

Xiao Lin

201 N Goodwin Ave, Urbana, IL 61801
Phone: 217-974-8925 | Website: scholar | Email: xiaol13@illinois.edu

RESEARCH INTEREST

I am a third-year Computer Science Ph.D. student from the University of Illinois Urbana-Champaign. I am deeply engaged in advancing **controllable and reliable AI systems under dynamic environments**. My research interests include sequential modeling, graph machine learning, large language models / agents, and time series analysis. Recently, I investigate **how to harness agentic reasoning** through the integration of memory, agentic reinforcement learning, and multi-agent systems.

EDUCATION

University of Illinois Urbana-Champaign

Ph.D. in Computer Science Advisor: Dr. Hanghang Tong

Champaign, IL, USA

Aug. 2023 – May. 2028 (expected)

Peking University

B.E. in Computer Science Advisor: Dr. Yisen Wang

Beijing, China

Sep. 2019 – Jun. 2023

INDUSTRY EXPERIENCE

Research Intern

Core Machine Learning Team, Meta. Manager: Dr. Yajuan Wang.

Sunnyvale, CA, USA

May. 2025 – Oct. 2025

- Researched on Mixture of Experts for long-sequence recommendation.

Research Intern

Mobile Intelligence Business Team, SenseTime. Manager: Dr. Tong Li.

Beijing, China

May. 2021 – Nov. 2021

- Researched on human posture detection and facial image generation.

SELECTED PUBLICATION

- P1. **Xiao Lin**, Philip Li, Zhichen Zeng, et al. ALERT: Zero-shot LLM Jailbreak Detection via Internal Discrepancy Amplification. ACL 2026.
- P2. **Xiao Lin**, Robert Tang, Weilin Cong, et al. Mixture of Sequence: Theme-Aware Mixture-of-Experts for Long-Sequence Recommendation. **WWW 2026 Oral**.
- P3. **Xiao Lin**, Zhichen Zeng, Tianxin Wei, Zhining Liu, Hanghang Tong. CATS: Mitigating correlation shift for multivariate time series classification. TMLR 2026.
- P4. **Xiao Lin**, Zhining Liu, Ze Yang, et al. MORALISE: A Structured Benchmark for Moral Alignment in Visual Language Models. ICML 2026.
- P5. **Xiao Lin**, Mingjie Li, Yisen Wang. MADE: Graph Backdoor Defense with Masked Unlearning. Under review.
- P6. **Xiao Lin**, Zhining Liu, Dongqi Fu, Ruizhong Qiu, Hanghang Tong. Backtime: Backdoor attacks on multivariate time series forecasting. **NeurIPS 2024 Spotlight. (top 2% paper)**
- P7. **Xiao Lin**, Jian Kang, Weilin Cong, Hanghang Tong. Bemap: Balanced message passing for fair graph neural network. Learning on Graphs Conference 2024.
- P8. Tianxin Wei, Ting-Wei Li, Zhining Liu, Xuying Ning, Ze Yang, Jiaru Zou, Zhichen Zeng, Ruizhong Qiu, **Xiao Lin**, Dongqi Fu, Zihao Li, Mengting Ai, Duo Zhou, Wenxuan Bao, Yunzhe Li, Gaotang Li, Cheng Qian, Yu Wang, Xiangru Tang, et al. Agentic reasoning for large language models. Under review.
- P9. Zhining Liu, Tianyi Wang, **Xiao Lin**, Penghao Ouyang, Gaotang Li, Ze Yang, et al. Do VLMs Have a Moral Backbone? A Study on the Fragile Morality of Vision-Language Models. Under review.
- P10. Zihao Li*, Zhichen Zeng*, **Xiao Lin***, Feihao Fang, Yanru Qu, Zhe Xu, Zhining Liu, Xuying Ning, Tianxin Wei, Ge Liu, Hanghang Tong, Jingrui He. Flow Matching Meets Biology and Life Science: A Survey. Under review.

