

Siqi Yang

yang1213853@gmail.com |  LinkedIn |  GitHub |  Portfolio | (+86) 130-7793-6116 (WeChat)

EDUCATION

• University of Illinois Urbana-Champaign

M.S. in Computer Science (Research Track)

GPA: 4.00/4.00

Aug.2025 - May.2027(Expected)

IL, USA

• Tsinghua University

B.Eng in Electronic Engineering

GPA: 3.86/4.00 (top 10%)

Sep.2021 - Jun.2025

Beijing, China

RESEARCH INTEREST

Machine learning researcher specializing in **coding agents** and **self-evolving AI systems**, with hands-on research experience spanning **agent harness design**, **reinforcement learning**, and **LLM post-training**.

PUBLICATIONS

Siqi Yang, Qianlan Yang, Yu-Xiong Wang, Saurabh Pujar, Martin Hirzel

What Does a Self-Evolved Harness Encode? Disentangling language and model compensation in coding agents

Preprint, Submitted to AAAI 2027

Qianlan Yang*, **Siqi Yang***, Martin Hirzel, Saurabh Pujar, Yu-xiong Wang

Interactive External Knowledge Acquisition for Software Engineering Agents

Submitted to NeurIPS 2026

Siqi Yang, Kai Yan, Alexander G. Schwing, Yu-xiong Wang.

Latent Wasserstein Adversarial Imitation Learning.

In *Proceedings of the International Conference on Learning Representations (ICLR 2026)*.

Chongkai Gao, Zhengrong Xue, Shuying Deng, Tianhai Liang, **Siqi Yang**, Lin Shao, Huazhe Xu.

RiEMann: Near Real-Time SE(3)-Equivariant Robot Manipulation without Point Cloud Segmentation.

In *Proceedings of the Conference on Robot Learning (CoRL 2024)*.

HONORS & AWARDS

• **Wang Dazhong Scholarship** | Top 3% Comprehensive Merit Scholarship, Tsinghua University

• **Tsinghua Academic Excellence Scholarship** | Awarded for top-tier academic standing

2022, 2024

• **Tsinghua Freshman Scholarship** | Admission scholarship for top students

2021

• **Silver Medal, 37th Chinese Physics Olympiad (CPhO)** | Top national ranking in physics olympiad

2020

EXPERIENCE

• IBM Research - Self Evolving Coding Agent Harness

May. 2026 – Aug. 2026

Research Intern

Yorktown Heights, NY

- Built a self-evolving coding-agent harness framework that routes every edit through a typed failure signal and records it as a falsifiable contract, making post-hoc attribution of each harness modification possible; first author, submitted to AAAI 2027.
- Ran a controlled 8-language $\times$ 3-model grid on Multi-SWE-Bench with one fixed evolution recipe and disjoint held-out splits, matching or exceeding the hand-designed mini-SWE-agent baseline in most cells while isolating language effects from model effects.

- Built automatic execution-defect detectors to explain where gains come from, showing that harness gains compensate recoverable process defects (non-compiling patches, test-file collisions, skipped target tests) whose composition shifts sharply by language, and identified two null regions the mechanism predicts.

• Interactive External Knowledge Acquisition for Software Engineering Agents Sep. 2025 – Apr. 2026

Research Assistant at UIUC, Advised by Prof. Yu-Xiong Wang

- Designed and implemented BRIDGE, a framework that augments **coding agents** with structured Internet-scale knowledge acquisition, separating local debugging from external retrieval via a dedicated retriever agent.
- Developed a multi-round belief-aligned communication protocol enabling **iterative context exchange** between the main agent and retriever, reducing unproductive local exploration by 16.3% in average trajectory length.
- Proposed an outcome-conditioned in-context reinforcement procedure that **self distills** behavioral lessons from contrastive trajectories to adapt frozen LM agents to new tool-use protocols without gradient updates.
- Achieved +7.4 pp improvement on SWE-bench Verified and consistent gains on contamination-resistant SWE-rebench across GPT-5-mini and Gemini-3-Flash backbones.

• Latent Wasserstein Adversarial Imitation Learning (LWAIL) Jul. 2024 – Mar. 2025

Research Assistant at UIUC, Project Leader, Advised by Prof. Yu-Xiong Wang



- Proposed and implemented a novel **Adversarial Imitation Learning** framework that performs state-only distribution matching via the Wasserstein distance in a learned latent space.
- Addressed the fundamental failure of Euclidean distance in capturing environment dynamics by **pre-training a dynamics-aware ground metric**.
- Achieved **SOTA performance** on locomotion benchmarks using only a single state-only expert trajectory, significantly outperforming prior KR-dual and f-divergence based LfO methods.
- Demonstrated that pre-training on just 1% unstructured, low-quality data can provide a robust embedding for highly data-efficient imitation learning.

SELECTED PROJECTS

• Collaborative Manipulation of Heterogeneous Robots and Robotic Arm Mar. 2024 - Jul. 2024

Research Assistant at TEA lab, Tsinghua University, Project Leader, Advised by Prof. Huazhe Xu.

- Engineered a synchronized multi-robot system in Mujoco by integrating perception, planning, and control modules, empowering heterogeneous robots to collaboratively perform complex object manipulation tasks that are infeasible for a single agent.
- Supported by the Tsinghua University Academic Advancement Program.

• Tsinghua University AI Challenge (THUAI) Jun. 2022 - Apr. 2023

- As a core developer in a large student-led team, contributed to both the C++ game core and front-end surface, delivering key features for the team-based programming competition and gaining significant experience in large-scale software collaboration and version control.

TEACHING AND ACADEMIC SERVICE

• CS598 ML for Software Engineering 2025 Fall

Teaching Assistant of a Graduate Level Course, focusing on AI Agents in Real-World Software Engineering

• CS527 Topic in Software Engineering 2026 Spring

Teaching Assistant of a Graduate Level Course, focusing on AI Agents and other topics in software engineering

• Reviewer: ICML 2026, ICPR 2026

SKILLS

- **Programming:** Python (Expert, PyTorch, NumPy, Scipy), C/C++ (Proficient), $\text{\LaTeX}$
- **ML:** Reinforcement Learning (Expert), Coding Agent, Robot Learning
- **Languages:** English (Fluent), Chinese (Native)