Guest Editorial
Communication-Efficient Distributed Learning Over Networks

Xuanyu Cao, *Senior Member, IEEE*, Tamer Başar, *Life Fellow, IEEE*, Suhas Diggavi, *Fellow, IEEE*, Yonina C. Eldar, *Fellow, IEEE*, Khaled B. Letaief, *Fellow, IEEE*, H. Vincent Poor, *Life Fellow, IEEE*, and Junshan Zhang, *Fellow, IEEE*

Distributed machine learning is envisioned as the bedrock of future intelligent networks, where agents exchange information with each other to train models collaboratively without uploading data to a central processor. Despite its broad applicability, a downside of distributed learning is the need for iterative information exchange between agents, which may lead to high communication overhead unaffordable in many practical systems with limited communication resources. To resolve this communication bottleneck, we need to devise communication-efficient distributed learning algorithms and protocols that can reduce the communication cost and simultaneously achieve satisfactory learning/optimization performance. Accomplishing this goal necessitates synergistic techniques from a diverse set of fields, including optimization, machine learning, wireless communications, game theory, and network/graph theory. This Special Issue is dedicated to communication-efficient distributed learning from multiple perspectives, including fundamental theories, algorithm design and analysis, and practical considerations.

The Special Issue received 100 submissions, which demonstrated the timeliness and importance of research on communication-efficient distributed learning. After a rigorous and selective two-round peer-review procedure, 25 high-quality papers were accepted.

The Special Issue starts with an overview article [A1], in which Cao et al. present a comprehensive survey of prevailing methodologies for communication-efficient distributed learning, including reduction of the number of communication rounds, compression and quantization of exchanged information, radio resource management, and game-theoretic mechanisms. The tutorial paper is then followed by 25 technical papers.

Xuanyu Cao and Khaled B. Letaief are with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong (e-mail: eexcao@ust.hk; eekhaled@ust.hk).

Tamer Başar is with the Department of Electrical and Computer Engineering, University of Illinois Urbana–Champaign, Urbana, IL 61801 USA (e-mail: basar1@illinois.edu).

Suhas Diggavi is with the Department of Electrical and Computer Engineering, University of California at Los Angeles, Los Angeles, CA 90095 USA (e-mail: suhas@ee.ucla.edu).

Yonina C. Eldar is with the Department of Mathematics and Computer Science, Weizmann Institute of Science, 7610001 Rehovot, Israel (e-mail: yonina.eldar@weizmann.ac.il).

H. Vincent Poor is with the Department of Electrical and Computer Engineering, Princeton University, Princeton, NJ 08544 USA (e-mail: poor@princeton.edu).

Junshan Zhang is with the Department of Electrical and Computer Engineering, University of California at Davis, Davis, CA 95616 USA (e-mail: jazh@ucdavis.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/JSAC.2023.3241848.

Digital Object Identifier 10.1109/JSAC.2023.3241848

I. Reducing the Number of Communication Rounds of Distributed Learning

In [A2], Hu et al. study an asynchronous federated learning, with periodic aggregation to ameliorate the straggler effect due to asynchronous updates. They also study scheduling policies based on wireless link quality to improve the aggregated model updates, providing theoretical convergence analysis and numerics.

In [A3], Chiu et al. study communication efficient decentralized learning inspired by an existing method called MATCHA, which showed that not all links are equally important for convergence. Therefore, the authors design communication patterns during decentralized learning by activating different links with different frequencies using Laplacian matrix sampling, and design communication payload to obtain communication efficiency without hurting convergence. They provide theoretical analysis and numerics to support this.

In [A4], Chen et al. exploit the correlation between federated learning accuracy loss and the synchronization frequency, and seek to fine-tune the synchronization frequency at training runtime. They propose a metric called gradient consistency, which can effectively reflect the training status despite the instability of realistic FL scenarios. A feedback-driven algorithm called Gradient-Instructed Frequency Tuning (GIFT), which adaptively increases or decreases the synchronization frequency based on gradient consistency, is then developed which improves FL accuracy by up to 10.7% with a time reduction of 58.1%.

