

Xiangqi (Shawn) Wang

PhD Student · LLM Continual Learning

xwang76@nd.edu | +1 (574) 323-0940 | Website | Google Scholar | LinkedIn

RESEARCH PROFILE

PhD student at the University of Notre Dame advancing **LLM continual learning**: enabling models to adapt from experience while preserving reliable reasoning and aligned behavior. Research spans test-time training, self-distillation, reinforcement learning, and adaptive agent systems, connecting principled algorithms with efficient deployment. First-author work includes **NeurIPS Spotlight**, **EMNLP Main and Findings**, **TMLR**, and **WSDM**.

RESEARCH & INDUSTRY EXPERIENCE

- SAP Remote, USA Fall 2026
Part-time iXp Research Intern
- IBM Research Yorktown Heights, NY May–Aug. 2026
Research Intern · AI Functions in SQL & LLM Systems
 - Developed **SAGE**, a unified algebra and adaptive execution framework for SQL AI functions, optimizing model selection and query execution for quality, cost, and latency.

RESEARCH CONTRIBUTIONS

- Continual adaptation.** Developed **PrimeTTT** for meta-learned test-time adaptation and **LGSD** for locality-constrained self-distillation, improving reasoning while limiting behavioral drift.
- Reliable reasoning & alignment.** Designed **CE-PO** for coherence-aware RL and **Lag-DPO** for multi-objective preference alignment, balancing truthfulness, anti-sycophancy, empathy, and creativity.
- Adaptive inference & online learning.** Built **AdaReasoner**, an RL controller for adaptive reasoning, and developed fairness-aware contextual bandits with sublinear-regret guarantees.

EDUCATION

- University of Notre Dame South Bend, IN Sep. 2024–Present
PhD Student, Computer Science and Engineering · Advisor: Prof. Xiangliang Zhang
- University of Science and Technology of China Hefei, China Sep. 2020–Jun. 2024
B.Sc. in Computer Science, First-Class Honors · School of the Gifted Young
Prior mentorship: Qi Liu (advisor); Huazheng Wang and Jian Kang (remote research); Zhuoran Yang (causal inference).

RESEARCH MENTORSHIP & TEACHING

- Student mentorship** Guide student-led projects from problem selection and method design through experiments and writing, turning research ideas into reproducible results and publication-ready manuscripts.
- Teaching assistant** Graduate Machine Learning, University of Notre Dame; Computer Programming A and Data Analysis and Practice, USTC.

TECHNICAL EXPERTISE

- Methods** Continual learning; test-time adaptation; meta-learning; RL post-training; preference optimization; causal inference; contextual and combinatorial bandits; semantic query processing.
- Libraries & tools** PyTorch, Transformers, VERL, LLaMA-Factory, SQL, Git, Docker.

PUBLICATIONS & RESEARCH MANUSCRIPTS

FIRST-AUTHOR / PUBLISHED & ACCEPTED

PrimeTTT: Priming Test-Time Training for LLMs via Meta-Learned Initialization

Xiangqi Wang, D. Aishan, Y. Zhou, Y. Huang, X. Zhang · [EMNLP 2026 Main](#) [paper]

Right Answers for the Right Reasons: Coherence-Enhanced Policy Optimization for Large Language Models

Xiangqi Wang, Y. Huang, Y. Zhou, X. Luo, K. Guo, Y. Ma, X. Zhang · [EMNLP 2026 Main](#) [paper]

Preprint title: [Causally-Enhanced Reinforcement Policy Optimization](#).

From Evasive Servants to Dignified Peers via Dual Preference Optimization

Xiangqi Wang, Y. Huang, H. Zhuang, K. Guo, X. Zhang · [Findings of EMNLP 2026](#) [paper]

Preprint title: [Dual Optimal: Make Your LLM Peer-like with Dignity](#).

AdaReasoner: Adaptive Reasoning Enables More Flexible Thinking

Xiangqi Wang, Y. Huang, Y. Wang, X. Luo, K. Guo, Y. Zhou, X. Zhang · [NeurIPS 2025 Spotlight](#) [paper]

Fair Online Influence Maximization

Xiangqi Wang, S. Zhang, J. E. A. Escamilla, Q. Wu, X. Zhang, J. Kang, H. Wang · [TMLR 2025](#) [paper]

WildlifeLookup: A Chatbot Facilitating Wildlife Management with Accessible Data and Insights

Xiangqi Wang, T. Yang, J. Rohr, B. Scheffers, N. Chawla, X. Zhang · [WSDM 2025](#) [paper]

FIRST-AUTHOR / SUBMISSIONS & PREPRINTS

Locality-Guided Self-Distillation

Xiangqi Wang, D. Aishan, K. Guo, Y. Huang, X. Zhang · [ICLR 2027 Submission](#)

SAGE: A Unified Algebra and Self-Adaptive Execution for AI Functions in SQL

Xiangqi Wang, N. H. Pham, O. Hassanzadeh, D. Subramanian, X. Zhang · [SIGMOD Submission](#)

