

Bolian Li

PhD Candidate

✉ li4468@purdue.edu

🐙 lblaoke

🌐 lblaoke.github.io

Education

PhD, Computer Science

Purdue University

Advisor: Ruqi Zhang

GPA: 3.8/4.0

2022 – 2026

West Lafayette, IN, USA

BE, Computer Science and Technology

Tianjin University

Advisor: Changqing Zhang

GPA: 3.8/4.0

2018 – 2022

Tianjin, China

Experiences

Apple Machine Learning Research

Research Intern, working with Cheng-Yu Hsieh and Ting-Yao Hu
MoE Optimization for Agentic RL

March 2025 – August 2026

Amazon

Applied Scientist Intern

Preference Alignment Generalization

May 2025 – August 2025

University of Illinois Urbana-Champaign

Research Intern, working with Hanghang Tong
Graph Representation Learning

May 2021 – November 2021

Institute of Computing Technology, Chinese Academy of Sciences

Research Intern, working with Shuhui Wang
Video Object Detection

January 2021 – February 2021

Research Interests

My research develops statistical foundations for stable and efficient post-training frameworks. I currently focus on improving the performance boundary of LLM RL, including RL entropy control, and agentic MoE RL. More broadly, I work on test-time scaling, preference alignment, multimodal LLM safety, and Bayesian/statistical methods for robust ML.

Technical Background

Model Training: Long-horizon agentic RL on LLMs up to 35B, with multi-node clusters and decoupled inference & training engines.

Framework Building: Building rollout functions and environment servers for interactive coding agents from scratch. Deeply refactoring RL frameworks like HF Transformers, verl, and slime.

Theoretical Analysis: Bayesian statistics, rejection sampling, and gradient analysis.

Publications

Uniform-Correct Policy Optimization: Breaking RLVR's Indifference to Diversity

Preprint, 2026

Anamika Lochab, Bolian Li, Ruqi Zhang

Addressing Performance Saturation for LLM RL via Precise Entropy Curve Control

Preprint, 2026

Bolian Li, Yifan Wang, Yi Ding, Anamika Lochab, Ananth Grama, Ruqi Zhang

SARL: Label-Free Reinforcement Learning by Rewarding Reasoning Topology

Preprint, 2026

Yifan Wang, *Bolian Li*, David Cho, Ruqi Zhang, Fanping Sui, Ananth Grama

Learning Self-Correction in Vision-Language Models via Rollout Augmentation

ICML 2026

Yi Ding, Ziliang Qiu, *Bolian Li*, Ruqi Zhang

Why Reasoning Fails to Plan: A Planning-Centric Analysis of Long-Horizon Decision Making in LLM Agents

Preprint, 2026

Zehong Wang, Fang Wu, Hongru Wang, Xiangru Tang, *Bolian Li*, Zhenfei Yin, Yijun Ma, Yiyang Li, Weixiang Sun, Xiusi Chen, Yanfang Ye

Structure-R1: Dynamically Leveraging Structural Knowledge in LLM Reasoning through Reinforcement Learning

Preprint, 2025

Junlin Wu, Xianrui Zhong, Jiashuo Sun, *Bolian Li*, Bowen Jin, Jiawei Han, Qingkai Zeng

From Personal to Collective: On the Role of Local and Global Memory in LLM Personalization

ACL 2026

Zehong Wang, Junlin Wu, Zhaoxuan Tan, *Bolian Li*, Xianrui Zhong, Zheli Liu, Qingkai Zeng

DRIFT: Learning from Abundant User Dissatisfaction in Real-World Preference Learning

ICLR 2026

Yifan Wang, *Bolian Li*, Junlin Wu, Zhaoxuan Tan, Zheli Liu, Ruqi Zhang, Ananth Grama, Qingkai Zeng

Reward-Shifted Speculative Sampling Is An Efficient Test-Time Weak-to-Strong Aligner

EMNLP 2025 Main

Bolian Li, Yanran Wu, Xinyu Luo, Ruqi Zhang

More is Less: The Pitfalls of Multi-Model Synthetic Preference Data in DPO Safety Alignment

COLM 2025

Yifan Wang, Runjin Chen, *Bolian Li*, David Cho, Yihe Deng, Ruqi Zhang, Tianlong Chen, Zhangyang Wang, Ananth Grama, Junyuan Hong

Cascade Reward Sampling for Efficient Decoding-Time Alignment

COLM 2025

*Bolian Li**, Yifan Wang*, Anamika Lochab*, Ananth Grama, Ruqi Zhang

Bayesian Computation in Deep Learning

MCMC Handbook

Wenlong Chen*, *Bolian Li**, Ruqi Zhang, Yingzhen Li

Stacey: Promoting Stochastic Steepest Descent via Accelerated ℓ_p -Smooth Nonconvex Optimization

ICML 2025

Xinyu Luo*, Site Bai*, *Bolian Li**, Petros Drineas, Ruqi Zhang, Brian Bullins

Making Reliable and Flexible Decisions in Long-tailed Classification

TMLR

Bolian Li, Ruqi Zhang

ETA: Evaluating Then Aligning Safety of Vision Language Models at Inference Time

ICLR 2025
Yi Ding, *Bolian Li*, Ruqi Zhang

Entropy-MCMC: Sampling from Flat Basins with Ease

ICLR 2024
Bolian Li, Ruqi Zhang

Long-tailed Classification from a Bayesian-decision-theory Perspective

AABI 2023
Bolian Li, Ruqi Zhang

Trustworthy Long-tailed Classifications

CVPR 2022
Bolian Li, Zongbo Han, Haining Li, Huazhu Fu, Changqing Zhang

Graph Communal Contrastive Learning

WWW 2022 Oral
Bolian Li, Baoyu Jing, Hanghang Tong

Services

Reviewer	CVPR 2026; ICML 2025, 2026; NeurIPS 2024, 2026; ICLR 2025; AISTATS 2024, 2025; UAI 2024, 2025
Guest Lecturer	Purdue CS 57800 - Statistical Machine Learning
Teaching Assistant	CS 57800; CS 37300; CS 38003; CS 17600

Awards

Silver Reviewer Award, <i>ICML</i>	2026
Tan Xuguang Alumni Scholarship, <i>Tianjin University</i>	2020
Tianjin University Vision Education Fund, <i>Tianjin University</i>	2019