Francis Y. Yan

Microsoft Building 99, room 2711
14820 NE 36th St
Redmond, WA 98052

francisyyan@gmail.com
https://francisyyan.org
Updated: Apr. 25, 2024

EMPLOYMENT

Microsoft Research — Redmond, WA
Senior Researcher, Victor Bahl's group — Aug. 2020 – Present

EDUCATION

Stanford University — Stanford, CA
Ph.D. in Computer Science — Sept. 2015 – June 2020

- Advisors: 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 highest honors) — Aug. 2011 – July 2015

- Valedictorian of Yao Class (founded by Turing Award laureate Andrew Yao)

B.A. in Economics — Sept. 2012 – July 2015

Massachusetts Institute of Technology — Cambridge, MA
Exchange Student in EECS Department — Feb. – May 2014

SERVICE

Program Committee Member	USENIX NSDI 2025
Program Committee Member	ACM SIGCOMM 2024
Program 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
Program 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 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)

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 Xuetang 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] Zhiyuan He, Aashish Gottipati, Lili Qiu, Francis Y. Yan, Xufang Luo, Kenuo Xu, Yuqing Yang. "LLM-ABR: Designing Adaptive Bitrate Algorithms via Large Language Models." Preprint *arXiv:2404.01617*, April 2024.

[2] Siva Kakarla, Francis Y. Yan, Ryan Beckett. "Diffy: Data-Driven Bug Finding for Configurations." To appear in *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, June 2024.

[3] 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)*, April 2024. 🏆 **Outstanding Paper Award.**

[4] Yihua Cheng, Ziyi Zhang, Hanchen Li, Anton Arapin, Yue Zhang, Qizheng Zhang, Yuhan Liu, 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)*, April 2024.

[5] Anuj Kalia, Nikita Lazarev, Leyang Xue, Xenofon Foukas, Bozidar Radunovic, Francis Y. Yan. "Towards Energy Efficient 5G vRAN Servers." Revise and Resubmit at *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, 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." To appear in *ACM Multimedia Systems Conference (MMSys)*, April 2024.

[7] Yongzhou Chen, Ammar Tahir, Francis Y. Yan, Radhika Mittal. "Octopus: In-Network Content Adaptation to Control Congestion on 5G Links." In *ACM/IEEE Symposium on Edge Computing (SEC)*, 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 *Proceedings of the Annual International Conference on Mobile Computing and Networking (MobiCom)*, 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." In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, September 2023.

[10] Nikita Lazarev, Tao Ji, Anuj Kalia, Daehyeok Kim, Ilias Marinos, Francis Y. Yan, Christina Delimitrou, Zhiru Zhang, Aditya Akella. "Resilient Baseband Processing in Virtualized RANs with Slingshot." In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, 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)*, 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." In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, April 2023.

[13] Zhengxu Xia*, Yajie Zhou*, Francis Y. Yan, Junchen Jiang (*equal contribution). "Genet: Automatic Curriculum Generation for Learning Adaptation in Networking." In *Proceedings of the ACM Special Interest Group on Data Communication (SIGCOMM)*, 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." In *Proceedings of the Asia-Pacific Workshop on Networking (APNet)*, July 2022. 🏆 **Best Paper Award.**

[15] Emily Marx, Francis Y. Yan, Keith Winstein. "Implementing BOLA-BASIC on Puffer: Lessons for the use of SSIM in ABR logic." Preprint *arXiv:2011.09611*, November 2020.

[16] Francis Y. Yan. "Practical Machine learning for Sequential Decision Problems on the Internet." Ph.D. thesis, *Stanford University*, June 2020.

[17] 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." In *USENIX Symposium on Networked Systems Design and Implementation (NSDI)*, February 2020. 🏆 **Community Award.** 🏆 **IRTF Applied Networking Research Prize.**

[18] 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." In *USENIX Annual Technical Conference (ATC)*, July 2018. 🏆 **Best Paper Award.**

[19] Francis Y. Yan, Songtao He, Yunxin Liu, Longbo Huang. "Optimizing Power Consumption of Mobile Games." In *Proceedings of the Workshop on Power-Aware Computing and Systems (HotPower)*, October 2015.

[20] 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." In *IEEE International Conference on Computer Communications (INFOCOM)*, April 2015.