In [A5], Yang et al. propose two client selection schemes, Distributed Stein Variational Gradient Descent (DSVGD) based on the Kernelized Stein Discrepancy (KSD) and the Hilbert Inner Product (HIP), for federated Bayesian learning. Comprehensive experiments using real datasets confirmed the performance gain of the proposed schemes. In particular, it is shown that the KSD-based scheme has smaller overhead but the HIP-based scheme has faster convergence.

In [A6], Ayache et al. consider the problem of parameter server that wishes to learn a model that fits data distributed on the nodes of a graph, with federated learning as a canonical application. They consider a random walk learning algorithm that addresses the issues of communication limitations between the nodes and the parameter server and of data heterogeneity across nodes. Casting this problem as a sleeping multi-armed bandit problem, the authors design a near-optimal sampling strategy with guaranteed convergence and that outperforms existing algorithms.

II. Compressing the Communications of Distributed Learning

In [A7], Ge et al. develop a sparsity-preserving sketch-based collective communication method, which preserves gradient sparsity and reduces communication costs via a bitmap-informed count sketch structure. It is shown that the method has the same convergence rate as vanilla data-parallel training and a much smaller communication overhead than those of state-of-the-art methods.

In [A8], Park and Choi consider communication-efficient federated learning by exploiting the correlations between consecutive local model update, and proposed a projection-based compression mechanism. They analyze the convergence behavior of such methods as well as provide evidence through numerics of its advantages over some benchmark methods.

In [A9], Jahani-Nezhad et al. study the communication overhead of secure aggregation for federated learning by using a group-based decentralized communication topology and secure multi-party computing (MPC) techniques, which they name SwiftAgg+ (building on existing SwiftAgg). This is mainly a theoretical paper that analyzes the privacy and communication load for the proposed schemes, and compares it to information-theoretic lower bounds, establishing it to be close to optimal for some cases.

In [A10], Chen et al. propose a service delay efficient Federated learning (SDEFL) scheme over mobile devices via the joint design of local computing control, weight quantization, and gradient quantization. In particular, it investigates the impacts of the weight quantization, gradient quantization, and local computing control strategies on the service delay, and provides the convergence rate of federated learning with compression from a theoretical perspective. Furthermore, it develops an optimization scheme to minimize the service delay of federated learning over heterogeneous devices. Extensive experimental results are presented to demonstrate the effectiveness and efficiency of the proposed SDEFL scheme, and its advantages over peer designs in various learning and wireless transmission settings.

III. Radio Resource Management for Distributed Learning

In [A11], Yang et al. propose to deploy split machine learning in a wireless MIMO communication network by utilizing the intricate interplay between MIMO-based over-the-air computation and neural networks. This can alleviate the communication load substantially.

In [A12], Guo et al. address the challenges presented by radio resource management (RRM) schemes in the context of spectrum sharing in cellular vehicle-to-everything (C-V2X), and propose a hybrid centralized-distributed RRM scheme and a distributed RRM scheme, designed toward maximizing the system capacity while guaranteeing the reliability of V2V links. Toward that end, they introduce a decoupling method that provides a theoretical lower bound so that channel allocation and power control can be optimized independently, which they proceed within the article. Specifically, the hybrid centralized-distributed RRM scheme they employ is based on graph matching and reinforcement learning, and the distributed RRM scheme employed requires only local channel state information with the hybrid-framework reinforcement learning. The article also includes numerical demonstrations of both schemes, out-performing competing ones.

In [A13], Du et al. propose a dynamic device scheduling mechanism for federated learning via AirComp, which can select qualified edge devices to transmit their local models with a proper power control policy. In this mechanism, the data importance measured by the gradient of local model parameters, channel condition, and energy consumption of devices are jointly considered. The simulation results validate that the proposed scheduling mechanism can achieve higher training accuracy and faster convergence rate, and is robust against different channel conditions.

