Zijian Zhang

Address: 5-147C, 200 Union St SE, Minneapolis, MN 55455
Email: zha00175@umn.edu, makikoqaq@gmail.com
Homepage: Google Scholar

Research Interests

My research interests focus on LLM Agents, especially post-training (RLVR/Agentic RL) and Multi-agent system for General/Coding Agents.

Education

University of Minnesota, MN, USA 2025.09 - 2030.06
Full-funded PhD Student in Department of Electrical and Computer Engineering
Advised by Prof. Mingyi Hong
Direction: LLM agents, Post-training, Multi-agent system.
GPA: 4/4

East China University of Science and Technology, Shanghai, China 2021.09 - 2025.06
Computer Science
GPA: 88/100

Publications

1. **StitchCUDA: An Automated Multi-Agent End-to-End GPU Programing Framework with Rubric-based Agentic Reinforcement Learning**
   Shiyang Li, Zijian Zhang, Winson Chen, Yuebo Luo, Mingyi Hong, Caiwen Ding
   ICML, 2026. **Co-First Author.**
2. **CudaForge: An Agent Framework with Hardware Feedback for CUDA Kernel Optimization**
   Zijian Zhang, Rong Wang, Shiyang Li, Yuebo Luo, Mingyi Hong, Caiwen Ding
   Preprint, 2026.
3. **InfantAgent-Next: A Multimodal Generalist Agent for Automated Computer Interaction**
   Bin Lei, Weitai Kang, Zijian Zhang, Winson Chen, Xi Xie, Shan Zuo, Mimi Xie, Ali Payani, Mingyi Hong, Yan Yan, Caiwen Ding
   NIPS, 2025
4. **GRAPE: Generalizing Robot Policy via Preference Alignment**
   Zijian Zhang, Kaiyuan Zheng, Zhaorun Chen, Joel Jang, Yi Li, Chaoqi Wang, Mingyu Ding, Dieter Fox, Huaxiu Yao
   ICRA 2026; ICLR Workshop, 2025.
5. **AnyPrefer: An Automatic Framework for Preference Data Synthesis**
   Yiyang Zhou, ..., Zijian Zhang, ..., Huaxiu Yao
   ICLR, 2025

Research Experience

PhD Student, Mingyi Hong's Lab @ University of Minnesota 2025.08 - 2030.05

- Introduced StitchCUDA, A multi-agent workflow for End-to-End CUDA code generation and optimization, enhanced by Rubric-based Agentic RL. It achieves SOTA performance in End-to-End CUDA code generation, **defeating GPT-5.2** by a 32B RL-based model, while avoiding reward-hacking and significantly reduce rollout overhead by decomposing atomic skills of Coding Agent.
  This paper has been submitted to ICML 2026.

- Introduced CudaForge, a simple, effective and low-cost multi-agent workflow for CUDA kernel generation and optimization. It achieves SOTA performance in KernelBench Levels 1-3, while only costs $0.3 in API and 25 minutes in single RTX 6000. This paper has been submitted to ICML 2026.
- Introduced InfantAgent-NEXT, which undertakes detailed modularization of agent workflows, tool selection, and tool execution, in favor of a modular architecture with a unified dialogue context. Paper has been accepted by NIPS 2025.
- Authored 1*_NIPS'25_.

Advisor: Prof. Mingyi Hong

Research Assistant, Huaxiu Yao's Lab @ University of North Carolina at Chapel Hill 2024.05 - 2025.01

- Introduced Anyprefer, a new method to improve VLA model through DPO and iterative training framework, eventually the robots perform better in several tasks
- Introduced GRAPE, a Trajectory-wise DPO for VLA model posttraining, which enhances the safety, efficiency, and success rate of the VLA model.
- Authored 1*_ICLR'25_, 1*_ICRA'26_.

Advisor: Prof. Huaxiu Yao

Research Assistant, Machine Learning Group @ Microsoft Research Asia 2024.10 - 2025.04

- Worked on efficiently training LLM from mixing NLP datasets by dynamic sampling.
- Worked on enhancing LLM's math reasoning ability by selecting high-quality CoT data via LLM-as-a-Judge.
- Worked on analyzing dynamic parameters (such as loss) during model training to better control the training process.

Advisor: Prof. Zhong Li

Research Assistant, InternLM2 Team @ Shanghai AI Lab 2023.11 - 2024.5

- Proposed an efficient data selection method to extract high-quality samples from the original SFT dataset of InternLM2. By fine-tuning InternLM2 on the top 10% of the highest-scoring examples, the model achieved superior performance compared to using the entire instruction dataset.
- Solved the Identity Attack problem in InternLM. Constructed a special finetune dataset to enhance the cognitive ability of the model. Introduced a benchmark to Evaluate the model's ability to resist Identity Attacks.
- Participated in the research and development of InternLM/InternLM2, which are well-known open source LLMs.

Advisor: Prof. Yining Li

Services

- Reviewer for ICML 2026 2025
- Reviewer for ICLR 2026 2025
- Reviewer for CVPR 2025 2024
- Reviewer for Neurips Workshop 2024 2024

Skills

- Programming Skills: C/C++, Python, PyTorch, LaTeX, Git
- Language Skills: Mandarin (native), English (TOFEL 100: R 27, L 24, S 22, W 27)
- Engineering Skills: Agentic RL / RLVR on VeRL framework

Awards

- Department Fellowship University of Minnesota

- Scholarship for Outstanding Students, First Prize — East China University of Science and Technology
- Outstanding student leaders — East China University of Science and Technology
- National Scholarship — Ministry of Education, China
- National Mathematics Competition for college students, First Prize — Chinese Mathematical Society
- Suzhou Industrial Park Scholarship — Suzhou government