scientists and recognizing the rising stars as "AI's 10 to Watch." For the biennial 2024 edition, we solicited submissions from individuals who had received their Ph.D. degree within the past 10 years. The selection committee selected 10 outstanding contributors from a pool of more than 30 highly competitive and strong nominations, which made the selection decisions rather difficult. After a careful and detailed selection process through many rounds of discussions via e-mails and live meetings, the committee voted unanimously on a short list of 10 top candidates who have all demonstrated outstanding achievements in different areas of AI. The selection was based solely on scientific quality, reputation, impact, and expert endorsements accumulated since receiving their Ph.D. It is our honor and privilege to announce the following 2024 class of "AI's 10 to Watch."

KAI-WEI CHANG, ASSOCIATE PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, UNIVERSITY OF CALIFORNIA, LOS ANGELES, AND AMAZON SCHOLAR, AMAZON ARTIFICIAL GENERAL INTELLIGENCE

Kai-Wei Chang[a] is an associate professor in UCLA Computer Science and an Amazon Scholar at Amazon Artificial General Intelligence (AGI). His research focuses on trustworthy natural language processing (NLP), with an emphasis on advancing reliable, responsible, and inclusive large language models and their multimodal variants. His research has been recognized by various paper awards, including the ACM Annual Conference on Knowledge Discovery and Data Mining (KDD) 2010 Best Paper Award, EMNLP 2017 Best Long Paper Award, ACL 2023 Outstanding Paper Award, and CVPR 2022 Best Paper Finalist. He is a Sloan fellow, an American Association for Artificial Intelligence (AAAI) senior member, and a recipient of the Google Research Scholar Award and the NSF CRII Award. His research is supported by the U.S. National Science Foundation (NSF), U.S. Department of Defense, National Institutes of Health, and industry partners like Google, Meta, Amazon, Taboola, OptumLabs, and Cisco.

Chang's pioneering work in trustworthy NLP addresses critical challenges such as fairness, robustness, explainability, inclusivity, and safety in language processing systems. He was among the first to quantify and mitigate gender stereotypes in word embeddings, establishing a foundation for broader studies on social biases in language understanding and generation systems. His current research focuses on adversarial robustness and AI model governance, ensuring that these models can be safely deployed in high-stakes applications, delivering both societal and practical impact. He has organized trustworthy NLP workshops since 2020 and delivered multiple tutorials on fairness, robustness, and multimodal alignment at major NLP and AI conferences. Through collective efforts from the research community and increasing recognition by industry, trustworthiness has become a mainstream research focus in NLP.

Chang's leadership extends beyond research. He was recently elected an officer of SIGDAT (ACL Special Interest Group) and is set to serve as its president in 2026. Over the previous eight years, he has mentored two postdoctoral students (postdocs), 12 Ph.D. students, and more than 70 master's students, many

[a]https://web.cs.ucla.edu/~kwchang/

Digital Object Identifier 10.1109/MIS.2024.3517893
Date of current version 7 February 2025.

of whom have gone on to become research scientists at companies such as Google, OpenAI, and Nvidia or faculty members at institutions like the University of Southern California, University of California, Merced, and Texas A&M. They continue to shape the future of trustworthy NLP and multimodal AI research. For more information, visit http://kwchang.net.

SIMON S. DU, ASSISTANT PROFESSOR, PAUL G. ALLEN SCHOOL OF COMPUTER SCIENCE AND ENGINEERING, UNIVERSITY OF WASHINGTON

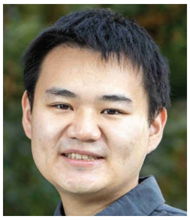

Simon S. Du[b] is an assistant professor at the Paul G. Allen School of Computer Science & Engineering at the University of Washington. Prior to starting as faculty, he was a postdoc at the Institute for Advanced Study of Princeton under the mentorship of Sanjeev Arora. He received his Ph.D. degree in machine learning (ML) from Carnegie Mellon University (CMU), where he was advised by Aarti Singh and Barnabás Póczos. Du's research has been recognized by a Sloan Research Fellowship, a Samsung AI Researcher of the Year Award, an Intel Rising Star Faculty Award, an NSF CAREER award, an Nvidia Pioneer Award, a Distinguished Dissertation Award honorable mention from CMU, among others. His notable contributions include proving the first global convergence result of gradient descent for optimizing deep neural networks, settling the sample complexity in reinforcement learning (RL), and establishing the necessary and sufficient conditions for RL in large state spaces. His current research focuses on multiagent RL and data selection algorithms for foundation models.