In [A14], Wu et al. consider a training latency reduction problem in Split learning (SL) over wireless networks. In contrast to existing approaches, the paper introduces a low-latency SL scheme, named Cluster-based Parallel SL (CPSL), which parallelizes the device-side model training in order to reduce training latency. In addition, it proposes a two-timescale resource management algorithm for the CPSL to minimize the training latency in wireless networks by taking network dynamics and device heterogeneity into account. Extensive simulation results are presented to validate the effectiveness of the proposed solutions compared with the existing solutions.

In [A15], Ben-Hur et al. consider binary classification via an ensemble of functions communicating real-value confidence levels. They propose a solution by optimizing the transmission gain and aggregation coefficients, and develop a post-training optimization algorithm to minimize the error probability.

In two-companion papers [A16], [A17], Liu et al. consider communication-efficient distributed learning of time-varying states over networks. They first propose encoding strategies for generating messages and present a sufficient condition for bounding the total inference errors of all agents over time. Then, they derive a necessary condition in terms of the sensing and communication capabilities of the network for bounding the learning error over time. By comparing the necessary condition with the sufficient condition, they show that the gap between the two conditions is small in some cases.

IV. Applications of Communication-Efficient Distributed Learning

In [A18], Schynol and Pesavento address coordinated weighted sum-rate maximization in multicell MIMO networks with intra- and inter-cell interference and local channel state at the base stations, and propose an unrolled architecture (which they name GCN-WMMSE) based on the classical weighted minimum mean squared error (WMMSE) algorithm. GCN-WMMSE uses ideas from graph signal processing and is agnostic to different wireless network topologies while exhibiting a low number of trainable parameters and high efficiency with respect to training data. The authors also provide favorable comparison of GCN-WMMSE to earlier unrolled WMMSE architectures despite its generality.

In [A19], Zhou et al. address the dynamic pricing problem for distributed machine learning jobs, while jointly taking the placement into consideration. The algorithm they propose estimates the unknown input by leveraging a multi-armed bandit online learning framework, and calculates rewards based on feedback of job runtime. Their algorithm consists of two subroutines: a dynamic pricing mechanism that determines the provider-profit maximizing price upon the arrival of each job, and a placement strategy that minimizes the runtime of accepted jobs. The authors also show that their algorithm achieves sub-linear regret with both the timespan and the total job number.

In [A20], Wang et. al consider the communication efficiency, low latency, and accuracy for distributed relative state estimation involving vision data in UAV networks. In particular, they investigate the node-based local graph model (NBM) and patch-based local graph model (PBM), and propose a communication-efficient initialization scheme together with a distributed state iteration method. The communication costs for the initialization and online iteration processes are then analyzed theoretically.

In [A21], Huang et. al design incentive mechanisms for the problem of information elicitation without verification (IEWV). They study a scenario where the data requester cannot access workers' heterogeneous information quality and costs ex-ante. They propose a continuum-armed bandit-based incentive mechanism that dynamically learns the optimal reward level from workers' reported information. The resulting inference problem is non-convex but is reformulated as a bi-convex problem. They then show that the proposed mechanism achieves a sub-linear regret and outperforms several celebrated benchmarks.

In [A22], Liu et al. study optimal resource allocation with batching and early exiting, which is an NP-complete integer program. A set of efficient algorithms are designed under the criterion of maximum throughput. The experimental results demonstrate that the proposed resource allocation algorithms can leverage integrated batching and early exiting to achieve 200% throughput gain over conventional schemes.

In [A23], Rodrigues and Kato propose a hybrid centralized and distributed learning solution to train a deep Q network (DQN) model in Multi-Access Edge Computing (MEC) equipped satellite networks. They then analyze the learning costs behind centralized and distributed learning and propose a hybrid solution that adaptively uses the advantages of both in cloud server-equipped satellite networks. The extensive results demonstrate that the proposal is not only efficient in solving machine learning tasks but also dynamic to react to different configurations while maintaining excellent performance.

