

Francis Y. Yan

Redmond, WA
Updated: Nov. 2024

fyy@illinois.edu
<https://francisyyan.org>

EMPLOYMENT

University of Illinois Urbana-Champaign	Urbana, IL
Assistant Professor of Computer Science	Jan. 2025 (<i>incoming</i>)
Affiliate Professor of Electrical & Computer Engineering	
Adjunct Professor of Computer Science	Aug. 2024 – Jan. 2025
Microsoft Research	Redmond, WA
Senior Researcher	Aug. 2020 – Present
<ul style="list-style-type: none">• Intelligent Networked Systems Group (led by Victor Bahl)	

EDUCATION

Stanford University	Stanford, CA
Ph.D. in Computer Science	Sept. 2015 – Jun. 2020
<ul style="list-style-type: none">• Advised by Keith Winstein and Philip Levis• Dissertation: Practical Machine Learning for Sequential Decision Problems on the Internet	
Tsinghua University	Beijing, China
B.S. in Computer Science (with <i>highest honors</i>)	Aug. 2011 – Jul. 2015
<ul style="list-style-type: none">• <i>Valedictorian</i> of Yao Class (founded by Turing Award laureate Andrew Yao)	
B.A. in Economics	Sept. 2012 – Jul. 2015
Massachusetts Institute of Technology	Cambridge, MA
Exchange Student in EECS	Feb. – May 2014

AWARDS AND HONORS

Outstanding Paper Award, USENIX NSDI	2024
Best Paper Award, APNet	2022
Applied Networking Research Prize, Internet Research Task Force	2021
Community Award, USENIX NSDI	2020
Best Paper Award, USENIX ATC	2018
Award of Excellence, Stars of Tomorrow Internship Program, Microsoft Research	2015
Outstanding Graduate of Tsinghua University	2015
Outstanding Graduate of Beijing, China	2015
National Scholarship, China	2014
Silver Prize, Yao Award, Tsinghua University	2014
Tsinghua University Comprehensive Scholarship	2013
National Endeavor Fellowship, China	2012
Tsinghua-Baidu Scholarship	2012

Scholarship for Tsinghua Xuetaang Talents Program	2011–2014
First Prize, National Senior High School Mathematical Olympiad, China	2010
First Prize, National Olympiad in Informatics in Provinces, China	2008

PUBLICATIONS

- [1] Anuj Kalia, Nikita Lazarev, Leyang Xue, Xenofon Foukas, Bozidar Radunovic, Francis Y. Yan. “Towards Energy Efficient 5G vRAN Servers.” *To appear in USENIX Symposium on Networked Systems Design and Implementation (NSDI '25)*, April 2025.
- [2] Zhiyuan He, Aashish Gottipati, Lili Qiu, Xufang Luo, Kenuo Xu, Yuqing Yang, Francis Y. Yan. “Designing Network Algorithms via Large Language Models.” *To appear in ACM Workshop on Hot Topics in Networks (HotNets '24)*, November 2024.
- [3] Siva Kakarla, Francis Y. Yan, Ryan Beckett. “Diffy: Data-Driven Bug Finding for Configurations.” *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI '24)*, June 2024.
- [4] Zibo Wang, Pinghe Li, Chieh-Jan Mike Liang, Feng Wu, Francis Y. Yan. “Autothrottle: A Practical Bi-Level Approach to Resource Management for SLO-Targeted Microservices.” *USENIX Symposium on Networked Systems Design and Implementation (NSDI '24)*, April 2024. **Outstanding Paper Award.**
- [5] Yihua Cheng, Ziyi Zhang, Hanchen Li, Anton Arapin, Yue Zhang, Qizheng Zhang, Yuhan Liu, Kuntai Du, Xu Zhang, Francis Y. Yan, Amrita Mazumdar, Nick Feamster, Junchen Jiang. “Grace: Loss-Resilient Real-Time Video through Neural Codecs.” *USENIX Symposium on Networked Systems Design and Implementation (NSDI '24)*, April 2024.
- [6] Sami Khairy, Gabriel Mittag, Vishak Gopal, Francis Y. Yan, Zhixiong Niu, Ezra Ameri, Scott Inglis, Mehrsa Golestaneh, Ross Cutler. “ACM MMSys 2024 Bandwidth Estimation in Real Time Communications Challenge.” *ACM Multimedia Systems Conference (MMSys '24)*, April 2024.
- [7] Yongzhou Chen, Ammar Tahir, Francis Y. Yan, Radhika Mittal. “Octopus: In-Network Content Adaptation to Control Congestion on 5G Links.” *ACM/IEEE Symposium on Edge Computing (SEC '23)*, December 2023.
- [8] Paramvir (Victor) Bahl, Matthew Balkwill, Xenofon Foukas, Anuj Kalia, Daehyeok Kim, Manikanta Kotaru, Zhihua Lai, Sanjeev Mehrotra, Bozidar Radunovic, Stefan Saroiu, Connor Settle, Ankit Verma, Alec Wolman, Francis Y. Yan, Yongguang Zhang (*authors in alphabetical order). “Accelerating Open RAN Research Through an Enterprise-scale 5G Testbed.” *Short paper in Annual International Conference on Mobile Computing and Networking (MobiCom '23)*, October 2023.
- [9] Zhiying Xu, Francis Y. Yan, Rachee Singh, Justin T. Chiu, Alexander M. Rush, Minlan Yu. “Teal: Learning-Accelerated Optimization of WAN Traffic Engineering.” *ACM Special Interest Group on Data Communication Conference (SIGCOMM '23)*, September 2023.
- [10] Nikita Lazarev, Tao Ji, Anuj Kalia, Daehyeok Kim, Ilias Marinou, Francis Y. Yan, Christina Delimitrou, Zhiru Zhang, Aditya Akella. “Resilient Baseband Processing in Virtualized RANs with Slingshot.” *ACM Special Interest Group on Data Communication Conference (SIGCOMM '23)*, September 2023.
- [11] Yueying Li, Daochen Zha, Tianjun Zhang, G. Edward Suh, Christina Delimitrou, Francis Y. Yan. “Mitigating Metastable Failures in Distributed Systems with Offline Reinforcement Learning.” *Short paper in International Conference on Learning Representations (ICLR '23)*, May 2023.