BO HAN, ASSISTANT PROFESSOR, HONG KONG BAPTIST UNIVERSITY

As trustworthiness is a key element to developing and deploying safe AI to improve our daily lives, Bo Han's[c] original and significant contributions focus on trustworthy ML (TML) under imperfect data in the following areas.

For the area of TML with noisy labels, his co-teaching algorithm has been elected the Most Influential Paper of Annual Conference on Neural Information Processing Systems, 2018 (NeurIPS'18) from Paper Digest, which well addresses the limitation of self-training MentorNet. The key idea of co-teaching is to train two deep neural networks simultaneously based on the memorization effects of deep neural networks and let them teach and filter each other every minibatch. The follow-up co-teaching+ algorithm bridges the "Update by Disagreement" strategy with the original co-teaching. More recently, to model real-world noise, Han's group investigated a harder noise named *instance-dependent* noise and proposed a fundamental solution called *confidence-scored instance-dependent* noise, where each instance-label pair is equipped with a confidence score.

For the area of TML against adversarial examples, the basic idea is to leverage the minimum-maximum optimization for obtaining adversarial robustness of the trained models proposed by Aleksander Madry. However, traditional adversarial training is conservative, sometimes hurting natural generalization. To address this issue well, the milestone works friendly adversarial training and geometry-aware instance-reweighted adversarial training have leveraged geometry information to balance adversarial robustness and natural generalization, which have been acknowledged and highlighted by parameter interpolation based adversarial training. This is an area with an extremely bright future because it relates to security and privacy of deploying deep neural networks.

For the area of TML under out-of-distribution (OOD) data, Han's OOD theory paper was awarded the Outstanding Paper Award of NeurIPS'22, which is the first study to provide the probably approximately correct theory for OOD detection. Furthermore, Han's group also contributed to a series of algorithms and applications in OOD detection and generalization, which themselves also set a novel frontier in the field. For the area of TML in federated environments, the group proposed trustworthy device–cloud deep learning, which has been deployed to Alibaba Taobao and Alipay, positively affecting hundreds of millions of customers.

Currently, Han's group is investigating two emerging areas of TML. The first is trustworthy foundation models because foundation models have shown impressive ability in real-life scenarios, including general problem solving and multimodal data processing. The other area is trustworthy causal learning because causality offers a rigorous way to investigate data-generating processes, and foundation models can effectively handle complex wild data. It is still an open problem how causality and foundation models can benefit mutually.

[b]https://www.cs.washington.edu/people/faculty/ssdu
[c]https://www.comp.hkbu.edu.hk/v1/?page=profile&id=bhanml

LINGJUAN LYU, HEAD OF PRIVACY-PRESERVING MACHINE LEARNING AND VISION FOUNDATION MODEL, SONY RESEARCH AND SONY AI

Lingjuan Lyu's[d] leadership in responsible AI and federated learning (FL) has profoundly influenced the AI community. She demonstrated exceptional foresight by proactively identifying a spectrum of privacy and security and intellectual property (IP) challenges and pioneering relevant countermeasures. Her series of works in using robust and invisible watermark techniques for protecting the IP of complex AI models and AI-generated content has gained widespread prominence. She was also able to resolve a series of practical and challenging problems in FL to propel it from theory to impactful practice, including collaborative fairness, privacy and security, domain and distribution shifts, and FL for financial technology and foundation models. Many of her proposed algorithms and frameworks in FL have been landed and commercialized in various industrial applications. She spearheads the development and open sourcing of one of the most impactful and vision-driven FL platforms, making a profound contribution to both the FL and computer vision communities by advancing innovation and practical applications.

SHIRUI PAN, PROFESSOR, GRIFFITH UNIVERSITY, AUSTRALIA

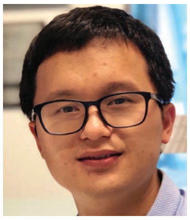

Dr. Shirui Pan[e] has gained international recognition as a rising star in AI and data science, particularly at the intersection of graph analytics and deep learning. Over the past few years, he has made significant contributions to advancing state-of-the-art techniques in graph neural networks (GNNs), knowledge graphs, and deep learning, with his work published in top-tier publications such as *NeurIPS*, *KDD*, *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *IEEE Transactions on Neural Networks and Learning Systems*, and *AAAI*. With more than 34,000 citations and an h-index of 71 achieved in just nine years since receiving his Ph.D. degree as of December 2024, his work is highly recognized by the research community. His notable accolades include the 2024 IEEE TNNLS Outstanding Paper Award and the ICDM-20 Best Student Paper Award, reflecting his exceptional contributions to the field.