In [A24], Zhang et al. consider the sensitivity to network conditions of real-time collaborative inference based on large volumes of image data. In particular, the present RTCoInfer, a framework for providing accurate and timely stream analytics in practical fluctuating networks. Extensive experiments show that this method achieves better efficiency than the prior state-of-the-art.

In [A25], Xie and Song also address the issues of communication bottleneck and data heterogeneity in federated learning, but in the context of reinforcement learning. In particular, they consider the imposition of a Kullback–Leibler-based penalty on the divergence between local and global policies, which achieves a favorable tradeoff between training speed (i.e., step-size) and convergence. The experiment results on two popular RL experiment platforms demonstrate the advantage of the proposed algorithm over existing methods in accelerating and stabilizing the training process with heterogeneous data.

In [A26], Xu et al. propose a multi-cell cluster-free NOMA framework that is designed to simultaneously mitigate both intra-cell and inter-cell interference. They treat the joint design problem as a sum rate maximization problem and develop a distributed auto-learning graph neural network (AutoGNN) architecture to achieve communication-efficient distributed scheduling.

ACKNOWLEDGMENT

The guest editorial team would like to thank all authors for their high-quality submissions and all reviewers for their time invested in improving the manuscripts. The Guest Editors are also grateful to Prof. Patrick Thiran for handling many aspects related to the Special Issue in his capacity as a Senior Editor.

APPENDIX: RELATED ARTICLES

[A1] X. Cao et al., "Communication-efficient distributed learning: An overview," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 851–873, Apr. 2023.

[A2] C.-H. Hu, Z. Chen, and E. G. Larsson, "Scheduling and aggregation design for asynchronous federated learning over wireless networks," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 874–886, Apr. 2023.

[A3] C.-C. Chiu, X. Zhang, T. He, S. Wang, and A. Swami, "Laplacian matrix sampling for communication-efficient decentralized learning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 887–901, Apr. 2023.

[A4] C. Chen et al., "GIFT: Towards accurate and efficient federated learning with gradient-instructed frequency tuning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 902–914, Apr. 2023.

[A5] J. Yang, Y. Liu, and R. Kassab, "Client selection for federated Bayesian learning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 915–928, Apr. 2023.

[A6] G. Ayache, V. Dasari, and S. E. Rouayheb, "Walk for learning: A random walk approach for federated learning from heterogeneous data," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 929–940, Apr. 2023.

[A7] K. Ge, K. Lu, Y. Fu, X. Deng, Z. Lai, and D. Li, "Compressed collective sparse-sketch for distributed data-parallel training of deep learning models," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 941–963, Apr. 2023.

[A8] S. Park and W. Choi, "Regulated subspace projection based local model update compression for communication-efficient federated learning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 964–976, Apr. 2023.

[A9] T. Jahani-Nezhad, M. Maddah-Ali, S. Li, and G. Caire, "SwiftAgg+: Achieving asymptotically optimal communication loads in secure aggregation for federated learning," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 977–989, 2023.

[A10] R. Chen, D. Shi, X. Qin, D. Liu, M. Pan, and S. Cui, "Service delay minimization for federated learning over mobile devices," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 990–1006, Apr. 2023.

