Yifei Liu

yifeliu@cs.stonybrook.edu • Phone available on request • Greater Seattle Area
github.com/Yifei-Liu • yifei-liu.github.io • linkedin.com/in/yifei-liu

Education

Stony Brook University — Stony Brook, NY
- Ph.D. in Computer Science (Advisor: Erez Zadok) — 08/2019 – 08/2025
- M.S. in Computer Science (earned en route to Ph.D. program) — 08/2019 – 12/2021
- GPA: 3.93 / 4.0

Huazhong University of Science and Technology — Wuhan, China
- M.Eng. in Computer System Architecture (Advisor: Ke Zhou) — 09/2016 – 06/2019

Huazhong Agricultural University — Wuhan, China
- B.Eng. in Computer Science and Technology — 09/2012 – 06/2016

Experience

Google LLC — Kirkland, WA
Software Engineer — 08/2025 – Present
- Work on Google Compute Engine (GCE) and AI/ML infrastructure in Google Cloud

File systems and Storage Lab (FSL), Stony Brook University — Stony Brook, NY
Research Assistant (C/C++, File Systems, Formal Verification) — 05/2020 – 08/2025
- Developed a model-checking framework that found 15+ Linux file-system bugs
- Built a user-space file system with state save/restore, 3–28× faster than others

Samsung Semiconductor, Inc. — San Jose, CA
Storage Systems Architect Intern (C++, Databases, Storage) — 05/2022 – 08/2022
- Offloaded aggregates to SmartSSD via custom PostgreSQL paths, boosting query performance

Wuhan National Laboratory for Optoelectronics — Wuhan, China
Research Assistant (Python, Deep Learning, Cloud Storage) — 09/2016 – 06/2019
- Built a storage system with DL hashing and a graph DB, cutting query latency by 82–94%

Tencent Cloud — Shenzhen, China
Backend Developer Intern (C++, Machine Learning, Storage) — 12/2015 – 08/2016
- Built infra for long-term disk data on 10K+ servers, predicting failures with 90%+ precision

Selected Publications

Summary: 4 journal articles, 10 conference/workshop papers, 2 posters, and 2 granted patents
Google Scholar (full list): scholar.google.com/citations?user=WNu87vQAAAAJ

Journal Articles

[1] M. Antunes, T. Estro, P. Bhandari, A. Gandhi, G. Kuenning, **Y. Liu**, et al. "Kneeliverse: A universal knee-detection library for performance curves." ***SoftwareX***, 2025.

[2] T. Estro, M. Antunes, P. Bhandari, A. Gandhi, G. Kuenning, **Y. Liu**, et al. "Accelerating Multi-Tier Storage Cache Simulations Using Knee Detection." ***Performance Evaluation***, 2024.

[3] Y. Liu, Y. Wang, K. Zhou, Y. Yang, and **Y. Liu**. "Semantic-aware Data Quality Assessment for Image Big Data." ***Future Generation Computer Systems***, 2020.

Conference and Workshop Papers

[1] **Y. Liu**, et al. "Enhanced File System Testing through Input and Output Coverage." In *the 18th ACM International Systems and Storage Conference* ***(SYSTOR)***, 2025.

[2] **Y. Liu**, et al. "Metis: File System Model Checking via Versatile Input and State Exploration." In *the 22nd USENIX Conference on File and Storage Technologies* ***(FAST)***, 2024.

[3] T. Estro, M. Antunes, P. Bhandari, A. Gandhi, G. Kuenning, **Y. Liu**, et al. "Guiding Simulations of Multi-Tier Storage Caches Using Knee Detection." In *the 31st International Symposium on the Modeling, Analysis, and Simulation of Computer and Telecommunication Systems* ***(MASCOTS)***, 2023.

[4] **Y. Liu**, et al. "Input and Output Coverage Needed in File System Testing." In *the 15th ACM Workshop on Hot Topics in Storage and File Systems* ***(HotStorage)***, 2023.

[5] W. Su, **Y. Liu**, et al. "Model-Checking Support for File System Development." In *the 13th ACM Workshop on Hot Topics in Storage and File Systems* ***(HotStorage)***, 2021.

[6] Y. Liu, H. Jiang, Y. Wang, K. Zhou, **Y. Liu**, et al. "Content Sifting Storage: Achieving Fast Read for Large-scale Image Dataset Analysis." In *the 57th Design Automation Conference* ***(DAC)***, 2020.

[7] Y. Liu, Y. Wang, K. Zhou, Y. Yang, **Y. Liu**, et al. "A Framework for Image Dark Data Assessment." In *the 3rd APWeb-WAIM joint conference on Web and Big Data* ***(APWeb-WAIM)***, 2019. **(Best Paper Runner-Up)**

Patents

[1] K. Zhou, Y. Liu, Y. Yang, H. Wang, C. Li, Y. Wang, **Y. Liu**. Method for valuation of image dark data based on similarity hashing. U.S. Patent US11,138,479B2, Granted: 10/05/2021.

[2] K. Zhou, **Y. Liu**, Y. Liu, Y. Wang, Y. Yang. Image query method and system based on content semantic metadata. Chinese Patent CN110413807B, Granted: 04/20/2021.

Skills

Programming Languages

- **Fluent (>= 10,000 LoC):** C, C++, Python, Bash
- **Intermediate (>= 2,000 LoC):** SQL, Java, MATLAB, Cypher, JavaScript, Promela, Prolog

Technologies

- **Databases:** MySQL, Neo4j, PostgreSQL, HBase, Db2
- **File and Storage:** Linux VFS and kernel file systems, NFS, OpenStack Swift, HDFS
- **Virtualization:** Docker, Kubernetes, QEMU, KVM, VMware ESXi
- **Tools:** CMake, GDB, Git, Hadoop, Spark, TensorFlow, Elasticsearch, bpftrace, LTTng

Projects

Metis: a versatile framework for file system model checking (C/C++) 2020 – 2024
RefFS: a fast and reliable file system for checking reference (C++) 2020 – 2024
IOCov: input and output coverage for file system testing (Python) 2022 – 2025
CM-IOCov: improving input coverage in file system crash testing (C/C++) 2024 – 2025

Talks

- ***Enhanced File System Testing through Input and Output Coverage***
  - ACM SYSTOR 2025
- ***Metis: File System Model Checking via Versatile Input and State Exploration***
  - USENIX FAST 2024, Graduate Research Day 2024
- ***Input and Output Coverage Needed in File System Testing***
  - ACM HotStorage 2023
- ***Model-Checking Support for File System Development***
  - ACM HotStorage 2021, Dutch Model Checking Day 2022

Service

Journal Reviewer

- ACM Transactions on Architecture and Code Optimization (TACO)
- Future Generation Computer Systems
- IEEE Access

Artifact Evaluation Committee

- USENIX OSDI '23, USENIX ATC '23

Teaching

Teaching Assistant for CSE376 Advanced Systems Programming in Unix/C S '20, S '21
Teaching Assistant for CSE306 Operating Systems F '19

Contest Awards

- Finalist, Interdisciplinary Contest in Modeling (MCM/ICM), USA, 2015.
- First Prize, National Postgraduate Mathematic Contest in Modeling, China, 2014.
- First Prize, Contemporary Undergraduate Mathematical Contest in Modeling, China, 2014.