Dr. Pan's pioneering work in GNNs has set benchmarks for both theoretical advancements and real-world applications. He introduced the first semisupervised deep learning method for attributed graphs [International Joint Conference on Artificial Intelligence (IJCAI), citations: more than 600] and brought adversarial training theory into network representation learning (IJCAI-2018, citations: more than 1000), redefining how complex networks are modeled. His recent publication in *Nature Machine Intelligence* (2024) presents an AI-driven tool for cost-effective drug discovery, receiving global attention for its ability to accelerate early-stage drug development. His comprehensive survey on GNNs (cited more than 10,000 times) offers a definitive taxonomy and road map, guiding future research and demonstrating his leadership in shaping the field.

In time-series analysis, Dr. Pan's influential works, such as "Connecting the Dots" (KDD 2020, citations: more than 1500) and "Graph WaveNet" (IJCAI 2019, citations: more than 2300), have transformed spatial–temporal modeling, with applications in traffic forecasting, energy systems, and business analytics. His research has been integrated into major frameworks like PyTorch Geometric and supported by funding from industry leaders such as Amazon and the Australian Research Council. Through his innovative research and leadership roles, including directing the TrustAGI Lab and organizing IEEE initiatives, Dr. Pan continues to advance AI and mentor the next generation of researchers.

XIAOQIAN QI, ASSISTANT PROFESSOR, THE UNIVERSITY OF HONG KONG

Humans navigate a dynamic 3-D world, constantly acquiring diverse skills and engaging in activities through perception, understanding, and interaction. Xiaoqian Li's[f] long-term research mission focuses on reconstructing and simulating this dynamic 3-D world, empowering AI systems to perceive, understand, and interact with it in a human-like way.

At the core of the work is the goal of rebuilding the dynamic 3-D world from real-world images and enabling its manipulation and recreation. This is challenging due to the ambiguities in extracting 3-D structures from 2-D images and the lack of effective 3-D representations. To overcome these obstacles, she developed data-driven methods for 3-D structure estimation

[d] https://ai.sony/people/Lingjuan-Lyu/
[e] https://experts.griffith.edu.au/37935-shirui-pan
[f] https://www.eee.hku.hk/people/xjqi/

and introduced new representations for world reconstruction. Notably, she pioneered a novel approach to single-view geometric estimation (depth and surface normal) using geometric constraints, yielding estimations consistent with real-world 3-D structures. Additionally, she advanced implicit and Gaussian representations for high-quality geometry reconstruction and 3-D/4-D image rendering. More recently, she designed hybrid neural architecture foundation models to generate high-resolution, consistent textures on 3-D objects.

In parallel with world reconstruction and creation, Li's research enhances AI's ability to comprehend and reason in 3-D. She pioneered the use of GNNs for analyzing unstructured 3-D spatial data and later developed adaptive convolutional operators and transformer architectures to improve processing speed and efficiency. More recently, she extended 3-D spatial intelligence to open-world environments, addressing challenges like diverse data distributions and large semantic vocabularies. This includes developing self-training methods for adapting 3-D models and repurposing image-based models as 3-D annotators, advancing 3-D spatial intelligence for real-world applications.

Beyond the realms of world reconstruction and comprehension, her work has had a broad interdisciplinary impact, driving AI advancements in fields such as medical diagnosis and scientific discovery. In collaborative projects, she has contributed to the development of networks for analyzing 3-D computerized tomography volumes to detect tumors and has helped to design AI-powered 3-D volume electron microscopy technologies that have the potential to overcome the physical limitations of biological imaging.

Looking ahead, Li envisions a future where AI systems not only interpret but also interact with the 3-D world through learning in realistic 3-D simulations, providing transformative solutions across industries such as health care, entertainment, architecture, and beyond.

