

Muxin (Murphy) Tian

✉ murphy.tian@mail.utoronto.ca |  LinkedIn |  Github |  Google Scholar |  Website

EDUCATION

University of Toronto, St George Campus.

Toronto, Canada

Bachelor of Computer Science, Arts & Science Internship Program

Sep 2021 – Jun 2026

RESEARCH EXPERIENCE

Tencent

Jan 2026 - Present

Senior Research Engineer

- Lead LLM Agent Post-training in Tencent SSV Times AI Lab.
- Research on Agentic-RL, exploring next agent paradigm to evolve product experience via large-scale RL.

RedNote

Aug 2025 - Jan 2026

Contract Research Scientist

- Initiated and led MobileBench, an industry-level benchmark reconstructing real mobile development workflows with authentic PRDs, Figma manuscripts, and production iOS/Android codebases—substantially more complex than SWE-Bench, where state-of-the-art code agents achieve only single-digit success rates.
- Formulated a diverse task distribution spanning UI layout, feature implementation, bug resolution and user experience optimization. Built a deterministic build-and-run evaluation pipeline with Appium-based checks.
- Organized collaboration with researchers from UofT, Berkeley, and UIUC, leading benchmark formulation, experimental analysis, and drafting the MobileBench paper, targeted for arXiv release this month.
- Contributed to internal on-device LLM inference benchmarking and GUI agent prototypes building, helping evaluate inference capability trade-offs between llama-cpp and MLC LLM on mobile hardware.

Harvard Systems Lab, Harvard University

Jun 2025 - Present

Research Intern, , Supervised by Prof. JunCheng Yang

- Collaborating on Nimbus, a hybrid LLM inference system that meets TTFT SLOs under bursty workloads by combining local GPU deployments with serverless APIs; implemented an offline simulator that replays real traffic traces to evaluate SLO-aware offloading policies and latency–cost trade-offs.
- Built an online Nimbus prototype as a full-stack inference gateway: integrating local vLLM/SGLang backends with multiple external providers, adding high-concurrency request handling, authentication, and rate limiting, and deploying Prometheus + Grafana observability plus a web front-end and Sphinx-based user docs.
- Cast serverless and local routing across heterogeneous providers—daily-quota subscriptions (Chutes), per-token APIs, concurrency-capped endpoints (GLM/Featherless), and local GPU deployments—as a unified cost-minimization problem, analyzing the offline optimum (ILP) and online routing policies using greedy and ski-rental algorithms.
- Operate the FreeInference service and collect anonymized LLM workload traces to support Nimbus’s trace-driven simulation and multi-provider routing experiments, with OSDI and ICML submissions in preparation.

ULab, UIUC

Feb 2025 - Feb 2026

Research Intern, , Supervised by Prof. JiaXuan You

- Co-leading OpenManus-RL, an agentic reinforcement learning framework for fine-tuning LLM agents on environments such as ALFWorld, WebShop, and GAIA; integrated benchmarks, unified rollout formats (ReAct / CoT / modular), and migrated training pipelines onto veRL / verl-agent for large-scale experiments.
- Explored online, offline, and modular RL formulations for agent finetuning, including reward redesign, training stability debugging, and baseline comparisons (ragen and verl-agent), in collaboration with OpenManus and AMD, with an ICML 2026 submission in preparation.
- Co-developed AgentDebug, a framework that analyzes LLM agent trajectories with a modular error taxonomy to localize root-cause failures and provide targeted feedback; contributed trajectory integration, debugging and ablation experiments, and writing for the ICLR 2026 submission.

Microsoft Research

Jan 2025 - Sep 2025

Research Intern, Mentored by Zhongxin Guo

- Developed AutoForge, a framework that formulates agentic system design as a search problem over architectures, task decompositions, and tools, using Monte Carlo Tree Search (MCTS) to explore large spaces of agent variants across several benchmarks (ARC-AGI, HumanEval, GSM8K, SWE-Bench, AlfWorld).
- Designed a trajectory-guided, hierarchical optimization loop (orchestration / decomposition / component / parameter) with a decision optimizer that selects where to optimize the system based on execution traces.
- Extended the search space to cost-aware assignments of heterogeneous base LLMs, using capability-specific leaderboards (reasoning, math, coding) and model pricing to choose cheaper models where yielding agentic systems that maintain performance close to the strongest model while reducing inference cost.

University of Toronto, MIE Department

May 2024 - August 2025

Contract Part-time, Software Engineer -> Research Intern, Supervised by Dr.

Daniela Rosu and Prof. Mark S. Fox

- Developed a Calendar interface using React.js and designed a GraphDB schema for Calendar in express.js.
- Transitioned to the research phase, developing BEDEO, a multi-agent ontology-based LLM recommendation framework that enhanced interpretability and recommendation quality for Canada Ontario residents.

