

LIAOYAQI WANG

Phone: (+1) 5102297160

Email: wangliaoyaqi@gmail.com

EDUCATION



University of Texas at Austin

Aug. 2026 - May. 2030

Doctor of Philosophy (PhD), Computer Science,



Johns Hopkins University

Aug. 2024 - May. 2026

Master of Science, Computer Science, GPA: 3.82/4.0

Coursework: *Trustworthy Machine Learning; Machine Learning Theory; Intro to Human Language Technology; Self Supervised Learning.*



University of California, Berkeley

Jan. 2023 - Jul. 2023

BGA Student, Computer Science, GPA: 3.67/4.0

Coursework: *Optimization Models in Engineering; Designing, Visualizing and Understanding Deep Neural Networks; Data Structure*



Xi'an Jiaotong University

Sep. 2020 - July. 2024

Bachelor of Engineering, Artificial Intelligence (Honor Track), GPA: 3.63/4.0

Coursework: *Machine Learning; Computer Vision; Natural Language Processing; Modern Control; Digital Signal Processing; Probability Theory and Stochastic Process, Reinforcement Learning*

PUBLICATION

(Equal contribution is marked with an asterisk)

- [1] **Liaoyaqi Wang**, Chunsheng Zuo, William Jurayj, Benjamin Van Durme, Anqi Liu. Process Supervision of Confidence Margin for Calibrated LLM Reasoning. *In Proceedings of the Third Conference on Language Modeling*.
- [2] Chunsheng Zuo, **Liaoyaqi Wang**, William Jurayj, William Fleshman, Benjamin Van Durme. Rethinking LoRA Memory Through the Lens of KV Cache Compression. (EMNLP26 under review)
- [3] Yeyao Tao, **Liaoyaqi Wang**, Elias Stengel-Eskin. Video-CER: Annotation-Free Grounded Video Understanding via Counterfactual Evidence Rewards (AAAI26 under review)
- [4] Huicong Zhang, **Liaoyaqi Wang**, Arda Uzunoglu, Tianjian Li, Marc Marone, Benjamin Van Durme. Know Your Sources: Data Selection Matters when Rewriting for Data-Constrained Pretraining. (EACL26 under review)
- [5] Yi Cao*, **Liaoyaqi Wang***, Jieneng Chen, Benjamin Van Durme, Alan Yuille, Paulette Clancy. ARIA: A Causal-Aware Framework for Rescuing LLM Reasoning in Trustworthy Materials Discovery. (KDD26 Accepted)
- [6] **Liaoyaqi Wang***, Zhengping Jiang*, Anqi Liu, and Benjamin Van Durme. 2025. Always tell me the odds: Fine-grained conditional probability estimation. *In Proceedings of the Second Conference on Language Modeling*.
- [7] Wenhao Zheng*, **Liaoyaqi Wang***, Dongshen Peng, Hongxia Xu, Yun Li, Hongtu Zhu, Tianfan Fu, and Huaxiu Yao. LIFTED: Multimodal clinical trial outcome prediction via large language models and mixture-of-experts. In Christos Christodoulopoulos, Tanmoy Chakraborty, Carolyn Rose, and Violet Peng, editors, *Findings of the Association for Computational Linguistics: EMNLP 2025*, pages 7503–7517, Suzhou, China, November 2025. Association for Computational Linguistics.
- [8] **Liaoyaqi Wang** and Minhao Cheng. 2024. GuardEmb: Dynamic watermark for safeguarding large language model embedding service against model stealing attack. *In Findings of the Association for Computational Linguistics: EMNLP 2024*, pages 7518–7534, Miami, Florida, USA. Association for Computational Linguistics.

