Chengxing Xie

☎ (+86)13957199364 • ✉ xiechengxing34@gmail.com • /yitianlian.github.io/

Education

Tsinghua University — **Beijing, China**
CS Ph.D. Student — *Aug. 2025 – Present*

- **Advisor:** Prof. Hongning Wang

Xidian University, — **Xi'An, China**
Undergraduate in Computer Science — *Sep. 2021 – June. 2025*

- **GPA:** 3.9/4.0
- **Rank:** 1/30 (TOP Class of Computer Science), (Transcript is here),
- **TOEFL: 110**

Research Interests

My recent research interests focus on the following areas:

- **Agentic RL**: I am particularly interested in enhancing the problem-solving capabilities of LLM agents, focusing on leveraging reinforcement learning to train them effectively. My research aims to improve their long-horizon planning abilities and memory-aware training strategies.
- **LLM RL Efficiency**: I am also focused on building efficient LLM RL training paradigms, specifically on optimizing training speed and aligning the RL algorithm with infra constraints.

Open-source Projects

I am currently actively working on the following projects:

- **slime (4.2k stars)**: slime is an LLM post-training framework designed for scalable RL training, supporting the post-training of **GLM-4.5** and **GLM-5**. I'm a **core contributor**, actively developing new RL algorithms, infrastructure features, and system capabilities.
- **SGLang**: Given the importance of LLM inference frameworks in RL training, I am also involved in the code design of RL-related features in SGLang.

Publication

1. **GLM-5: from Vibe Coding to Agentic Engineering**
   - I'm one of the **core contributors** of GLM-5.
2. **GLM-4.5: Agentic, Reasoning, and Coding (ARC) Foundation Models**
   - I'm one of the **core contributors** of GLM-4.5.
3. **Can Large Language Model Agents Simulate Human Trust Behavior?**
   - **Chengxing Xie**, Canyu Chen, Feiran Jia, Ziyu Ye, Shiyang Lai, Kai Shu, Jindong Gu, Adel Bibi, Ziniu Hu, David Jurgens, James Evans, Philip Torr, Bernard Ghanem, Guohao Li
   - Accepted in **NeurIPS 2024**, with **200+** citations. The code is Here.
4. **SWE-Fixer: Training Open-Source LLMs for Effective and Efficient GitHub Issue Resolution**
   - **Chengxing Xie**, Bowen Li, Chang Gao, He Du, Wai Lam, Difan Zou, Kai Chen
   - Accepted in **ACL 2025 Finding**.
   - Our method achieves the **highest Best@1 score** on SWE-Bench among all open-source model-based approaches. The code is Here.

Research Experience

GLM Models — **Zhipu AI**
Research Intern — *April. 2025 – Present*

- **Advisor:** Zhenyu Hou
- **For GLM-5:** Lead architect and primary developer of the unified multi-task RL training infrastructure, enabling seamless onboarding of heterogeneous agent tasks and fine-grained control over multi-task data flow. Designed and built a high-performance Token-in-Token-out gateway that bridges diverse agent workloads. Architected the RL fault-tolerance mechanisms within slime to ensure stability in long-horizon training, and established comprehensive profiling and performance metrics for PD-disaggregated RL systems. Also led the design of rollout DP-routing to further accelerate large-scale RL training.
- **For GLM-4.5:** Core contributor to the development of GLM-4.5 and an early builder of slime from its inception, participating in overall system architecture design. Propose the design of the server-based rollout architecture and buffer component tailored for agentic RL, integrated slime with multiple internal agent frameworks, and conducted extensive research on fully asynchronous RL training, including stability analysis and mitigation strategies.

Repo-level Code LLM **Openmmlab, Shanghai AI Lab**
Research Intern *July. 2024 – April. 2025*

- **Advisor:** Bowen Li, Kai Chen
- My work is focused on improving the performance of open-source LLMs on the SWE-Bench benchmark, which evaluates the ability of LLMs to solve real-world GitHub issues. We propose **SWE-Fixer**, an open-source model-based approach that achieves competitive performance compared to proprietary model-based methods and reaches **the highest Best@1 score** among all existing open-source model-based solutions.

Human-Like Reasoning Framework for Multi-Phases Planning Task **HKU**
Research Assistant *March. 2024 – July. 2024*

- **Advisor:** Prof. Difan Zou
- We developed a human-like planning framework for LLM agents to tackle complex tasks like multi-phase travel planning, integrating Strategy and Knowledge Blocks for improved information gathering and planning, achieving a 10x performance improvement with GPT-4-Turbo, as detailed in our paper here.

LLM Trust Behaviors Project **KAUST**
Visiting Student *May. 2023 – Jan. 2024*

- **Advisor:** Prof. Bernard Ghanem, Dr. Guohao Li
- Current Research Project: *Uncovering the Trust Behaviors of Large Language Model Agents* . In the realm of LLM applications, there is a hypothesis that these models are adept at simulating human behavior. My research delves into the nuances of this claim, specifically focusing on whether LLM agents can simulate human trust behaviors. This project aims to critically assess and understand the ability of LLMs to mirror such complex human behavior. The paper is **Here**, and the code is **Here**.

SenseTime Research **SenseTime (Xi'An)**
Algorithm Intern, AI for Health Team *Feb. 2022 – Nov. 2022*

- **Advisor:** Dr. Qigong Sun
- I contributed to three projects: developing a real-time cough detection algorithm for SenseTime products, designing a privacy-preserving sleep quality evaluation system.