University of Toronto

May 2023 - May 2024

Research Intern, Supervised by Prof. Ravi Shekhar and Prof. Annie Lee

- Led the creation of a multilingual text simplification dataset for English, Sinhala, Tamil, and Thai.
- Implemented and fine-tuned mT5 using the MUSS framework for unsupervised text simplification, constructing paired simplified–raw datasets to overcome supervised data scarcity.
- Conducted quantitative and qualitative analysis comparing LLM-based and unsupervised simplification methods, and contributed substantially to paper writing (submitted to LREC 2025).

WORK EXPERIENCE

Microsoft

Oct 2024 - Jan 2025

Software Engineer Intern

- Designed and implemented an internal front-end component library using Typescript and Lit for Microsoft Edge, creating docs using Storybook for team members to ensure cross-browser compatibility and usability.
- Developed Microsoft Edge Mobile on both Android and iOS platforms, addressing responsive layout inconsistencies across devices to enhance cross-platform UI reliability.

RedNote

Jul 2024 - Oct 2024

Mobile Software Engineer, iOS Platform Intern

- Drove and led the group's first on-device learning project, leveraging Naive Bayes and Microsoft's Phi-3 small language models to build an on-device chat risk monitoring API, reducing response time by **~70%** compared to cloud solutions.
- Resolved frequent app startup crashes on low-end devices by addressing iOS virtual memory limitations.
- Followed Figma designs to implement multiple iOS chat pages in Swift and been successfully released.

OPENSOURCE EXPERIENCE

ChromiumOS

 21k stars

GSOC 2025 Contributor, mentored by Sarthak Kukreti

- Enhanced the farfetchd service by implementing tracing and replay support for disk I/O during app startup.
- Built a tracing API to log page-level disk accesses and a replay mechanism for preloading data.
- Validated performance gains via tast-based Go tests with measurable reductions in cold start latency.

OpenManus-RL

 3.9k stars

Project Leader

- Contributed to the Agent-Finetuning module of **OpenManus**, a reinforcement learning framework designed to fine-tune LLM-based agents under Alfworld, Webshop and Gaia Environments.
- Building online learning RL framework based on veRL and verl-agent; integrating several benchmarks and generating rollout trajectories in ReAct, CoT and modular formats; building baseline comparison with Ragen.

libCacheSim | A high-performance C++ library for cache simulation

 270 stars

Contributor

- Bridged the gap between high-performance C++ and the JavaScript by developing a JS binding.

- Authored and published the libcachesim package to the npm registry, allowing for installation in Node.js.

SKILLS

- **Languages:** Java, Python, C/C++, JavaScript, TypeScript, Go, SQL, Swift, Objective-C, Kotlin, HTML5/CSS3.
- **Frameworks & Libraries:** React, Node.js, Lit, Spring Boot, Spring Cloud, Django, Flask, Spring Data JPA.
- **Databases & Cloud Platforms:** MySQL, PostgreSQL, MongoDB, AWS, Google Cloud Platform.
- **Tools & Systems:** Git, Jenkins, CI/CD, Docker, Kubernetes, Kafka, RabbitMQ, Nginx, Prometheus, Grafana

RELEVANT COURSES

Operating System, Database System, Distributed System, Algorithm Design and Analysis, Computer Networking, Software Engineering, Web Development, Compiler, Machine Learning, Deep Learning, Reinforcement Learning, Image Understanding, Natural Language Computing, Business of Software

Selected Publications

*denotes equal contribution

1. **Muxin Tian**, Yiyang Zhai, Juncheng Yang
“RouteWise: Latency–Cost Optimization for Multi-Provider LLM Routing”.
2. **Muxin Tian***, Zhe Wang*, Zhenwei Tang, Blair Yang, Kunlun Zhu, Honghua Dong, Hanchen Li, Guangjing Wang, Jiakuan You
“SWE-Bench Mobile: Can Large Language Model Agents Develop Industry-Level Mobile Applications?”
3. Kunlun Zhu*, **Muxin Tian***, Zijia Liu*, Bingxuan Li*, Yingxuan Yang, Jiakuan Zhang, Pengrui Han, Qipeng Xie, Fuyang Cui, Weijia Zhang, Xiaoteng Ma, Xiaodong Yu, Gowtham Ramesh, Jialian Wu, Zicheng Liu, Pan Lu, James Zou, Jiakuan You
“Where LLM Agents Fail and How They can Learn From Failures”.
4. Hannah Liu*, **Muxin Tian***, Iqra Ali, Haonan Gao, Qiaoyiwen Wu, Blair Yang, Uthayasanker Thayasivam, Annie Lee, Pakawat Nakwijit, Surangika Dayani Ranathunga, Ravi Shekhar
“OasisSimp: An Open-source Asian-English Sentence Simplification Dataset”. (LREC 2026)

Services

- ACL 2025, 2026 Reviewer
- ICML 2026 Reviewer