- P11. Zhichen Zeng, Qi Yu, **Xiao Lin**, Ruizhong Qiu, Xuying Ning, Tianxin Wei, Yuchen Yan, Jingrui He, Hanghang Tong. Harnessing Consistency for Robust Test-Time LLM Ensemble. Under review.
- P12. Zhichen Zeng, Wenxuan Bao, **Xiao Lin**, Ruizhong Qiu, et al. Subspace Alignment for Vision-Language Model Test-time Adaptation. Under review.
- P13. Shuhang Lin, Zhencan Peng, Lingyao Li, **Xiao Lin**, Xi Zhu, Yongfeng Zhang. Cache Mechanism for Agent RAG Systems. Under review.
- P14. Zhichen Zeng, Ruizhong Qiu, Wenxuan Bao, Tianxin Wei, **Xiao Lin**, Yuchen Yan, et al. Pave Your Own Path: Graph Gradual Domain Adaptation on Fused Gromov-Wasserstein Geodesics. Under review.
- P15. Zhichen Zeng, Mengyue Hang, Xiaolong Liu, Xiaoyi Liu, **Xiao Lin**, Ruizhong Qiu, Tianxin Wei, Zhining Liu, et al. Hierarchical LoRA MoE for Efficient CTR Model Scaling. Under review.
- P16. Duo Zhou, Yuji Zhang, Tianxin Wei, Ruizhong Qiu, Ke Yang, **Xiao Lin**, Cheng Qian, Jingrui He, Hanghang Tong, Heng Ji, Huan Zhang. Geometric-Disentanglement Unlearning. Under review.
- P17. Yuanchen Bei, Tianxin Wei, Xuying Ning, Yanjun Zhao, Zhining Liu, **Xiao Lin**, Yada Zhu, et al. Mem-Gallery: Benchmarking Multimodal Long-Term Conversational Memory for MLLM Agents. ACL 2026.
- P18. Baoyu Jing, Sanhorn Chen, Lecheng Zheng, Boyu Liu, Zihao Li, Jiaru Zou, Tianxin Wei, Zhining Liu, Zhichen Zeng, Ruizhong Qiu, **Xiao Lin**, Yuchen Yan, Qi Yu, Dongqi Fu, Jingrui He, Hanghang Tong TRQA: Time Series Reasoning Question And Answering Benchmark. Under review.
- P19. Zihao Li, **Xiao Lin**, Zhining Liu, Jiaru Zou, et al. Language in the flow of time: Time-series-paired texts weaved into a unified temporal narrative. ICLR 2026.
- P20. Zhining Liu, Ze Yang, **Xiao Lin**, Ruizhong Qiu, Tianxin Wei, et. al. Breaking Silos: Adaptive Model Fusion Unlocks Better Time Series Forecasting. ICML 2025.
- P21. Ruizhong Qiu, Jun-Gi Jang, **Xiao Lin**, Lihui Liu, Hanghang Tong. TUCKET: A tensor time series data structure for efficient and accurate factor analysis over time ranges. VLDM 2025.
- P22. Dongqi Fu, Yada Zhu, Zhining Liu, Lecheng Zheng, **Xiao Lin**, Zihao Li, Liri Fang, Katherine Tieu, et al. ClimateBench-M: A Multi-Modal Climate Data Benchmark with a Simple Generative Method. CIKM 2025.

PROJECT EXPERIENCES

Methodology Advancements for Generalizable Insights into Complex Systems <i>Research Assistant</i>	DARPA <i>Spring 2026 – Present</i>
NetSafe: Towards a Computational Foundation of Safe Graph Neural Networks <i>Research Assistant</i>	NSF <i>Fall 2024 – Present</i>
Sequential User Understanding <i>Research Assistant</i>	Meta <i>Fall 2025</i>
Ethical Understanding of Vision Language Models <i>Research Assistant</i>	Fidelity <i>Spring 2025 – Present</i>
AI for Future Agricultural Resilience, Management, and Sustainability <i>Research Assistant</i>	NIFA <i>Fall 2023 – Fall 2024</i>
Safety Alignment of Large Language Models <i>Research Assistant</i>	VISA Research <i>Fall 2024 – Fall 2025</i>
Programming Fabric for Geospatial Representation, Fusion and Federation <i>Research Assistant</i>	IBM <i>Fall 2023 – Present</i>
Modeling Adversarial Activities <i>Research Assistant</i>	DARPA <i>Fall 2023 – Spring 2024</i>

MENTORING EXPERIENCE

- Mr. Philip Li (Junior, UIUC CS)** 2025 – 2026
- Mentored him on research in LLM safeguard mechanisms against jailbreak attacks, leading to an ARR submission to ACL 2026.
 - Currently guiding him in exploring the dynamics of collaboration in multi-agent systems, with promising research outcomes.
 - Currently, successfully admitted to the Ph.D. program in CS at UIUC.
- Mr. Tianyi Wang (Junior, UIUC CS)** 2025 – 2026
- Mentored him on research in VLM robustness under moral-aware perturbations, leading to an ARR submission to ACL 2026.
 - Currently, successfully admitted to the M.S. program in CS at UIUC.
- Mr. Haotian Li (Junior, UIUC ECE)** 2025 – 2026
- Mentoring him on research in LLM moral value alignment evaluation, resulting in a submission currently under review at ICML 2026.
 - Currently guiding him in formulating research questions and setting up initial experiments to build a strong foundation for future research.

TEACHING EXPERIENCE

- Teaching Assistant** Champaign, IL, USA
CS 412 Data Mining Spring 2026
- Teaching Assistant** Champaign, IL, USA
CS 410 Text Information Retrieval Spring 2025

AWARDS

- A1. Outstanding Graduate Award (top 2%) 2023
- A2. Academic Excellence Award (top 10%) 2022
- A3. Third-class Scholarship (top 20%) 2022
- A4. Academic Excellence Award (top 10%) 2022
- A5. Excellent Scientific Research Award (top 10%) 2021
- A6. Third-class Scholarship (top 20%) 2021
- A7. First Prize in the Chinese Physics Olympiad (top 0.1%) 2018, 2019

SERVICE

Conference reviewer: ICML, ICLR, NeurIPS, CVPR, ECCV, KDD, CIKM, TheWebConf, DSAA, LOG.
Journal reviewer: ACM Computing Survey, IEEE Transaction on Computational Biology and Bioinformatics.