Katherine Mohr

kmohr@cs.stanford.edu linkedin.com/in/katherinemohr katherinemohr cs.stanford.edu/~kmohr

Education

Stanford University *Stanford, CA*

PhD in Computer Science, advised by Fred Kjolstad and Keith Winstein *Sept 2025 - June 2030 (expected)*

- **GPA:** 4.03/4.0
- **Relevant Coursework:** Programming Languages, Formal Methods for Computer Systems, Efficient ML Inference
- **Awards:** EDGE Fellowship, School of Engineering PhD Fellowship

Massachusetts Institute of Technology *Cambridge, MA*

Master of Engineering & Bachelor of Science in Electrical Engineering and Computer Science *Sept 2019 - Aug 2024*

- **GPA:** 5.0/5.0
- **Teaching:** Programming Language Design, Dynamic Computer Language Eng, Graphics, Differential Equations
- **Awards:** 2023 SuperUROP Outstanding Research Award

Work Experience

Five Rings *New York, NY*

Software Developer *Aug 2024 - Sept 2025*

- Operated as a C++ generalist in a high-performance trading environment, building internal developer tooling, delivering trader-requested features, and refactoring latency-sensitive systems to enhance speed and efficiency across the codebase

Adobe *San Francisco Bay Area*

Software Developer Intern *June 2023 - Aug 2023*

- Implemented and tested different tactics for managing competing threads and prioritizing renders to improve efficiency and make background rendering less noticeable by Premiere Pro users

Computer Scientist Intern (C++) *May 2022 - Aug 2022*

- Built an end-to-end GPU inference pipeline to prevent slowdowns due to extraneous copies between the CPU and GPU

LinkedIn *San Francisco, CA*

Software Engineer Intern *June 2021 - Aug 2021*

- Redesigned messaging delivery logic and plugin schemas to improve maintainability and security

UnifyID (acquired by Prove) *Redwood City, CA*

Software Engineering Intern *Aug 2020 - Jan 2021*

- Created a pipeline to train gait models for individual users and improved methods for optimizing score policy configs
- Developed an automated performance regression pipeline to monitor the performance of ML models after version changes

Publications

[1] A. Srivatsan, Y. Deng, **K. Mohr**, et al. "Continuation-Centric Computing with Arca." *OSDI*, 2026.

[2] H. S. Shim, **K. Mohr**, and P. Levis. "The Kernel's Write: Application Read-Only Memory." *DIMES*, 2026.

Research Experience

Stanford Systems and Networking Research Group & Stanford Compilers Group *Advisors: Keith Winstein & Fred Kjolstad*

PhD Student Researcher *Sept 2025 - Present*

- Implemented a mechanism for efficiently migrating continuations between servers in Arca [1], an operating system for serverless computing that uses continuations as its isolation primitive

Stanford Information Networks Group (SING) *Advisor: Phil Levis*

Rotation Student *Apr 2026 - June 2026*

- Wrote a Linux kernel patch to use lower-cost emerging memory with slow writes as main memory alongside DRAM [2], implementing read-mostly page placement, copy-on-write migration to DRAM, and endurance-aware write scheduling

Compilers at MIT (COMMIT) Group *Advisor: Saman Amarasinghe*

Graduate Researcher *Sept 2023 - Aug 2024*

- Developed an on-stack replacement technique to transfer packet processing from optimized kernel code to an unoptimized user-space application, preventing extraneous side effects in the fallback case

Undergraduate Researcher (SuperUROP) *Sept 2022 - June 2023*

- Worked on LakePlacid, a system of compilers for optimizing networking applications with profile-guided optimization