[A11] Y. Yang et al., "Over-the-air split machine learning in wireless MIMO networks," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1007–1022, Apr. 2023.
[A12] C. Guo, C. Wang, L. Cui, Q. Zhou, and J. Li, "Radio resource management for C-V2X: From a hybrid centralized-distributed scheme to a distributed scheme," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1023–1034, Apr. 2023.
[A13] J. Du, B. Jiang, C. Jiang, Y. Shi, and Z. Han, "Gradient and channel aware dynamic scheduling for over-the-air computation in federated edge learning systems," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1035–1050, Apr. 2023.
[A14] W. Wu et al., "Split learning over wireless networks: Parallel design and resource management," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1051–1066, Apr. 2023.
[A15] Y. Ben-Hur, A. Goren, D.-E. Klang, Y. Kim, and Y. Cassuto, "Ensemble classification with noisy real-valued base-functions," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1067–1080, Apr. 2023.
[A16] Z. Liu, A. Conti, S. Mitter, and M. Z. Win, "Communication-efficient distributed learning over networks: Sufficient conditions for accuracy," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1081–1101, Apr. 2023.
[A17] Z. Liu, A. Conti, S. Mitter, and M. Z. Win, "Communication-efficient distributed learning over networks: Necessary conditions for accuracy," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1102–1119, Apr. 2023.
[A18] L. Schynol and M. Pesavento, "Coordinated sum-rate maximization in multicell MU-MIMO with deep unrolling," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1120–1134, Apr. 2023.
[A19] R. Zhou, X. Zhang, J. C. S. Lui, and Z. Li, "Dynamic pricing and placing for distributed machine learning jobs: An online learning approach," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1135–1150, Apr. 2023.
[A20] S. Wang, Y. Wang, X. Bai, and D. Li, "Communication efficient, distributed relative state estimation in UAV networks," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1151–1166, Apr. 2023.
[A21] C. Huang, H. Yu, J. Huang, and R. A. Berry, "An online inference-aided incentive framework for information elicitation without verification," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1167–1185, Apr. 2023.
[A22] Z. Liu, Q. Lan, and K. Huang, "Resource allocation for multiuser edge inference with batching and early exiting," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1186–1200, Apr. 2023.
[A23] T. Rodrigues and N. Kato, "Hybrid centralized and distributed learning for MEC-equipped satellite 6G networks," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1201–1211, Apr. 2023.
[A24] Z. Zhang et al., "RTCoInfer: Real-time collaborative CNN inference for stream analytics on ubiquitous images," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1212–1226, Apr. 2023.
[A25] Z. Xie and S. Song, "FedKL: Tackling data heterogeneity in federated reinforcement learning by penalizing KL divergence," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1227–1242, Apr. 2023.
[A26] X. Xu, Y. Liu, Q. Chen, X. Mu, and Z. Ding, "Distributed auto-learning GNN for multi-cell cluster-free NOMA communications," *IEEE J. Sel. Areas Commun.*, vol. 41, no. 4, pp. 1243–1258, Apr. 2023.

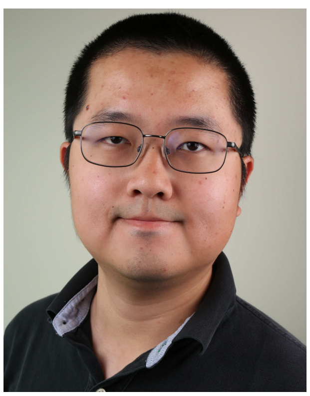

Xuanyu Cao (Senior Member, IEEE) received the B.S. degree in electrical engineering from Shanghai Jiao Tong University, China, in 2013, and the M.S. and Ph.D. degrees in electrical engineering from the University of Maryland, College Park, MD, USA, in 2016 and 2017, respectively. From August 2017 to October 2021, he was successively a Post-Doctoral Research Associate with the Department of Electrical Engineering, Princeton University, and the Coordinated Science Laboratory, University of Illinois at Urbana–Champaign. Since October 2021, he has been an Assistant Professor with the Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology, Hong Kong. His research interests include distributed/online optimization, communication-efficient distributed learning, federated learning over wireless networks, and game theory and network economics. He is an Editor of IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and IEEE TRANSACTIONS ON VEHICULAR TECHNOLOGY, a TPC Member of ACM MobiHoc 2022 and 2023, and the Lead Guest Editor of the Special Issue on Communication-Efficient Distributed Learning Over Networks in IEEE JOURNAL ON SELECTED AREAS IN COMMUNICATIONS.