RESEARCH EXPERIENCE

Johns Hopkins University

Position: Graduate Research Assistant

Enhancing Calibrated Reasoning with Margin Process Reward

Advisor: Prof. Anqi Liu, Prof. Benjamin Van Durme

Nov. 2025 – Apr. 2026

- Proposed a calibration-aware RL framework that jointly optimizes reasoning correctness and confidence reliability via margin-based process reward over intermediate reasoning states.
- Designed a lightweight probe-based confidence estimator using Monte Carlo forced-completion as training targets.
- Demonstrated improved calibration across in-domain math and out-of-domain code, science, and logic benchmarks while preserving reasoning accuracy.
- Validated downstream utility of calibrated confidence for conformal risk control and confidence-weighted answer aggregation.

Scientific Feasibility Assessment Agent

Advisor: Prof. Benjamin Van Durme, Prof. Anqi Liu

May. 2025 – Present

- Architected a modular multi-agent system (using LangGraph/LangChain) that orchestrates iterative loops of literature search, reflection, and summarization to validate materials science claims.

- Identified and formalized the "Contextual Tunneling" failure mode, demonstrating that naive RAG integration can degrade reasoning in scientific domains.
- Led the evaluation strategy and system optimization, achieving top-ranking performance in the DARPA dry run, outperforming five competing research institutions.

Enhancing Fine-grained Probabilistic Calibration and Reasoning in Large Language Models

Advisor: Prof. Benjamin Van Durme, Prof. Anqi Liu

Sep. 2024 – Apr. 2025

- Developed synthetic data for generous domain and long context in LongChain beyond human annotation dataset.
- Proposed a training framework enabling direct probability modeling with modern LLM backends.
- Constructed an objective and comprehensive evaluation suite by designing subproblems from existing datasets and frameworks that admit an intuitive probabilistic interpretation.

The University of North Carolina at Chapel Hill

Multimodal Clinical Outcome Prediction via LLMs and Mixture-of-Experts

Advisor: Prof. Huaxiu Yao

Jun. 2024 – Aug. 2024

- Developed unified noise-resilient encoders to extract information from language descriptions transformed from multi-modal data.
- Employed a Sparse Mixture-of-Expert framework to refine representations and dynamically integrated different modality representations for prediction.
- Achieved State-Of-The-Art performance in predicting clinical trial outcomes on the Hint and CTO datasets.

Predicting Gene Expression from Ultra-Long Sequences using Selective State-Space Models

Advisor: Prof. Huaxiu Yao

Feb. 2024 – Jun. 2024

- Constructed a bi-directional Mamba model, a selective state-space sequence model, for gene expression prediction, effectively handling long gene sequences over 100kb.
- Incorporated biological prior information, such as peak, exon and gene locations, to enhance model performance.
- Collaborated with a multidisciplinary team to integrate the sequence model into existing bioinformatics pipelines.

Hong Kong University of Science and Technology

Dynamic Backdoor Watermarking for Ownership Protection via Text Embeddings

Advisor: Prof. Minhao Cheng

Aug. 2023 – Feb. 2024

- Developed a method to protect the ownership of text embeddings by injecting dynamic backdoor watermark.
- Balanced watermark utility and detectability, even when facing out-of-distribution data.
- Outperformed current protective approaches in better balancing embedding utility and watermark verification.

JOB EXPERIENCE

Johns Hopkins University, Data Science and AI Institute

Feb. 2026 - Present

Position: Graduate Research Assistant

China Telecom, Shaanxi Branch

Jun. 2023 - Aug. 2023

Position: Software Developer

SERVICES

Reviewer of ICML 2026, Secondary Reviewer of NeurIPS 2025, AACL 2025, ICLR 2025, ARR March 2026, COLM 2026

CODE SKILLS

Languages: Python, JAVA, C++, MATLAB, Git, Latex

Frameworks: PyTorch, Hugging Face Transformers, LangGraph, LangChain, vLLM, ROS

HONORS AND AWARDS

2025 LLM Hackathon for Applications in Materials and Chemistry, Visionary Awards (top 30)

Sep. 2025

National College Students Innovation Training Program, Excellent

Jun. 2023

National Optoelectronic Design Competition, Provincial Second Prize

Jul. 2022

Mathematical Contest in Modeling, Provincial Second Prize

Jan. 2021