SARATH SREEDHARAN, ASSISTANT PROFESSOR, DEPARTMENT OF COMPUTER SCIENCE, COLORADO STATE UNIVERSITY

Prof. Sarath Sreedharan[g] is an assistant professor at Colorado State University. His core research focuses on designing and developing human-aware decision-making systems, which, among other capabilities, can generate behaviors that align with human expectations. He received his Ph.D. degree from Arizona State University in 2022, where he focused on developing a conceptual and theoretical foundation for understanding and generating explanations in the context of sequential decision-making problems. His Ph.D. research demonstrated the utility of leveraging human mental models to generate effective explanations. His doctoral dissertation received one of the 2022 Dean's Dissertation Awards for Ira A. Fulton Schools of Engineering and was awarded an honorable mention for the ICAPS 2023 Outstanding Dissertation Award. Since then, he has focused on developing formal methods to support modeling and understanding more general human–AI interaction challenges. He has developed algorithms that allow AI systems to more effectively interpret user instructions, model human trust, and provide proactive support and has even developed methods to generate deceptive behaviors. The algorithms he has helped to develop have been applied in various application domains, including decision support, intelligent tutoring, and cybersecurity. He is the lead author of a Morgan Claypool monograph on explainable human–AI interactions. He was selected as a DARPA Riser Scholar for 2022 and featured as a Highlighted New Faculty at AAAI 2023. His research has won multiple awards, including the Best Systems Demo and Exhibit Award at the 2020 International Conference on Automated Planning and Scheduling and the Best Paper Award at the 2024 IEEE International Conference on Trust, Privacy and Security in Intelligent Systems, and Applications. He was also recognized as an Outstanding Program Committee member for AAAI 2020, a highlighted reviewer for the 2022 International Conference on Learning Representations, a Distinguished Program Committee member for IJCAI 2022 and 2023, and a top reviewer at NeurIPS 2022.

XINCHAO WANG, PRESIDENTIAL YOUNG PROFESSOR, NATIONAL UNIVERSITY OF SINGAPORE

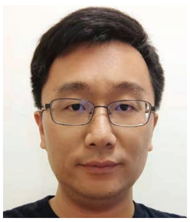

ML models, nowadays typically in the form of deep neural networks, have transformed many aspects of our daily lives. The success of state-of-the-art models, however, comes with significant costs: they require vast amounts of data as well as extensive training on thousands of GPUs for weeks or even months. Such a cumbersome process, apart from being extremely resource-intensive and environmentally unfriendly, significantly restricts the models' applicability on edge devices and precludes individual users or

[g]https://compsci.colostate.edu/person/?id=D0C6587C630B05A2E517779758A5B507&sq=t

small research labs with limited computational power from contributing to the community.

Wang's[h] research focuses on efficient ML, with the goal of developing faster, smaller, and, in many cases, more interpretable models, while reducing data dependency. Along this line, his work can be broadly categorized into three streams: efficient models, efficient strategies, and efficient data.

His work on efficient models explores lightweight architectures with fewer parameters and reduced computational requirements, enabling deployment on resource-limited platforms. Examples include MetaFormer, a framework that demonstrates that the computationally expensive token mixture, long considered a core component for the success of transformers, turns out to be effectively dispensable. As of now, MetaFormer has become one of the most widely adopted efficient transformer backbones across a broad range of domains. Wang's work on efficient strategies aims to develop novel schemes for building efficient models from existing ones. For instance, he developed the widely used structural pruning scheme DepGraph for constructing compact models from pretrained ones. DepGraph introduces, for the first time, a fully automatic method for tracing pruning dependencies across layers in a neural network, and it readily works for all major types of networks. Wang's work on efficient data focuses on condensing datasets into compact sets so that models trained on the smaller sets can approximate the performance of those trained on the original data. To this end, he introduced an innovative approach to factorizing datasets into bases and hallucinators, significantly enhancing the compression ratio while largely preserving the performance of the trained models.

JIAJUN WU, ASSISTANT PROFESSOR, STANFORD UNIVERSITY

AI researchers aim to develop machines with human-level scene understanding: from a single image, humans can explain what we see, reconstruct the scene in 3-D, predict what will happen, and plan our actions accordingly. Such "physical" understanding remains far from what current AI tools may achieve, despite impressive progress on large "foundation" models.

Toward this goal, Jiajun Wu's[i] core research contribution is a principled way to build efficient and versatile machines that learn to see, reason about, and interact with the physical world as humans do. The critical insight is to search for the "core knowledge" that learning systems need by exploiting the world's causal structure. To endow learning systems with causality and common sense, Wu and colleagues have developed AI methods that integrate top-down differentiable/neural simulation engines for computer graphics, physics, language, and human cognition, with bottom-up recognition models and perception and sensing systems. Bridging various disciplines of AI, Wu's research has led to significant progress on critical problems in perception, dynamics modeling, and cognitive reasoning, directly kindling a few now-popular subfields, such as 3-D generation, visual prediction, and multisensory learning.