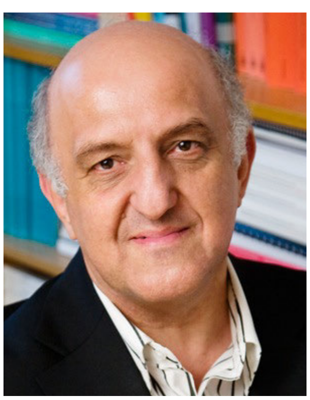

Tamer Başar (Life Fellow, IEEE) received the B.S. degree in electrical engineering from Robert College, Istanbul, Turkey, and the M.S., M.Phil., and Ph.D. degrees from Yale University. Since 1981, he has been with the University of Illinois at Urbana–Champaign (UIUC), where he is currently the Swanlund Endowed Chair Emeritus and the Center for Advanced Study (CAS) Professor Emeritus of electrical and computer engineering. He is also with the Coordinated Science Laboratory, Information Trust Institute, and the Department of Mechanical Science and Engineering, UIUC. At UIUC, he has also served as the Director of CAS from 2014 to 2020, the Interim Dean of Engineering in 2018, and the Interim Director of the Beckman Institute from 2008 to 2010. He has more than 1000 publications in systems, control, communications, optimization, networks, and dynamic games, including books on dynamic non-cooperative game theory, robust control, network security, wireless communication and networks, and stochastic networked control systems. His current research interests include stochastic teams, games, and networks; risk-sensitive estimation and control; mean-field game theory; multi-agent systems and learning; data-driven distributed optimization; epidemics modeling and control over networks; strategic information transmission, spread of disinformation, and deception; security and trust; energy systems; and cyber-physical systems. He is a member of the U.S. National Academy of Engineering and a fellow of IFAC and SIAM. He has served as the President of the IEEE Control Systems Society (CSS), the International Society of Dynamic Games (ISDG), and the American Automatic Control Council (AACC). He received the Wilbur Cross Medal from Yale University in 2021. He has also received several awards and recognitions over the years, including the highest awards of IEEE CSS, IFAC, AACC, and ISDG, the IEEE Control Systems Award, and a number of international honorary doctorates and professorships. He was the Editor-in-Chief of *Automatica* from 2004 to 2014. He is currently an editor of several book series.

Suhas Diggavi (Fellow, IEEE) received the bachelor's degree from the Indian Institute of Technology (IIT) Delhi and the Ph.D. degree from Stanford University.

He worked as a Principal Member Research Staff at AT&T Shannon Laboratories and directed the Laboratory for Information and Communication Systems (LICOS) at EPFL. He is currently a Professor of electrical and computer engineering at the University of California at Los Angeles (UCLA), Los Angeles, CA, USA, where he also directs the Information Theory and Systems Laboratory. He has eight issued patents. His research interests include information theory and its applications to several areas, including machine learning, security and privacy, wireless networks, data compression, cyber-physical systems, bio-informatics, and neuroscience.

Dr. Diggavi has received several recognitions for his research from IEEE and ACM, including the 2013 IEEE Information Theory Society and Communications Society Joint Paper Award, the 2021 ACM Conference on Computer and Communications Security (CCS) Best Paper Award, the 2013 ACM International Symposium on Mobile Ad Hoc Networking and Computing (MobiHoc) Best Paper Award, and the 2006 IEEE Donald Fink Prize Paper Award. He was selected as a Guggenheim Fellow in 2021. He also received the 2019 Google Faculty Research Award, the 2020 Amazon Faculty Research Award, and the 2021 Facebook/Meta Faculty Research Award. He served as an IEEE Distinguished Lecturer and also served on the Board of Governors for the IEEE Information Theory Society from 2016 to 2021. He has been an Associate Editor of IEEE TRANSACTIONS ON INFORMATION THEORY, IEEE/ACM TRANSACTIONS ON NETWORKING, and other journals and special issues, and in the program committees of several conferences. He has also helped organize IEEE and ACM conferences, including serving as the Technical Program Co-Chair for the 2012 IEEE Information Theory Workshop (ITW), the Technical Program Co-Chair for the 2015 IEEE International Symposium on Information Theory (ISIT), and the General Co-Chair for ACM Mobihoc 2018. More information can be found at http://licos.ee.ucla.edu.

