

Xiao Liang

[✉ Email](#) | [🐙 Github](#) | [🌐 Website](#) | [🌐 LinkedIn](#) |

Education

University of California, Los Angeles (UCLA) *Sept. 2025 - June. 2029 (expected)*
Ph.D | Advisor: [Kai-Wei Chang](#), [Ying Nian Wu](#) *California, USA*

Tsinghua University (THU) | *GPA: 3.98/4.0* *Sept. 2022 - June. 2025*
M.Sc. | *Data Science and Information Technology* *Beijing, China*

Sun Yat-sen University (SYSU) | *GPA: 4.11/5.0* *Sept. 2018 - June. 2022*
B.Sc. | *Electronic Information Science and Technology* *Guangzhou, China*

Publications

RL and Reasoning

- [1] **Xiao Liang***, Zhong-Zhi Li*, Yeyun Gong, Yelong Shen, Ying Nian Wu, Zhijiang Guo, Weizhu Chen. *Beyond Pass@1: Self-Play With Variational Problem Synthesis Sustains RLVR*. **ICLR 2026**. [\[Paper\]](#) [\[Code\]](#) [\[Page\]](#)
- [2] **Xiao Liang***, Zhong-Zhi Li*, Yeyun Gong, Yang Wang, Hengyuan Zhang, Yelong Shen, Ying Nian Wu, Weizhu Chen. *SwS: Self-aware Weakness-driven Problem Synthesis in Reinforcement Learning for LLM Reasoning*. **NeurIPS 2025**. [\[Paper\]](#) [\[Code\]](#) [\[Page\]](#)
- [3] **Xiao Liang**, Zhong-Zhi Li, Zhenghao Lin, Eric Hancheng Jiang, Hengyuan Zhang, Yelong Shen, Kai-Wei Chang, Ying Nian Wu, Yeyun Gong, Weizhu Chen. *Training LLMs for Divide-and-Conquer Reasoning Elevates Test-Time Scalability*. **ACL 2026 Main**. [\[Paper\]](#) [\[Code\]](#)
- [4] Boyan Duan, **Xiao Liang**, Shuai Lu, Yaoxiang Wang, Yelong Shen, Kai-Wei Chang, Ying Nian Wu, Mao Yang, Weizhu Chen, Yeyun Gong. *Gold-Medal-Level Olympiad Geometry Solving with Efficient Heuristic Auxiliary Constructions*. **ACL 2026 Main**. [\[Paper\]](#) [\[Code\]](#) [\[Benchmark\]](#)
- [5] Zhong-Zhi Li*, **Xiao Liang***, Zihao Tang, Lei Ji, Peijie Wang, Haotian Xu, Xing W, Haizhen Huang, Weiwei Deng, Yeyun Gong, Ying Nian Wu, Zhijiang Guo, Xiao Liu, Fei Yin, Cheng-Lin Liu. *TL;DR: Too Long, Do Re-weighting for Efficient LLM Reasoning Compression*. **ACL 2026 Main**. [\[Paper\]](#) [\[Code\]](#)
- [6] Xumeng Wen, Zihan Liu, Shun Zheng, Zhijian Xu, Shengyu Ye, Zhirong Wu, **Xiao Liang**, Yang Wang, Junjie Li, Ziming Miao, Jiang Bian, Mao Yang. *Reinforcement Learning with Verifiable Rewards Implicitly Incentivizes Correct Reasoning in Base LLMs*. **ICLR 2026**. [\[Paper\]](#)
- [7] Fanding Huang*, Guanbo Huang*, Xiao Fan, Yi He, **Xiao Liang**, Xiao Chen, Qinting Jiang, Faisal Nadeem Khan, Jingyan Jiang, Zhi Wang. *Semantic-Space Exploration and Exploitation in RLVR for LLM Reasoning*. **ACL 2026 Findings**. [\[Paper\]](#) [\[Code\]](#) [\[Project\]](#)
- [8] Yiyao Yu, Yuxiang Zhang, Dongdong Zhang, **Xiao Liang**, Hengyuan Zhang, Xingxing Zhang, Ziyi Yang, Mahmoud Khademi, Hany Awadalla, Junjie Wang, Yujiu Yang, Furu Wei. *Chain-of-Reasoning: Towards Unified Mathematical Reasoning in Large Language Models via a Multi-Paradigm Perspective*. **ACL 2025 Main**. [\[Paper\]](#)
- [9] Zhong-Zhi Li, Duzhen Zhang, Ming-Liang Zhang, Jiaxin Zhang, Zengyan Liu, Yuxuan Yao, Haotian Xu, Junhao Zheng, Pei-Jie Wang, Xiuyi Chen, Yingying Zhang, Fei Yin, Jiahua Dong, Zhiwei Li, Bao-Long Bi, Ling-Rui Mei, Junfeng Fang, **Xiao Liang**, Zhijiang Guo, Le Song, Cheng-Lin Liu. *From*

