Qiushui Xu

(+1) 814-321-6058 xqsisme@gmail.com Homepage Google Scholar LinkedIn

Professional summary

Machine Learning Scientist with expertise in LLM post-training and agentic reasoning. Proven track record in transforming complex domain challenges into scalable ML problems, specializing in high-fidelity synthetic data generation and RL-driven alignment to enhance model reasoning, safety, and system-level autonomy.

Education

Penn State University Aug 2021 – Expected 2026
Ph.D. in Industrial Engineering - Operations Research *State College, PA*

Lanzhou University Aug 2016 – Jun 2020
B.S. in Mathematics, Cuiying Honors College *Gansu, China*

Experience

Google, X Lab (Tapestry) May 2025 – May 2026
AI Resident Intern (Advisor: Rob Sare)

- Formulated long-context document parsing as an **agentic reasoning task**. Developed a multi-stage orchestration framework to decompose complex information extraction into iterative retrieval and verification steps, achieving **93.9% precision** in zero-shot scenarios.
- Developed a robust **Text-to-Code-to-Diagram** pipeline featuring a **self-correction loop** and a unique **hybrid evaluation framework**—integrating **deterministic, simulation-based, and LLM-as-a-judge** metrics—to effectively mitigate structural hallucinations and achieve **90% accuracy** in domain-specific tasks.
- Orchestrated a high-fidelity synthetic data engine to generate 3k high-quality reasoning trajectories. Implemented a **multi-dimensional filtering strategy** to ensure data quality. Leveraged these samples for **Supervised Fine-Tuning** via LoRA, significantly improving Qwen3-8B's domain-specific success rate from **30% to 80%**.
- Architected an enterprise-grade internal platform and ensured **Safety and Policy compliance** through constrained decoding and output-filtering guardrails. Managed the full lifecycle from backend optimization to a production-ready interface, effectively automating manual drafting workflows with high reliability.

Microsoft Research, Machine Learning Group May 2024 – Aug 2024
Research Scientist Intern (Advisor: Song Lei)

- Conducted in-depth analysis of **PLM-initialized Decision Transformer** by generating heatmaps of key, query, and value matrices, identifying a novel phenomenon termed **Markov heads**.
- Evaluated the influence of Markov and Non-Markov heads using **Mixture of Attention**, and introduced **GPT-DTMA** to enable adaptive learning across environments without training from scratch.
- Proposed **In-context Compositional Q-Learning** (ICQL), a novel framework that leverages the **in-context learning** capabilities of linear Transformers to infer local linear Q-functions from small, retrieved transition sets.
- Validated that **ICQL** improves the ability to learn accurate value functions in offline settings: improving performance in kitchen tasks by up to **16.4%**, and in Gym and Adroit tasks by up to **8.6%** and **6.3%**.

Penn State University Aug. 2021 – May 2025
PhD Research Assistant (Advisor: Necdet Serhat Aybat)

- Designed and implemented high-efficiency algorithms to tackle large-scale deep learning challenges.
- Proposed innovative **tuning-free accelerated first-order algorithms** for deterministic and stochastic scenarios.
- Integrated adaptive learning rates to achieve outstanding theoretical and practical performance, while pioneering backtracking conditions and establishing convergence results for minimax problems.
- Executed experiments on distributed robust optimization for neural networks using **PyTorch**, achieving state-of-the-art performance surpassing adaptive algorithms like **Adam**.

Selected Publications

In-Context Compostional Q-Learning for Offline Reinforcement Learning
Qiushui Xu*, Yuhao Huang, Yushu Jiang, Wenliang Zheng, Lei Song, Jinyu Wang, Jiang Bian*
Accepted by ***International Conference on Learning Representation*** 2026

Unveiling Markov heads in Pretrained Language Models for Offline Reinforcement Learning
Qiushui Xu**, Wenhao Zhao*, Linjie Xu, Lei Song, Jinyu Wang, Chunlai Zhou, Jiang Bian*
Accepted by ***International Conference on Machine Learning*** 2025

A Stochastic GDA Method With Backtracking For Solving Nonconvex Concave Minimax Problems
Qiushui Xu*, Xuan Zhang, Necdet Serhat Aybat and Mert Gurbuzbalaban*