

Gaokai Zhang

+1 412-214-3200 | gaokaiz@andrew.cmu.edu | linkedin.com/in/gaokaizhang

EDUCATION

Carnegie Mellon University

M.S. in Intelligent Information Systems

Pittsburgh, Pennsylvania

Dec 2026 (Expected)

University of Illinois at Urbana-Champaign

B.S. in Computer Engineering, GPA: 3.89/4.0

Champaign, Illinois

May 2025 (Graduated with Highest Honors)

Zhejiang University

B.E. in Electrical and Computer Engineering, GPA: 3.95/4.0

Hangzhou, China

May 2025

SKILLS

Languages & Tools: Python, C/C++, Bash, Git, SQL, Slurm, RSpec, Linux, Docker, ReactJS

Frameworks: PyTorch, Hugging Face Transformers, Spark, verl, OpenRLHF, DeepSpeed, Megatron-LM

ML/Systems: Distributed Training, Supervised Fine-tuning, Reinforcement Learning, Cloud (Azure, CloudLab)

PROFESSIONAL EXPERIENCE

Coding Agents & Data Synthesis

Graduate Research Assistant supervised by Prof. Daniel Fried from CMU

Pittsburgh, Pennsylvania

Oct 2025 – Jan 2026

- Co-developed **Hybrid-Gym**, a scalable synthetic training framework for coding agents targeting cross-task generalization, improving Qwen2.5-Coder-32B by **25.4%** on SWE-Bench Verified and **7.9%** on SWT-Bench Verified.
- Designed transferable tasks isolating reasoning, repository exploration, and implementation skills, reducing repository setup cost by **30×** compared to prior datasets (e.g., SWE-Smith) while maintaining downstream performance.

Microsoft Research Asia

Machine Learning Engineer Intern in Systems Research Group supervised by Dr. Li Lyna Zhang

Beijing, China

July 2024 – Aug 2025

- Initiated a multi-stage reinforcement learning recipe for long-context reasoning with Group Relative Policy Optimization (GRPO); boosted Qwen2.5-7B-Instruct by **48%** on long-context benchmarks, surpassing SOTA 32B Large Reasoning Models (LRMs). Accepted as **ICLR 2026 Oral**, delivering oral presentation at the conference.
- Collaborated with mentors and colleagues to scale context windows of Large Language Models (LLMs); extended Llama3-8B from 8K to 128K using only 10B tokens—**80×** fewer than Meta’s approach—while retaining 98.5% short-context accuracy.
- Led development of a parallel pipeline for browsing-history retrieval and filtering, reducing 8B-token data curation time by **80%** (35h → 7h); partnered with customer-oriented teams to deliver production-ready recommendation LLMs with 70% user-interest prediction accuracy against GPT-4o, supporting Microsoft Asia-Pacific R&D Group in product delivery.

LLM Systems & Cloud Optimization

Research Intern supervised by Prof. Fan Lai from UIUC

Champaign, Illinois

Nov 2024 – Aug 2025

- Coordinated with peers to design an MCTS planner for cloud machine learning under Service-Level Objectives (SLOs), integrating multi-fidelity profiling data from collaborators; cut search time for training an 80-layer model on 32 GPUs from 100,000s with Metis’ exhaustive search to only **80s**.
- Benchmarked and validated profiling via Megatron-LM across diverse RSpec configurations; communicated findings with peers to resolve heterogeneity mismatches across GPU (V100, A40, A100, H200, etc.) deployments.

PUBLICATIONS

- Siyuan Wang*, **Gaokai Zhang***, Li Lyna Zhang, Ning Shang, Fan Yang, Dongyao Chen, Mao Yang. LoongRL: Reinforcement Learning for Advanced Reasoning over Long Contexts. *ICLR 2026*. **Oral**.
- Yiqing Xie, Emmy Liu, **Gaokai Zhang**, Nachiket Kotalwar, Shubham Gandhi, Sathwik Acharya, Xingyao Wang, Carolyn Rosé, Graham Neubig, Daniel Fried. Hybrid-Gym: Training Coding Agents to Generalize Across Tasks. Under review at *ICML 2026*.
- Ning Shang*, Li Lyna Zhang*, Siyuan Wang, **Gaokai Zhang**, Gilsinia Lopez, Fan Yang, Weizhu Chen, Mao Yang. LongRoPE2: Near-Lossless LLM Context Window Scaling. *ICML 2025*.
- Jason Vega, Junsheng Huang*, **Gaokai Zhang***, Hango Kang*, Minjia Zhang, Gagandeep Singh. Stochastic Monkeys at Play: Random Augmentations Cheaply Break LLM Safety Alignment. (*: equal contribution)