Wu's research has been recognized through many prestigious awards, including eight best paper awards or candidates/nominations from the top computer vision (the Conference on Computer Vision and the Pattern Recognition and International Conference on Computer Vision), graphics (SIGGRAPH Asia), and robotics conferences (the International Conference on Robotics and Automation, Conference on Robot Learning, and IEEE/RSJ International Conference on Intelligent Robots and Systems); two Young Investigator Program Awards from the U.S. Air Force Office of Scientific Research and the U.S. Office of Naval Research; an NSF CAREER award; dissertation awards from the Association for Computing Machinery, AAAI, and the Massachusetts Institute of Technology; the 2020 Samsung AI Researcher of the Year; and faculty research awards from J.P. Morgan, Samsung, Amazon, and Meta.

JUNCHI YAN, PROFESSOR, SCHOOL OF AI, SHANGHAI JIAO TONG UNIVERSITY

Dr. Junchi Yan[j] has made contributions to AI4Science & Science4AI (including math, quantum, and so on). As one of early researchers in AI for optimization, he developed both AI-native and AI-aided approaches for solving complex problems, especially combinatorial optimization (CO) tasks including mixture integer programming, the satisfiability problem, and so on, where he further pushes the frontier

[h]https://cde.nus.edu.sg/ece/staff/wang-xinchao/

[i]https://engineering.stanford.edu/people/jiajun-wu

[j]https://thinklab.sjtu.edu.cn/

of AI4Crypto by developing neural networks to predict plaintext–ciphertext satisfiability. He further developed quantum-native and quantum-classic hybrid parameterized circuits (trainable models) for solving constrained tasks, ranging from CO (to replace quantum approximate optimization algorithm) to computational chemistry problems like molecular ground-state prediction. He also gives theoretical results on the trainability of some quantum circuits as well as specifically on Hamming weight preserving circuits in terms of the necessary and sufficient conditions for its subspace universality.

Dr. Yan's research also spans impact on applications in various areas, including electronic design automation (EDA) and autonomous driving (AD). In AI4EDA, he developed the early works and subsequent improvements for placement and routing, and also the RL pipeline for netlist reduction. In AD, his team developed the first-ever neural planner (using model-based RL: think2drive) that managed to solve all the cases in CARLA Leaderboard v2, which paves the way for end-to-end AD with model-based RL in a closed-loop manner. A benchmark for end-to-end AD methods was further developed, called *bench2drive*, by his team. He is the pioneering researcher in the area of prediction-and-optimization paradigm, with applications to both CO and AD.

He has also made diverse contributions to AI, such as ML for temporal point process, time-series transformer backbone (crossformer) and pretraining (up2me), rotated object detection, self-supervised and generative models (including quantum ones, e.g., a quantum variational autoencoder for 3-D molecule generation). His vision is to develop neural-symbolic systems for decision making and scientific computing.

ACKNOWLEDGMENTS

IEEE Intelligent Systems thanks selection committee members Chiranjib Bhattacharyya (Indian Institute of Science, Bangalore, India), Michael Fisher (University of Manchester, U.K.), Shi-Min Hu (Tsinghua University, China), Sarit Kraus (Bar-Ilan University Ramat Gan, Israel), and Lin Padgham (Royal Melbourne Institute of Technology, Melbourne, Australia), who devoted many hours to studying the nomination materials and discussing and selecting the best young members of our AI communities. We would like to warmly congratulate these awardees; they are truly outstanding as bright rising stars in their research areas and sincerely deserve this honor. We would also like to mention that the final selection result very much reflects the pool of submissions.

There was an overwhelming number of nominations from the ML, data mining, and computer vision areas, and geographically, only very few from Europe and Africa. Although the public typically and often thinks of AI as deep learning, we all know that this is not true. We would therefore like to encourage all young scientists working in all the areas of AI—in particular, those underrepresented, less-popular, classic areas of AI—in all parts of the world to consider applying for the next edition of this honor. Specifically, for those candidates who were nominated this year and were fewer than five years from receiving their Ph.D. degree, please do reapply. The selection was based on overall cumulative achievements, not on the ratio of "achievements per time." Jurgen Dix and Zhongfei (Mark) Zhang were Selection Committee cochairs and Editorial Board members.

JURGEN DIX is a professor of computer science at Clausthal University of Technology, 38678, Clausthal, Germany. Contact him at dix@tu-clausthal.de.

ZHONGFEI (MARK) ZHANG is a professor of computer science at Binghamton University, State University of New York, Binghamton, NY, 13902, USA. Contact him at zhongfei@cs.binghamton.edu.