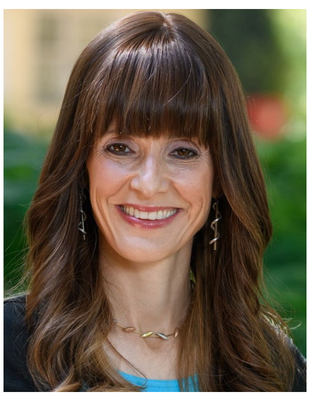

Yonina C. Eldar (Fellow, IEEE) received the first B.Sc. degree in physics and the second B.Sc. degree in electrical engineering from Tel-Aviv University (TAU), Tel-Aviv, Israel, in 1995 and 1996, respectively, and the Ph.D. degree in electrical engineering and computer science from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2002. She was a Visiting Professor at Stanford University. She is currently a Professor with the Department of Mathematics and Computer Science, Weizmann Institute of Science, Rehovot, Israel, where she is the Head of the Center for Biomedical Engineering. She is also a Visiting Professor at MIT, a Visiting Scientist at the Broad Institute, and an Adjunct Professor at Duke University. She is a member of the Israel Academy of Sciences and Humanities and a EURASIP Fellow. She was a Horev Fellow of the Leaders in Science and Technology Program at the Technion—Israel Institute of Technology and selected as one of the 50 most influential women in Israel and Asia. She was also a member of the Young Israel Academy of Science and Humanities and the Israel Committee for Higher Education. She has received many awards for excellence in research and teaching, including the IEEE Signal Processing Society Technical Achievement Award in 2013, the IEEE/AESS Fred Nathanson Memorial Radar Award in 2014, and the IEEE Kiyo Tomiyasu Award in 2016. She is the Editor-in-Chief of *Foundations and Trends in Signal Processing* and serves on many IEEE committees.

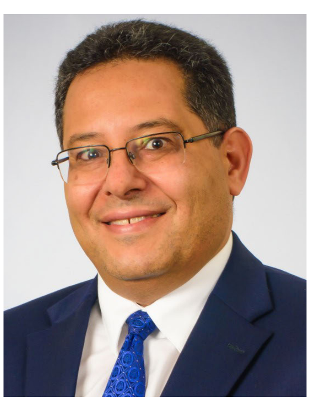

Khaled B. Letaief (Fellow, IEEE) received the B.S. (Hons.), M.S., and Ph.D. degrees in electrical engineering from Purdue University, West Lafayette, IN, USA, in 1984, 1986, and 1990, respectively, and the Ph.D. (Honoris Causa) degree from the University of Johannesburg, South Africa, in 2022.

From 1990 to 1993, he was a Faculty Member at The University of Melbourne, Australia. Since 1993, he has been with The Hong Kong University of Science and Technology (HKUST), where he has held many administrative positions, including Acting Provost, Dean of Engineering, Head of the Department of Electronic and Computer Engineering, Director of the Wireless IC Design Center, Founding Director of the Huawei Innovation Laboratory, and Director of the Hong Kong Telecom Institute of Information Technology. From September 2015 to March 2018, he joined Hamad Bin Khalifa University (HBKU), Qatar, as a Provost to help establish a research-intensive university in Qatar in partnership with strategic partners that include Northwestern University, Carnegie Mellon University, Cornell University, and Texas A&M University.

Dr. Letaief is currently serving as a member of the IEEE Board of Directors, the U.S. National Academy of Engineering, The National Academy of Sciences, India, and the Hong Kong Academy of Engineering Sciences, and a fellow of The Hong Kong Institution of Engineers. He was a recipient of many distinguished awards and honors, including the 2007 IEEE Communications Society Joseph LoCicero Publications Exemplary Award, the 2010 Purdue University Outstanding Electrical and Computer Engineer Award, the 2011 IEEE Communications Society Harold Sobol Award, the 2016 IEEE Marconi Prize Paper Award in Wireless Communications, the 2019 IEEE Communications Society and Information Theory Society Joint Paper Award, the 2021 IEEE Communications Society Best Survey Paper Award, the 2022 IEEE Communications Society Edwin Howard Armstrong Achievement Award, and more than 19 IEEE Best Paper Awards. He is also recognized by Thomson Reuters as an ISI Highly Cited Researcher. He was listed among the 2020 top 30 of AI 2000 Internet of Things Most Influential Scholars. He is well recognized for his dedicated service to professional societies and IEEE, where he has served in many leadership positions. These include the Founding Editor-in-Chief of the prestigious IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS. He has also served as the President of the IEEE Communications Society from 2018 to 2019, the world's leading organization for communications professionals with headquarter in New York City and members in 162 countries. He is an Internationally Recognized Leader of wireless communications and networks.