- [12] Michael Rudow, Francis Y. Yan, Abhishek Kumar, Ganesh Ananthanarayanan, Martin Ellis, K.V. Rashmi. “Tambur: Efficient loss recovery for videoconferencing via streaming codes.” *USENIX Symposium on Networked Systems Design and Implementation (NSDI '23)*, April 2023.
- [13] Zhengxu Xia*, Yajie Zhou*, Francis Y. Yan, Junchen Jiang (*equal contribution). “Genet: Automatic Curriculum Generation for Learning Adaptation in Networking.” *ACM Special Interest Group on Data Communication Conference (SIGCOMM '22)*, August 2022.
- [14] Jeongyoon Eo, Zhixiong Niu, Wenxue Cheng, Francis Y. Yan, Rui Gao, Jorina Kardhashi, Scott Inglis, Michael Revow, Byung-Gon Chun, Peng Cheng, Yongqiang Xiong. “OpenNetLab: Open Platform for RL-based Congestion Control for Real-Time Communications.” *Asia-Pacific Workshop on Networking (APNet '22)*, July 2022. **Best Paper Award.**
- [15] Francis Y. Yan. “Practical Machine learning for Sequential Decision Problems on the Internet.” *Ph.D. thesis, Stanford University*, June 2020.
- [16] Francis Y. Yan, Hudson Ayers, Chenzhi Zhu, Sadjad Fouladi, James Hong, Keyi Zhang, Philip Levis, Keith Winstein. “Learning *in situ*: a randomized experiment in video streaming.” *USENIX Symposium on Networked Systems Design and Implementation (NSDI '20)*, February 2020. **Community Award. IRTF Applied Networking Research Prize.**
- [17] Francis Y. Yan, Jestin Ma, Greg D. Hill, Deepti Raghavan, Riad S. Wahby, Philip Levis, Keith Winstein. “Pantheon: the training ground for Internet congestion control research.” *USENIX Annual Technical Conference (ATC '18)*, July 2018. **Best Paper Award.**
- [18] Francis Y. Yan, Songtao He, Yunxin Liu, Longbo Huang. “Optimizing Power Consumption of Mobile Games.” *Workshop on Power-Aware Computing and Systems (HotPower '15)*, October 2015.
- [19] Dongxiao Yu, Yuexuan Wang, Francis Y. Yan, Jiguo Yu, Francis C.M. Lau. “Speedup of Information Exchange using Multiple Channels in Wireless Ad Hoc Networks.” *IEEE International Conference on Computer Communications (INFOCOM '15)*, April 2015.
- [20] Francis Y. Yan, Dongxiao Yu, Yuexuan Wang, Jiguo Yu, Francis C.M. Lau. “Bounded information dissemination in multi-channel wireless networks.” *Journal of Combinatorial Optimization*, October 2014.

BOOK CHAPTERS

1. *Design of Deep Learning Systems: Theory and Practice*: college textbook written with Microsoft Research Asia.

PATENTS

1. “Artificial Intelligence for Intent-Based Networking,” filed with Ryan Beckett and Victor Bahl in May 2023. Granted U.S. Patent No. 11,968,088 in Apr. 2024.
2. “Data Streaming Protocols in Edge Computing,” filed with Ganesh Ananthanarayanan and Yuanchao Shu in Jun. 2021. Granted U.S. Patent No. 11,831,698 in Nov. 2023.
3. “Data Stream Prioritization for Communication Session,” filed with Landon Cox and Shadi Noghabi in Jun. 2021. Granted U.S. Patent No. 11,632,404 in Apr. 2023.
4. “Media Server Management for Communication Session,” filed with Landon Cox in Jun. 2021. Granted U.S. Patent No. 11,601,478 in Mar. 2023.

