MINGHAO LI

SEC 4.431, 150 Western Ave, Boston, MA 02134
minghaoli@g.harvard.edu

EDUCATION

Harvard University
Ph.D., Computer Science — **Sept 2026 (Expected)**
Master of Science, Computer Science — **Feb 2025**
GPA 3.734/4.0

Cornell University
Bachelor of Science, Computer Science — **May 2021**
summa cum laude, GPA 4.0/4.0

RESEARCH INTERESTS

Systems and Networking, Machine Learning Systems

PUBLICATIONS

Li, M., Avdiukhin, D., Shahout, R., Ivkin, N., Braverman, V. & Yu, M.
FIELDING: Clustered Federated Learning with Data Drift
The 29th International Conference on Artificial Intelligence and Statistics (AISTATS'26)

Chen, P., **Li, M.,** Wan, Z., Hsiao, Y., Yu, M., Reddi, V. & Liu, Z.
OctoCache: Caching Voxels for Accelerating 3D Occupancy Mapping in Autonomous Systems
ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS'25)

Li, M., Basat, R., Vargaftik, S., Lao, C., Xu, K., Mitzenmacher, M. & Yu, M.
THC: Accelerating Distributed Deep Learning Using Tensor Homomorphic Compression
USENIX Symposium on Networked Systems Design and Implementation (NSDI'24)

Duan, Y., Zhao, X., Pan, Y., Li, S., **Li, M.**, Xu, F. & Zhang, M.
Towards Automated Safety Vetting of Smart Contracts in Decentralized Applications
ACM Conference on Computer and Communications Security (CCS'22)

Wu, H., Tian, X., **Li, M.**, Liu, Y., Ananthanarayanan, G., Xu, F. & Zhong, S.
PECAM: Privacy-Enhanced Video Streaming & Analytics via Securely-Recoverable Transformation
ACM International Conference on Mobile Computing and Networking (MobiCom'21)

Wu, H., Tian, X., Gong, Y., Su X., **Li, M.** & Xu, F.
DAPter: Preventing User Data Abuse in Deep Learning Inference Services
ACM International World Wide Web Conference (WWW'21)

UNDER REVIEWS

Li, M., Golden, A., Hsia, S., Kuchnik, M., Gangidi, A., Zhang, X., Shetty, A. J., DeVito, Z., Chu, W., He, D., Zhang, H., Hao, Y., Pang, R., Zeng, J. H., Zhang, Y., Yu, M. & Wu, C.
ScaleAcross Explorer: Exploring Communication Optimization for Scale-Across AI Model Training

RESEARCH & RELATED EXPERIENCE

Research Assistant/Ph.D. Student (Advisor: Minlan Yu) **Sept 2021 – Present**
Computer Science Department, Harvard University, Cambridge, MA

- Built Federated Learning systems with client clustering and selection strategies to mitigate statistical heterogeneity and reach the target accuracy faster. Addressed edge devices data drift and new devices arrival through timely re-clustering with low overheads. Accepted at AISTATS'26.
- Built systems that accelerate distributed Deep Neural Networks training. Focused on reducing the inter-machine synchronization overhead through gradient compression and in-network aggregation at the emerging programmable network devices. Accepted at NSDI'24.
- Built efficient path planning for autonomous systems in the cloud-edge scenario with data points caching on edge devices. Accepted at ASPLOS'25.

Research Scientist Intern (Manager: Carole-Jean Wu) **June 2025 – Feb 2026**
Meta Platforms, Inc., Menlo Park, CA

- Built a configuration optimizer for scale-across AI model training workloads distributed across multiple data center buildings and regions. Given model architecture specification, batch sizes, network topology, and hardware specifications, the optimizer searches for the configuration across the stack of parallelism placement, parallelism schedules, and network technologies that minimizes iteration time. This project is under submission.

Research Assistant (Advisors: Elaine Shi, Robbert Van Renesse) **May 2019 – May 2021**
Computer Science Department, Cornell University, Ithaca, NY

- Worked on the privacy-enhanced video streaming and analytics system PECAM. PECAM makes recoverable video transformation that eliminates visual details while preserving enough information for analytics tasks.
- Worked on the security tool DAPPSCOPE for decentralized applications (DApp) running on blockchains. Defined the high-level specifications used by DAPPSCOPE to check against a DApp's function model graphs and discover discrepancies between the UI and the contract code.
- Researched on constructing function model graphs from Ethereum smart contract data and applying node classification algorithms to the graphs. Evaluated the algorithms' performances by checking whether they computed potentially malicious nodes as similar.

COURSEWORK

- **At Harvard University**: Advanced Computer Networks, Big Data Systems, Algorithms at Ends of the Wire, High Performance Computing, Systems Security, Systems and Control, Artificial Intelligence.
- **At Cornell University**: Cloud Computing, Practicum in Operating Systems, Operating Systems, Computer Networks, Computer Architecture, Embedded Systems, Introduction to Database Systems, Introduction to Analysis of Algorithms, Data Structures and Functional Programming, Object-Oriented Programming and Data Structures, Digital Logic and Computer Organization

TEACHING EXPERIENCE

Teaching Fellow – Networking at Scale (Spring 2023)
Computer Science Department, Harvard University, Cambridge, MA

- Developed course projects. Led discussion sections and held weekly office hours to help students comprehend course materials and complete projects. Graded projects, homework, and exams.

Course Consultant – Operating Systems; Discrete Structures (Fall 2020, Spring 2021; Fall 2018)
Computer Science Department, Cornell University, Ithaca, NY

- Held weekly office hours to help students comprehend course materials and complete homework. Developed homework and handouts. Supervised study groups. Graded homework and exams.