Ziyao Mou

Johns Hopkins University
Computer Science
3400 North Charles Street
Baltimore, MD 21218

Email: zmou1@jhu.edu
Homepage: https://ziyaomou.github.io
Github:https://github.com/ZiyaoMou
Linkedin: https://www.linkedin.com/in/ziyao-mou/

Education

Johns Hopkins University *Aug. 2024–Dec. 2025 (expected)*
M.S.E. in Computer Science
Advisor: Eric Nalisnick
Master's Thesis: Learning to Defer under Expert Drift.

Huazhong University of Science and Technology *Sept. 2018–June 2022*
B.Eng. in Computer Science
Bachelor's Thesis: Design and Implementation of a Resource Scheduling System for Workflow Applications on a Serverless Platform.

Research

Ziyao Mou, Andrea Wynn, Eric Nalisnick. *Feb. 2025-Present*
Learning to Defer under Expert Drift. [pdf](Under Review)

- Designed an expert-aware learning-to-defer framework that models temporal expert drift by conditioning the deferral head on recent expert outcomes via a sequence model.
- Learned history-aware strategies that proactively route hard or critical cases to humans when reliability dips, improving safety under time-varying expertise.
- Evaluated on image classification, medical diagnosis, and hate-speech detection with simulated and real expert drift.
- Achieved higher system accuracy and better deferral efficiency than prior L2D baselines.

Ziyao Mou, Ha Manh Bui, Anqi Liu. *June 2025-Present*
Multiple Domain Recalibration for Model-based Reinforcement Learning.

- Designed and implemented a multi-domain calibration framework for model-based reinforcement learning to improve robustness under domain shifts.
- Compared mix-up training and multi-calibration strategies to align epistemic uncertainty across environments, improving both in-domain reliability and out-of-domain generalization in model-based reinforcement learning.
- Conducted experiments in the MuJoCo continuous-control suite, demonstrating faster convergence, improved uncertainty calibration, and higher cumulative rewards.

Selected Projects

Ziyao Mou, Shuwen Bai *Oct. 2025 – Jan. 2026*
Risk-calibrated Robust Adversarial Reinforcement Learning. [pdf]
Instructors: Dr. Vladimir Braverman, Dr. Zih-Yun (Sarah) Chiu

- Developed a calibrated risk model by learning a logistic failure predictor and applying histogram binning to correct miscalibration, yielding reliable risk estimates for gating adversarial actions.
- Designed a risk-aware adversary gating mechanism that selectively applies perturbations only in high-risk states, improving stability and robustness during adversarial policy training.
- Achieved consistent performance improvements over standard RARL across benchmarks, with faster convergence and higher returns.

Ziyao Mou, Yiran Zhong, Elena Kote *Sept. 2024–Jan. 2025*
Frequency and Representation-Based Exploration of Mathematical Reasoning in Large Language Models [pdf]
Instructors: Dr. Daniel Khashabi

- Built a 57,600-sample arithmetic dataset with multi-format numeric/textual representations to isolate reasoning vs. memorization effects in LLaMA-2-7B.
- Conducted large-scale frequency analysis using the Infini-gram trillion-token index, revealing strong correlations between log-frequency and model accuracy, especially for multi-digit operations.
- Demonstrated that number representation significantly impacts performance while operator representation does not, providing evidence toward symbolic reasoning capabilities rather than pure memorization.

Employments

Software Development Engineer *Jul. 2023–Jul. 2024*
Alibaba Group, Taobao Search Team
Led full-stack development of Taobao's AI-powered shopping assistant, delivering a cross-platform user interface with intelligent features such as product comparison, recommendations, and travel planning.

Software Development Engineer *Jul. 2022–Jul. 2023*
Alibaba Group, Platform Infrastructure Team
Built a user experience reporting & telemetry platform that ingests multimedia feedback and logs, with a BERT-based sentiment clustering pipeline and screenshot similarity detection.

Software Development Intern *Jun. 2021–Aug. 2021*
Alibaba Group, Buyer Base Team
Designed and implemented an edge-computing streaming rendering full-stack architecture using Serverless technologies to improve responsivenessr on low-end and older devices.

Teaching

Johns Hopkins University
Course Assistant, Machine Learning: Artificial Intelligence System Design and Development, Fall 2025

Awards

Bronze Award, National Innovation and Entrepreneurship Competition(Challenge Cup)
Team Lead, AidocX
Huazhong University of Science and Technology, China
First Prize, Physics Competition for High School Students
Zhejiang University Affiliated High School, China

Skills

Programming: Python, Java, C/C++, SQL, TypeScript, HTML/CSS.
Cloud & DevOps & Big Data Technologies: AWS (S3, Lambda), Docker, Kubernetes, MongoDB, Redis.
Machine Learning & Data: PyTorch, NumPy, Pandas, Scikit-learn.
Tools & Version Control: Git.

Languages

English (Advanced), Mandarin Chinese (Native).