

XIANG ZHENG

Atlanta, GA | (470) 838-9440 | xzheng353@gatech.edu
linkedin.com/in/xiang-zheng-gatech | xiangzheng2002.github.io

EDUCATION

Georgia Institute of Technology

Master of Science in Computer Science

Atlanta, GA

Aug 2026 – May 2028 (Expected)

University of California, Berkeley

EECS Visiting Student, GPA: 4.0/4.0

Berkeley, CA

Aug 2024 – Aug 2025

South China University of Technology

Bachelor of Engineering in Artificial Intelligence, GPA: 3.89/4.0, Rank: 3/80

Guangzhou, China

Sep 2021 – Jul 2025

EXPERIENCE

Microsoft Research, DK1 Group (advisors: Dr. Fangkai Yang, Dr. Lu Wang)

Research Intern

Beijing, China

Apr 2026 – Present

- Built **ReGuide**, a controllable planning framework that adapts classifier-free guidance to Tiny Recursive Models via exponential tilting and cross-window latent-state reuse; achieved a **SOTA D4RL score of 111** on Hopper with **10× fewer parameters** than the baseline, while analyzing why TRMs can outperform diffusion models on reasoning tasks.

verl – Open-source RL framework for LLM post-training

Open Source Contributor

Jun 2026 – Aug 2026

- Contributed to multimodal continuous-token rollouts for AgentLoop, including draft implementation and an end-to-end DeepEyes-v1/Qwen2.5-VL-7B GRPO recipe on **8×A100**; merged upstream in #6804.
- Contributed to **On-Policy Self-Distillation (OPSD)** and led reference-parity validation, reproducing official TRL training dynamics and Avg@12 results across AIME24/25 and HMMT25 (#6909).

Berkeley AI Research (BAIR), UC Berkeley, advised by Prof. Michael W. Mahoney

Research Assistant

Berkeley, CA

Aug 2024 – Aug 2025

- Built ODE-1000, a 1,000-task benchmark, and symbolic-equivalence diagnostics; performed full-parameter fine-tuning on five Qwen/Llama models using 900 examples across **8 GPUs**; for example, on Qwen3-0.6B, increased execution rate from **27% to 97%** and accuracy among successful runs from **63.0% to 87.6%**, while reducing mean relative L_2 error from **0.371 to 0.104**; MATH-AI @ NeurIPS 2025 (co-first author).
- Improved in-context operator learning across Transformer architectures by **50%** via scientific data augmentation, with stronger generalization on forward and inverse ODE/PDE tasks; co-first-authored the resulting [technical report](#).

Westlake University, advised by Prof. Tailin Wu

Research Assistant

Hangzhou, China

Jul 2023 – Jun 2024, Jul 2025 – Mar 2026

- Built a multi-resolution wavelet-domain diffusion framework with inference-time guidance for indirect control of physical systems; released the first 2D incompressible fluid benchmark for this setting.
- Designed a BIC-based scoring framework that constructs per-token DAGs from discrete-diffusion probabilities and folds causal structure into the sampling process, improving both accuracy and sampling throughput.
- Developed a discrete diffusion model that learns inter-token masking schedules via maximum-entropy reduction, yielding more interpretable generation order.
- Publications: [WDNO \(ICLR 2025\)](#) and [CL-DiffPhyCon \(ICLR 2025\)](#).

TECHNICAL SKILLS

Languages: Python, C++

ML / LLM: PyTorch, supervised fine-tuning (SFT), LoRA/PEFT, GRPO, model evaluation

Distributed Training & Inference: FSDP, verl, SGLang, vLLM

AWARDS

BGA Scholarship (Top 2.5%), UC Berkeley, 2024 | **Outstanding Student Trainee**, Tencent, 2024 | **First-Class Academic Scholarship** (Top 5%), SCUT, 2022