System 1 to System 2: A Survey of Reasoning Large Language Models. **TPAMI**. [Paper] [Code]

- [10] Eric Hancheng Jiang*, **Xiao Liang***, Yikai Zhang, Yingjia Wan, Mengting Li, Haikang Deng, Alexander K Taylor, Ying Nian Wu, Andrea L. Bertozzi, Kai-Wei Chang, Raghu Meka, Matthew Sottile, Nanyun Peng, Amit Sahai, Terence Tao, Wei Wang. *Toward Research-Level Formal Mathematics with Large Language Models*. Preprint 2026. [Page]
- [11] Jiawen Xie, Pengyu Cheng, **Xiao Liang**, Yong Dai, Nan Du. *Chunk, Align, Select: A Simple Long-sequence Processing Method for Transformers*. **ACL 2024 Main**. [Paper] [Code]
- [12] Eric Hanchen Jiang, Levina Li, Rui Sun, **Xiao Liang**, Yubei Li, Yuchen Wu, Haozheng Luo, Hengli Li, Zhi Zhang, Zhaolu Kang, Kai-Wei Chang, Ying Nian Wu. *Agent Q-Mix: Selecting the Right Action for LLM Multi-Agent Systems through Reinforcement Learning*. Preprint 2026. [Paper] [Code]

Post-training and Inference

- [13] Yi Cheng, **Xiao Liang**, Yeyun Gong, Wen Xiao, Song Wang, Yuji Zhang, Wenjun Hou, Kaishuai Xu, Wenge Liu, Wenjie Li, Jian Jiao, Qi Chen, Peng Cheng, Wayne Xiong. *Integrative Decoding: Improve Factuality via Implicit Self-consistency*. **ICLR 2025**. [Paper] [Code]
- [14] Hengyuan Zhang*, Shiping Yang*, **Xiao Liang**, Chenming Shang, Yuxuan Jiang, Chaofan Tao, Jing Xiong, Hayden Kwok-Hay So, Ruobing Xie, Angel X. Chang, Ngai Wong. *Find Your Optimal Teacher: Personalized Data Synthesis via Router-Guided Multi-Teacher Distillation*. **ACL 2026 Main**. [Paper] [Code]
- [15] Hengyuan Zhang, Zhihao Zhang, Mingyang Wang, Zunhai Su, Yiwei Wang, Qianli Wang, Shuzhou Yuan, Ercong Nie, Xufeng Duan, Qibo Xue, Zeping Yu, Chenming Shang, **Xiao Liang**, Jing Xiong, Hui Shen, Chaofan Tao, Zhengwu Liu, Senjie Jin, Zhiheng Xi, Dongdong Zhang, Sophia Ananiadou, Tao Gui, Ruobing Xie, Hayden Kwok-Hay So, Hinrich Schütze, Xuanjing Huang, Qi Zhang, Ngai Wong. *Locate, Steer, and Improve: A Practical Survey of Actionable Mechanistic Interpretability in Large Language Models*. **ACL 2026 Findings**. [Paper] [Code]
- [16] Hengyuan Zhang, Xinrong Chen, Yingmin Qiu, **Xiao Liang**, Ziyue Li, Guanyu Wang, Weiping Li, Tong Mo, Wenyue Li, Hayden Kwok-Hay So, Ngai Wong. *GuiLoMo: Allocating Expert Number and Rank for LoRA-MoE via Bilevel Optimization with Guided Selection Vectors*. **EMNLP 2025 Findings**. [Paper] [Code]
- [17] Eric Hancheng Jiang*, Guancheng Wan*, Sophia Yin*, Mengting Li*, Yuchen Wu, **Xiao Liang**, Xinfeng Li, Yizhou Sun, Wei Wang, Kai-Wei Chang, Ying Nian Wu. *Dynamic Generation of Multi LLM Agents Communication Topologies with Graph Diffusion Models*. **ACL 2026 Main**. [Code]
- [18] **Xiao Liang**, Xinyu Hu, Simiao Zuo, Jimi He, Yu Wang, Victor Ye Dong, Yeyun Gong, Kushal S. Dave, Yi Liu, Qiang Lou, Shao-Lun Huang, Jian Jiao. *What You See Is What You Get: Entity-Aware Summarization for Reliable Sponsored Search*. **NeurIPS SafeGenAI 2025**. [Paper]
- [19] Tao Shi*, **Xiao Liang***, Yaoyuan Liang, Xinyi Tong, Shao-Lun Huang. *SSLCL: An Efficient Model-Agnostic Supervised Contrastive Learning Framework for Emotion Recognition in Conversations*. Preprint 2023. [Paper] [Code]
- [20] Hengyuan Zhang, Xinrong Chen, Zunhai Su, **Xiao Liang**, Jing Xiong, Wendong Xu, He Xiao, Chaofan Tao, Wei Zhang, Ruobing Xie, Lei Jiang, Hayden Kwok-Hay So, Ngai Wong. *Beyond Outliers: A Data-Free Layer-wise Mixed-Precision Quantization Approach Driven by Numerical and Structural Dual-Sensitivity*. Preprint 2026. [Paper]