[21] Francis Y. Yan, Dongxiao Yu, Yuexuan Wang, Jiguo Yu, Francis C.M. Lau. "Bounded information dissemination in multi-channel wireless networks." In *Journal of Combinatorial Optimization*, October 2014.

Book Chapters

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

Patents

1. "Data Streaming Protocols in Edge Computing," filed with Ganesh Ananthanarayanan and Yuanchao Shu in June 2021. Granted as U.S. Patent No. 11,831,698 in Nov. 2023.
2. "Data Stream Prioritization for Communication Session," filed with Landon Cox and Shadi Noghabi in June 2021. Granted as U.S. Patent No. 11,632,404 in Apr. 2023.
3. "Media Server Management for Communication Session," filed with Landon Cox in June 2021. Granted as U.S. Patent No. 11,601,478 in Mar. 2023.
4. "Loss Recovery Using Streaming Codes in Forward Error Correction," filed with Michael Rudow, Ganesh Ananthanarayanan, and Martin Ellis in Sept. 2021. Granted as U.S. Patent No. 11,489,620 in Nov. 2022.
5. "Automatically Detecting Anomalies in Complex Configurations," filed with Ryan Beckett and Siva Kakarla in May 2023.
6. "Bi-Level Machine-Learning-Assisted Management of Computing Resources," filed with Zibo Wang, Pinghe Li, Mike Liang in May 2023.
7. "Power Control for Energy-Efficient 5G vRAN," filed with Anuj Kalia, Xenofon Foukas, and Bozidar Radunovic in May 2023.
8. "CPU Power Management for Virtualized Radio Access Networks," filed with Anuj Kalia, Xenofon Foukas, and Bozidar Radunovic in May 2023.
9. "Artificial Intelligence for Intent-Based Networking," filed with Ryan Beckett and Victor Bahl in May 2023.
10. "Wireless Parameter Limits for Predicted vRAN Resource Loads," filed with Anuj Kalia, Sanjeev Mehrotra, and Victor Bahl in May 2022.
11. "Determining Reference Signal Transmission Times," filed with Neil Agarwal, Manikanta Kotaru, and Victor Bahl in May 2022.

Work Experience

Microsoft Research — Redmond, WA
Research Intern, Mobility and Networking Research Group — June – Sept. 2019

Google — Mountain View, CA
Software Engineering Intern, Congestion Control Group — June – Sept. 2016

Baidu — Beijing, China
Research Intern, Big Data Lab — June – Aug. 2015

Microsoft Research — Beijing, China
Research Intern, Wireless and Networking Group — Mar. – June 2015

The University of Hong Kong — Hong Kong
Research Assistant, Francis C.M. Lau's group — July – Aug. 2014

Teaching

Stanford University — Stanford, CA
Teaching Assistant for *CS140: Operating Systems* — Winter & Spring 2019

Massachusetts Institute of Technology Cambridge, MA
Grader for *6.207J: Networks* Spring 2014

SELECTED TALKS

1. **Guest Lecture:** "Learning *in situ*: a randomized experiment in video streaming," *Columbia University*, Nov. 2023.
2. **Guest Lecture:** "Machine Learning for Networked Systems," *University of Texas at Austin*, Nov. 2023.
3. "Teal: Traffic Engineering Accelerated with Learning," *Azure WAN team*, July 2022.
4. **Guest Lecture:** "Practical Reinforcement Learning on the Internet," *Carnegie Mellon University*, Mar. 2022.
5. **Guest Lecture:** "Learning *in situ*: a randomized experiment in video streaming," *Columbia University*, Oct. 2021.
6. **Panel Discussion:** "Networking Meets Cloud and Edge Applications," *Microsoft Research Summit*, Oct. 2021.
7. "Grand Challenge on Bandwidth Estimation for Real-Time Communications," *ACM MMSys*, Istanbul, Turkey, Sept. 2021.
8. "Learning *in situ*: a randomized experiment in video streaming," *IETF 110*, Mar. 2021.
9. "Practical Machine Learning for Sequential Decision Problems on the Internet," *University of California, Irvine*, Feb. 2021.
10. "Practical Machine Learning for Sequential Decision Problems on the Internet," *Princeton University*, Sept. 2020.