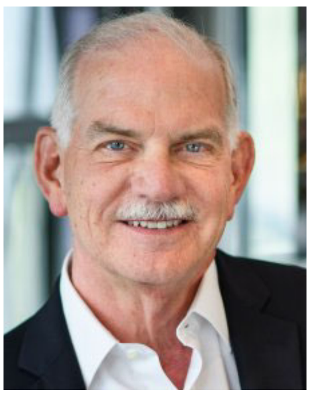

H. Vincent Poor (Life Fellow, IEEE) received the Ph.D. degree in electrical engineering and computer science (EECS) from Princeton University in 1977. From 1977 to 1990, he was on the faculty of the University of Illinois at Urbana–Champaign. Since 1990, he has been on the faculty at Princeton, where he is currently the Michael Henry Strater University Professor. From 2006 to 2016, he served as the Dean of Princeton's School of Engineering and Applied Science. He has also held visiting appointments at several other universities, including most recently at Berkeley and Cambridge. His research interests include information theory, machine learning, and network science, and their applications in wireless networks, energy systems, and related fields. Among his publications in these areas is the recent book *Machine Learning and Wireless Communications* (Cambridge University Press, 2022). He is a member of the National Academy of Engineering and the National Academy of Sciences and a foreign member of the Chinese Academy of Sciences, the Royal Society, and other national and international academies. He received the IEEE Alexander Graham Bell Medal in 2017.

Junshan Zhang (Fellow, IEEE) received the Ph.D. degree from the School of Electrical and Computer Engineering (ECE), Purdue University, West Lafayette, IN, USA, in 2000.

He was a Faculty Member with the School of Electrical, Computer and Energy Engineering (ECEE), Arizona State University, from 2000 to 2021. He is currently a Professor with the Department of Electrical and Computer Engineering, University of California at Davis, Davis, CA, USA. Building on his research findings, he co-founded Smartiply Inc., a fog computing startup company, delivering boosted network connectivity and embedded artificial intelligence. His research interests include information networks and data science, including edge intelligence, reinforcement learning, continual learning, network optimization and control, game theory, with applications in connected and automated vehicles, 5G and beyond, wireless networks, IoT data privacy/security, and smart grids.

Prof. Zhang was a recipient of the NSF CAREER Award in 2003, the ONR Young Investigator Award in 2005, and the IEEE WIRELESS COMMUNICATION Technical Committee Recognition Award in 2016. His papers have won a few awards, including the Best Paper Awards at IEEE ICC 2008 and ICC 2017, the Best Paper Runner-Up Awards of IEEE INFOCOM 2009 and IEEE INFOCOM 2014, the Kenneth C. Sevcik Outstanding Student Paper Award of ACM SIGMETRICS/IFIP Performance 2016, and the Best Student Paper Award from WiOPT 2018. He was the General Chair of ACM/IEEE SEC 2017 and WiOPT 2016. He was a Distinguished Lecturer of the IEEE Communications Society. He is currently serving as the Editor-in-Chief for the IEEE TRANSACTIONS ON WIRELESS COMMUNICATIONS and a Senior Editor for IEEE/ACM TRANSACTIONS ON NETWORKING. He was the TPC Co-Chair for a number of major conferences in communication networks, including IEEE INFOCOM 2012 and ACM MOBIHOC 2015.