Pre-training

- [21] **Xiao Liang***, Xinyu Hu*, Simiao Zuo, Yeyun Gong, Qiang Lou, Yi Liu, Shao-Lun Huang, Jian Jiao. *Task Oriented In-Domain Data Augmentation*. **EMNLP 2024**. [Paper]

- [22] Kailai Yang*, Xiao Liu, Lei Ji, Hao Li, **Xiao Liang**, Zhiwei Liu, Yeyun Gong, Peng Cheng, Mao Yang. *Data Mixing Agent: Learning to Re-weight Domains for Continual Pre-training*. **ACL 2026 Main**. [Paper]
- [23] Microsoft Sigma Team. *SIGMA: Differential Rescaling of Query, Key and Value for Efficient Language Models*. Preprint 2025. [Paper]
- [24] Microsoft SIGMA v-Team. *SIGMA: An AI-Empowered Training Stack on Early-Life Hardware*. Preprint 2025. [Paper] [Code]
- [25] Microsoft SIGMA v-Team. *SIGMA-MoE-Tiny: Technical Report*. Preprint 2025. [Paper] [Code] [Page]

Multimodal Understanding

- [26] **Xiao Liang***, Tao Shi*, Yaoyuan Liang, Te Tao, Shao-Lun Huang. *Exploring Iterative Refinement with Diffusion Models for Video Grounding*. **ICME 2024**. [Paper] [Code]
- [27] Yaoyuan Liang*, **Xiao Liang***, Yansong Tang, Zhao Yang, Ziran Li, Jingang Wang, Wenbo Ding, Shao-Lun Huang. *CoSTA: End-to-End Comprehensive Space-Time Entanglement for Spatio-Temporal Video Grounding*. **AAAI 2024**. [Paper]
- [28] Yaoyuan Liang, Zhuojun Cai, Jian Xu, Guanbo Huang, Yiran Wang, **Xiao Liang**, Jiahao Liu, Ziran Li, Jingang Wang, Shao-Lun Huang. *Unleash Region Understanding in Intermediate Layers for MLLM-based Referring Expression Generation*. **NeurIPS 2024**. [Paper] [Code]

(* indicates equal contribution)

Professional Experience

Microsoft Research & CoreAI

Nov. 2023 - Aug. 2025

Mentor: *Dr. Yeyun Gong, Dr. Yelong Shen, Dr. Weizhu Chen*

- Pre-training and continual learning for large language models (LLMs).
- Reinforcement learning to enhance reasoning in LLMs.

Tencent Artificial Intelligence Lab (AI Lab, NLP)

Mar. 2023 - Sep. 2023

Mentor: *Dr. Pengyu Cheng*

Shenzhen, China

- Utilizing reinforcement learning for long sequence modeling in Transformers.
- Dialog generation and task-specific instruction data generation.

Awards and Honors

- Graduate Dean's Scholar Award (GDSA) of UCLA, Amount: \$14,500 Sep. 2025
- Outstanding Master Graduate Thesis of Tsinghua University. May. 2025
- Second Prize Scholarship of Tsinghua University (top 10%). 2023 - 2024
- Outstanding Graduate Thesis of Sun Yat-sen University (top 5%). Jun. 2022
- Outstanding Graduate Student of Sun Yat-sen University (top 10%). Jun. 2022
- Second Prize Scholarship of Sun Yat-sen University (top 15%) for 3 times. 2019 - 2022

Skills

Programming Languages: Python | MATLAB | Shell | LATEX | C++/C | HTML

Programming Tools: Git | PyTorch | Docker | Linux Ops | Vim

English Skills: Toefl 102

Academic Services

Reviewer of :

- ACL Rolling Review (ARR) {Jan 26, Oct 25}
- International Conference on Learning Representations (ICLR) {26}
- Annual Conference on Neural Information Processing Systems (NeurIPS) {24,25}
- International Conference on Machine Learning (ICML) {25}
- Empirical Methods in Natural Language Processing (EMNLP) {24}