

Longtao Zheng

Contact Information

Website: ltzheng.github.io
Email: ltzheng01@gmail.com
Google Scholar: [Longtao Zheng](#)
GitHub: [ltzheng](#)

Professional Summary

Researching LLM agents and LLM RL. Built real-world computer environments and agent frameworks for long-horizon computer tasks. Developing RL recipes for self-evolving coding agents.

Education

Nanyang Technological University Singapore
Ph.D. in Computer Science, advised by Bo An 2022.8 – 2026.2

University of Science and Technology of China China
Bachelor in Computer Science 2018.8 – 2022.6

Experience

Bytedance Singapore
Research Scientist 2026.3 – Present
RL for coding agent

- Built RL systems for training black-box LLM agents (e.g., Codex), with token-prefix-based-segment design, allowing RL for long-horizon agents with arbitrary context management.
- Trained agents to summarize context with RL for better test-time scaling.
- Developing end-to-end RL recipes for harness optimization and self-evolving coding agents.

Research Intern, worked with [Qian Liu](#) and [Zejun Ma](#) 2025.2 – 2026.3
Multi-turn RL (300+ GitHub stars, ICLR 2026 top 1% score)

- Built the end-to-end agentic RL system.
- Proposed SimpleTIR, a simple trajectory filtering strategy to stabilize multi-turn LLM RL.
- Hill-climbed on AIME from collapsed training to 7B/32B SoTA.

Skywork AI Singapore
Research Intern, worked with [Shuicheng Yan](#) 2024.1 – 2024.11
Real-world computer environments (200+ GitHub stars)

- Built a real-world computer use environment featuring universal observation space (video/image/text) and action space (GUI/CLI/API).
- Developed an online benchmark task suite and three diagnostic datasets (GUI grounding, video pre-training, and LLM judge) to evaluate computer use agents.

Long-horizon video game agent (2K+ GitHub stars)

- Designed an inference-time LLM agent framework for playing hour-long AAA video games with keyboard and mouse by only watching the screen.
- Developed the memory architecture for key event memorization and skill retrieval.
- Solved one-hour main mission in Red Dead Redemption 2 fully autonomously.

Long video generation (1K+ GitHub stars)

- Processed and filtered thousands of hours of raw videos to high-quality video-audio datasets.
- Trained a diffusion model to generate videos based on input image and audio, achieving SoTA on anti-drifting, visual quality, and multimodal synchronization.

Multi-agent RL

- Developed multi-agent RL algorithms with automatic curriculum learning for complex, sparse-reward environments like Google Research Football.
- Selected as Top Intern.

Publications

14. Dr. Kernel: Reinforcement Learning Done Right for Triton Kernel Generations
Wei Liu, Jiawei Xu, Yingru Li, **Longtao Zheng**, Tianjian Li, Qian Liu, Junxian He
ICML 2026
13. The Optimal Token Baseline: Variance Reduction for Long-Horizon LLM-RL
Yingru Li, Jiawei Xu, Ziniu Li, Jiakai Liu, Wei Liu, Yuxuan Tong, **Longtao Zheng**, Zhenghai Xue, Yaxiang Zhang, Tianle Cai, Ge Zhang, Qian Liu, Baoxiang Wang
ICML 2026
12. SimpleTIR: End-to-End Reinforcement Learning for Multi-Turn Tool-Integrated Reasoning
Zhenghai Xue*, **Longtao Zheng***, Qian Liu, Yingru Li, Zejun Ma, Bo An (*Equal contribution)
ICLR 2026
11. Towards Efficient Online Tuning of VLM Agents via Counterfactual Soft RL
Lang Feng, Weihao Tan, Zhiyi Lyu, **Longtao Zheng**, Haiyang Xu, Ming Yan, Fei Huang, Bo An
ICML 2025
10. Cradle: Empowering Foundation Agents Towards General Computer Control
Cradle Team (Core contributor)
ICML 2025
9. AgentStudio: A Toolkit for Building General Virtual Agents
Longtao Zheng*, Zhiyuan Huang*, Zhenghai Xue, Xinrun Wang, Bo An, Shuicheng Yan (*Equal contribution)
ICLR 2025
8. FinAgent: A Multimodal Foundation Agent for Financial Trading
Wentao Zhang, ..., **Longtao Zheng**, Xinrun Wang, Bo An
KDD 2024
7. Synapse: Trajectory-as-Exemplar Prompting with Memory for Computer Control
Longtao Zheng, Rundong Wang, Xinrun Wang, Bo An
ICLR 2024
6. True Knowledge Comes from Practice: Aligning LLMs with Embodied Environments via RL
Weihao Tan, Wentao Zhang, Shanqi Liu, **Longtao Zheng**, Xinrun Wang, Bo An
ICLR 2024
5. MEMO: Memory-Guided Diffusion for Expressive Talking Video Generation
Longtao Zheng*, Yifan Zhang*, Hanzhong Guo, Jiachun Pan, Zhenxiong Tan, Jiahao Lu, Chuanxin Tang, Bo An, Shuicheng Yan (*Equal contribution)
TMLR (J2C Certificate)
4. Multi-Agent Multi-Game Entity Transformer
Rundong Wang, ..., **Longtao Zheng**, ..., Bo An
DAI 2024 (Best Paper Award)
3. Controlling Type Confounding in Ad Hoc Teamwork with Instance-wise Teammate Feedback Rectification
Dong Xing, Pengjie Gu, Qian Zheng, Xinrun Wang, Shanqi Liu, **Longtao Zheng**, Bo An, Gang Pan
ICML 2023

Manuscripts

2. Dr. MAS: Stable Reinforcement Learning for Multi-Agent LLM Systems
Lang Feng, **Longtao Zheng**, Shuo He, Fuxiang Zhang, Bo An
arXiv preprint arXiv:2602.08847
1. Towards Skilled Population Curriculum for Multi-Agent Reinforcement Learning
Rundong Wang*, **Longtao Zheng***, Wei Qiu, Bowei He, Bo An, Zinovi Rabinovich, Yujing Hu,
Yingfeng Chen, Tangjie Lv, Changjie Fan (*Equal contribution)
arXiv preprint arXiv:2302.03429

**Professional
Service**

Conference Reviewer: ICLR, NeurIPS, ICML, CVPR, ICCV, WACV

Journal Reviewer: TMLR, IJCV