Licheng Liu

Mathematics, B.S. Candidate, Imperial College London
Ph.D. in Computer Science, University of Wisconsin–Madison — Expected Sep 2026 – Jun 2031
licheng.liu22@imperial.ac.uk | (+86)18913966201 | lichengliu03.github.io

PREPRINT

(* denotes equal contribution or alphabetical ordering)

1. **Real-Time Audio Tool Bench: Evaluating Tool Use in Multi-Turn Voice Interaction**
   Licheng Liu*, Jingsong Sun*, Huanzhi Mao, and Joseph E. Gonzalez.
   In submission to NeurIPS 2026.
2. **Let's Try Again: Eliciting Multi-Turn Reasoning in Language Models via Simplistic Feedback** [Paper] [Github] [Model] [Website]
   Licheng Liu*, Zihan Wang*, Linjie Li, Chenwei Xu, Yiping Lu, Han Liu, Avirup Sil, and Manling Li.
   NeurIPS 2025 Workshop on MTI-LLM.
3. **CocoaBench: An Evaluation Framework for General Agents with Compositional Cognitive Abilities** [Paper] [Website]
   CocoaBench Team: Shibo Hao, Zhining Zhang, Zhiqi Liang, Tianyang Liu, Yuheng Zha, Qiyue Gao, Jixuan Chen, Zilong Wang, Zhoujun Cheng, Haoxiang Zhang, Junli Wang, Hexi Jin, Boyuan Zheng, Kun Zhou, Yu Wang, Feng Yao, **Licheng Liu**, Yijiang Li, Zhifei Li, Zhengtao Han, Pracha Promthaw, Tommaso Cerruti, Xiaohan Fu, Ziqiao Ma, Jingbo Shang, Lianhui Qin, Julian McAuley, Eric P. Xing, Zhengzhong Liu, Rupesh Kumar Srivastava, and Zhiting Hu.
   In submission.
4. **RAGEN: Understanding Self-Evolution in LLM Agents via Multi-Turn Reinforcement Learning** [Paper] [Github] [Website]
   Zihan Wang*, Kangrui Wang*, Qineng Wang*, Pingyue Zhang*, Linjie Li*, Zhengyuan Yang, Xing Jin, Kefan Yu, Minh Nhat Nguyen, **Licheng Liu**, Eli Gottlieb, Yiping Lu, Kyunghyun Cho, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li.
   COLM 2025 Workshop on AI Agents (AIA). In submission to TMLR.

PUBLICATIONS

1. **[Spotlight, top 3.2% of 21,575] Online Prediction with Limited Selectivity** [Paper]
   Licheng Liu and Mingda Qiao.
   Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS) 2025.
2. **[Oral, top 0.7% of 23,918] RAGEN-2: Reasoning Collapse in Agentic RL** [Website] [Paper]
   Zihan Wang*, Chi Gui*, Xing Jin*, Qineng Wang*, **Licheng Liu***, Kangrui Wang, Shiqi Chen, Linjie Li, Zhengyuan Yang, Pingyue Zhang, Yiping Lu, Jiajun Wu, Li Fei-Fei, Lijuan Wang, Yejin Choi, and Manling Li.
   Forty-Third International Conference on Machine Learning (ICML) 2026.

EDUCATION

Imperial College, London, United Kingdom — Oct 2022 | June 2026
Mathematics, B.S. Candidate, Imperial College London
Selected Courses: Scientific computation, Time series analysis, Stochastic simulation, Network Science, Principles of programming, Introduction to computation

University of Wisconsin-Madison, Madison, United States — Aug 2026 | June 2031 (Expected)
Computer Science, Ph.D. Candidate (Incoming), University of Wisconsin-Madison
Research interests: LLM post-training

ACADEMIC EXPERIENCE

MLL Lab, Northwestern University
Research Intern - Advisor: Prof. Manling Li — Mar 2025 | Present
Project 1: A Simple "Try Again" Can Elicit Multi-Turn LLM Reasoning

- Identified a key failure mode of single-turn RLVR for reasoning: models **stop revising across turns** and instead **repeat the same answer** under feedback, which breaks interactive problem solving. Introduced an evaluation metric PassSeq@k for multi-turn revision behavior to understand the failure mode.

- Proposed **Unary Feedback as Observation (UFO)**: a multi-turn RL formulation that injects minimal unary feedback (e.g., "Please try again") as the environment observation, enabling multi-turn training **directly on static single-turn datasets** with minimal changes to existing RL setups.
- Conducted large-scale experiments across Math/STEM/QA/general reasoning, showing UFO **preserves single-turn performance** while improving **multi-turn reasoning accuracy by up to 14%**, and promotes self-reflection and adaptation to prior failures.
- Designed reward shaping aligned with real multi-turn objectives: **turn-wise reward decay** to encourage solving in fewer turns (minimality), and an **answer repetition penalty** to encourage diverse exploration when wrong (diversity).
- Released open-source code and trained models for reproducible multi-turn RL research. First-authored paper accepted at ICML AI4Math Workshop 2025, and also presented at NeurIPS LAW and NeurIPS MTI-LLM Workshops.

Project 2: Understanding Reasoning Collapse in Multi-turn Agent Reinforcement Learning

- Proposed a theoretical framework to **explain and interpret reasoning collapse** in closed-loop multi-turn RL, modeling it as the interaction of task signal, prompt-agnostic regularizers (e.g., KL/entropy), and reward noise.
- Developed an **SNR / reward-variance mechanism**: argued that low reward variance implies weak task gradient signal and, under non-vanishing reward noise, yields noise-dominated updates that accumulate as drift toward prompt-agnostic templates.
- Introduced a conceptual taxonomy distinguishing multiple forms of collapse, enabling clearer evaluation and ablations.

Project 3: RAGEN

- Contributed to the RAGEN framework to train LLM reasoning agents in interactive, stochastic environments (2400+ stars on Github).
- Assisted in the theoretical development and paper writing.

University of California, Berkeley
Research Collaborator - Advisors: Prof. Joseph Gonzalez, Prof. Ion Stoica Ongoing

- Contributing to the construction of an **Audio Tool Benchmark**, which evaluates whether models can correctly recognize user needs, invoke appropriate tools, and complete task execution after receiving spoken requests in audio-interaction scenarios.
- Beyond standard explicit-instruction tasks, designed **implicit-instruction tasks** that test whether models can infer latent user intent from speech content and trigger the correct tool calls through non-direct instructions.
- Responsible for the **infrastructure** of the benchmark, implementing the end-to-end evaluation pipeline and system; the project is ongoing.

Massachusetts Institute of Technology
Research Intern - Advisor: Postdoctoral Associate Mingda Qiao May 2024 | May 2025

- Studied **online prediction under limited selectivity**, a setting where the forecaster may only start predictions on a given subset of timesteps, generalizing selective prediction to more realistic forecasting constraints.
- Introduced the **Prediction with Limited Selectivity (PLS)** model and formulated the goal as minimizing the optimal achievable prediction error on a given PLS instance.
- Developed an **instance-dependent complexity measure** (**approximate uniformity**) and used it to derive near-optimal, instance-by-instance bounds on the optimal prediction error.
- Proved both **upper and lower bounds** in PLS, and provided an **average-case analysis** showing the bounds match **with high probability** on randomly generated PLS instances (tight up to constants).
- First-authored paper accepted to **NeurIPS 2025** as a **Spotlight**.

SERVICE

Reviewer: AI4MATH@ICML, MTI-LLM@NeurIPS, external reviewer@ICLR