CO-AUTHORED / PUBLISHED & ACCEPTED

AutoDavis: Automatic and Dynamic Evaluation Protocol of Large Vision-Language Models on Visual Question-Answering

H. Bao, Y. Huang, Y. Wang, J. Ye, Xiangqi Wang, et al. · [KDD 2026](#) [paper]

ProbeLLM: Automating Principled Diagnosis of LLM Failures

Y. Huang, Z. Jiang, Y. Ma, Y. Jiang, Xiangqi Wang, et al. · [ICML 2026](#) [paper]

DyFlow: Dynamic Workflow Framework for Agentic Reasoning

Y. Wang, Z. Xu, Y. Huang, Xiangqi Wang, et al. · [NeurIPS 2025](#) [paper]

Dynamic workflow construction for tool-using and multi-agent reasoning, supported by theoretical analysis.

CLIPerase: Efficient Unlearning of Visual-Textual Associations in CLIP

T. Yang, L. Dai, Xiangqi Wang, M. Cheng, Y. Tian, X. Zhang · [ACL 2025](#) [paper]

Contributed to implementation and research visualization.

Dissecting Logical Reasoning in LLMs: A Fine-Grained Evaluation and Supervision Study

Y. Zhou, J. Ye, Z. Ling, ..., Xiangqi Wang, X. Zhang · [Findings of EMNLP 2025](#) [paper]

On the Trustworthiness of Generative Foundation Models: Guideline, Assessment, and Perspective

Y. Huang, C. Gao, S. Wu, H. Wang, Xiangqi Wang, et al. · [ICLR 2026](#) [paper]

Contributed to the assessment framework, methodology, and large-scale experiments.

Think it Image by Image: Multi-Image Moral Reasoning of Large Vision-Language Models

C. Gao, Y. Huang, Xiangqi Wang, S. Wu, N. Chawla, X. Zhang · [CIKM 2025](#)

Contributed to manuscript development and experimental analysis.

Exposing and Patching the Flaws of Large Language Models in Social Character Simulation

Y. Huang, Z. Yuan, Y. Zhou, K. Guo, Xiangqi Wang, et al. · [COLM 2025](#) [paper]

Related preprint: [Social Science Meets LLMs: How Reliable Are Large Language Models in Social Simulations?](#) (2024). Contributed simulation scenarios, reliability metrics, and analysis.

CO-AUTHORED PREPRINTS & ONGOING WORK

Getting Better at Working With You: Compiling User Corrections into Runtime Enforcement for Coding Agents

Y. Zhou, K. Guo, H. Zhuang, Xiangqi Wang, Y. Huang, et al. · Preprint, 2026 [\[paper\]](#)

SenseMath: Do LLMs Have Number Sense? Evaluating Shortcut Use, Judgment, and Generation

H. Zhuang, Xiangqi Wang, Y. Shen, Y. Cheng, X. Zhang · Preprint, 2026 [\[paper\]](#)

ValueLence: A Dashboard for In-Depth Value Probing of LLMs

Y. Huang, J. Ye, Z. Liu, Y. Li, Xiangqi Wang, et al. · [\[code\]](#)

Contributed to the evaluation and analysis pipeline.

OPEN-SOURCE SYSTEMS & RESEARCH PROJECTS

CausalIRL Extended VERL with Jacobian-based causal scores and a parallel reward backend for policy optimization. [\[link\]](#)

AdaReasoner Official implementation with a modular reasoning controller and reproducible training and inference pipelines. [\[link\]](#)

Fair Online Influence Maximization Combinatorial bandit codebase with scalable optimization, online experiments, and evaluation on public social graphs, including Pokec. [\[link\]](#)

Knowledge Network of Wildlife Built a searchable national wildlife knowledge interface and graph-RAG chatbot for species, habitat, and policy queries, supporting an NSF-funded AI-for-science project. [\[link\]](#)

NSF Proto-OKN & ACCESS Contributed to data and knowledge integration; GPU research infrastructure through [NSF ACCESS](#). [\[link\]](#)

HONORS, AWARDS & RESEARCH SUPPORT

NSF DISCOVER ACCESS research computing allocation (2025).

OpenAI Researcher Access Program — research API credits (2025).

Anthropic Research computing support.

Competitive awards Chinese Mathematics Competition, Provincial First Prize (2022); ICPC China Regional, competition prize (2021).

Academic recognition USTC Outstanding Student (annual award); SGY 87-00 Scholarship nominee.

ACADEMIC SERVICE

Conference reviewer NeurIPS, ICML, ICLR, KDD, ACL Rolling Review (2025–2026); WWW (2025); KDD (2024).

Journal reviewer Transactions on Machine Learning Research (2025); IEEE Transactions on Computational Social Systems.

Additional service IEEE NASC Competition.

TALKS & PARTICIPATION

Talks & participation NSF Proto-OKN Annual Meeting; SASC Symposium; ACM WSDM Tutorial; EMNLP 2024; Third Reinforcement Learning Conference (remote).