5. “Loss Recovery Using Streaming Codes in Forward Error Correction,” filed with Michael Rudow, Ganesh Ananthanarayanan, and Martin Ellis in Sept. 2021. Granted U.S. Patent No. 11,489,620 in Nov. 2022.
6. “Generating Adaptive Bitrate Data Streaming Neural Network Code by Large Language Model,” filed with Zhiyuan He, Lili Qiu, and Xufang Luo in Jun. 2024.
7. “Techniques for Detecting Anomalies in Data Files,” filed with Ryan Beckett and Siva Kakarla in Apr. 2024.
8. “Automatically Detecting Anomalies in Complex Configurations,” filed with Ryan Beckett and Siva Kakarla in May 2023.
9. “Bi-Level Machine-Learning-Assisted Management of Computing Resources,” filed with Zibo Wang, Pinghe Li, Mike Liang in May 2023.
10. “Power Control for Energy-Efficient 5G vRAN,” filed with Anuj Kalia, Xenofon Foukas, and Bozidar Radunovic in May 2023.
11. “CPU Power Management for Virtualized Radio Access Networks,” filed with Anuj Kalia, Xenofon Foukas, and Bozidar Radunovic in May 2023.
12. “Wireless Parameter Limits for Predicted vRAN Resource Loads,” filed with Anuj Kalia, Sanjeev Mehrotra, and Victor Bahl in May 2022.
13. “Determining Reference Signal Transmission Times,” filed with Neil Agarwal, Manikanta Kotaru, and Victor Bahl in May 2022.

SERVICE

Award Committee Member	IRTF Applied Networking Research Prize 2025
Program Committee Member	USENIX NSDI 2025
Artifact Evaluation Chair	ACM SIGCOMM 2024
Program Committee Member	ACM SIGCOMM 2024
Award Committee Member	Machine Learning and Systems Rising Stars 2024
Award Committee Member	IRTF Applied Networking Research Prize 2024
Steering Committee Member	ACM/IRTF Applied Networking Research Workshop (2024–2026)
Program Committee Member	USENIX NSDI 2024
Award Committee Member	Machine Learning and Systems Rising Stars 2023
Program Committee Chair	ACM/IRTF Applied Networking Research Workshop 2023
Program Committee Member	IEEE ICNP 2023
Program Committee Member	ACM CoNEXT 2023
Award Committee Member	IRTF Applied Networking Research Prize 2023
Poster Session Chair	USENIX NSDI 2023
Program Committee Member	USENIX NSDI 2023
Networking Area Chair	Journal of Systems Research (2022–2023)
Award Committee Member	IRTF Applied Networking Research Prize 2022
External Reviewer	USENIX NSDI 2022
Editorial Board Member	Journal of Systems Research (2021–2022)
Challenge Chair	ACM MMSys 2021
Reviewer	IEEE/ACM Transactions on Networking (2019–)
Reviewer	ACM SIGCOMM Computer Communication Review (2019)

INTERNSHIPS

Microsoft Research <i>Research Intern, Mobility and Networking Research Group</i>	Redmond, WA Jun. – Sept. 2019
Google <i>Software Engineering Intern, Congestion Control Group</i>	Mountain View, CA Jun. – Sept. 2016
Baidu <i>Research Intern, Big Data Lab</i>	Beijing, China Jun. – Aug. 2015
Microsoft Research <i>Research Intern, Wireless and Networking Group</i>	Beijing, China Mar. – Jun. 2015
The University of Hong Kong <i>Research Intern, Systems and Networking Group</i>	Hong Kong Jul. – Aug. 2014

TEACHING

Stanford University <i>Teaching Assistant for CS140: Operating Systems</i>	Stanford, CA Winter & Spring 2019
--	--------------------------------------

SELECTED INVITED TALKS

1. “Practical Machine Learning for Networked Systems.”
Peking University, Sept. 5, 2024.
Tsinghua University, Sept. 4, 2024.
2. “Machine Learning for Networked Systems.”
Princeton University, Nov. 13, 2024 (guest lecture).
University of Texas at Austin, Nov. 13, 2023 (guest lecture).
3. “Teal: Traffic Engineering Accelerated with Learning.”
Microsoft Research Asia Shanghai, May 11, 2023.
Azure Networking, Jul. 12, 2022.
Microsoft Research India, May 16, 2022.
4. “OpenNetLab 2.0 – The Next-Gen Platform for AI-assisted Networking.”
Microsoft Research Summit, Oct. 21, 2021.
5. “Networking Meets Cloud and Edge Applications” (panel discussion).
Microsoft Research Summit, Oct. 19, 2021.
6. “Grand Challenge on Bandwidth Estimation for Real-Time Communications,”
ACM MMSys '21, Sept. 29, 2021.
7. “Learning *in situ*: a randomized experiment in video streaming.”
Columbia University, Nov. 28, 2023 (guest lecture).
Columbia University, Oct. 28, 2021 (guest lecture).
IETF 110, Mar. 8, 2021.
8. “Practical Machine Learning for Sequential Decision Problems on the Internet.”
Carnegie Mellon University, Mar. 28, 2022 (guest lecture).
University of California, Irvine, Feb. 19, 2021.
Microsoft Research Redmond, September 2020.
Princeton University, Sept. 23, 2020.