Kaiyuan Liu

Seattle lky04@cs.washington.edu kaiyuanliu04.github.io Scholar Kaiyuan Liu

Education

B.S./M.S. **University of Washington, Seattle** 2022 ~2027
Combined B.S./M.S. Program in Computer Science
Double Major: Computer Science, Mathematics

- GPA: 3.93/4.0; Honor Math Sequence; Dean's List
- †Departmental Honors in Computer Science

Awards

ICPC World Final Honor, ICPC Foundation 2024.09
CRA Undergraduate Research Award Honorable Mention, CRA 2025.09
Shenoy Undergraduate Research Fellowship, Simons Foundation 2024 ~2025
CSE Award for Excellence Scholarship, University of Washington 2024 ~2025
ICPC North American Championship 12th Place, ICPC Foundation 2024.05
UW Winter Programming Contest Champion, University of Washington 2023 ~2024
Tinker Research Grant, 5000$ credit, Thinking Machines Lab 2025.11

Publications

LiveCodeBench Pro: How Do Olympiad Medalists Judge LLMs in Competitive Programming?
Zihan Zheng*, Zerui Cheng*, Zeyu Shen*, Shang Zhou*, **Kaiyuan Liu***, *et al.* *Equal contribution
The Thirty-Ninth Annual Conference on Neural Information Processing Systems (NeurIPS 2025)
arxiv.org/abs/2506.11928

AutoCode: LLMs as Problem Setters for Competitive Programming
Shang Zhou*, Zihan Zheng*, **Kaiyuan Liu***, *et al.* *Equal contribution
arxiv.org/abs/2510.12803
The Fourteenth International Conference on Learning Representations (ICLR 2026)

Evaluating LLM Agents as Human Simulators in Climate Social Dilemmas
Kaiyuan Liu*, Xiaoxuan Hou*, Jiayi Yuan, Natasha Jaques *Equal contribution
OpenReview Link
ICLR 2026 MALGAI Workshop *Oral*

FrontierCS: Evolving Challenges for Evolving Intelligence
Qiuyang Mang*, Wenhao Chai*, Zhifei Li*, Huanzhi Mao*, Shang Zhou*, *et al.* *Equal contribution
arxiv.org/abs/2512.15699
Forty-Third International Conference on Machine Learning

Harbor Adapters and Harbor-Mix: Infrastructure and a Curated Meta-Dataset for Large-Scale Agentic Evaluation
NeurIPS 2026 Under Review

Agents' Last Exam
NeurIPS 2026 Under Review

OpenDeepThink: Parallel Reasoning via Bradley-Terry Aggregation
Shang Zhou, Wenhao Chai, **Kaiyuan Liu**, Huanzhi Mao, Qiuyang Mang, Jingbo Shang
OpenReview Link
NeurIPS 2026 Under Review

FrontierSmith: Synthesizing Open-Ended Coding Problems at Scale

Runyuan He, Qiuyang Mang, Shang Zhou, **Kaiyuan Liu**, Hanchen Li, Huanzhi Mao, *et al.*

OpenReview Link

NeurIPS 2026 Under Review

Blogpost

Your Next Long-Context Recipe: Open-Ended Problems

Qiuyang Mang, **Kaiyuan Liu**, Zhifei Li, Wenhao Chai, Hanchen Li, Lin Shi, Zhixuan Zhu, Crystal Zhou, Alex Dimakis

frontier-cs.org/blog/harbor/

Frontier-CS Blog, May 12, 2026

Experience

Social RL Lab, Research Assistant — Seattle, WA, 2025.01 ~Present

- Working on various projects related to multi-agent, reinforcement learning, and large language models.
- Mentored by Dr. Natasha Jaques.

Allen Institute, Research Assistant — Seattle, WA, 2024.09 ~2025.05

- Applied Deep Learning and Reinforcement Learning to analyze representations from artificial models and neural data from Multi-Armed Bandit mouse foraging tasks. Explored brain representations and dynamics in decision making

VecML, Machine Learning Engineer — Seattle, WA, 2024.06 ~2024.09

- VecML is a startup company focusing on Machine Learning System Infrastructure.
- Tested and developed vector databases for Retrieval-Augmented Generation.
- Designed and implemented an Memory-Disk Hybrid Architecture for fast Top-K Nearest Neighborhood Search

University of Washington, Teaching Assistant — Seattle, WA, 2023 ~2024

- Assisted in undergraduate algorithm courses for both major (CSE 421) and non-major (CSE 417) students.
- Served as leading TA in CSE 421 2024 Winter.

Professional Service

Reviewer, NeurIPS Workshop 2025, NeurIPS 2026, AAAI 2026, AAMAS 2026
ICPC Coach, University of Washington ICPC 2026

Projects

Neural Translator: Predicting Neural Activity from Other Brain Regions — Comp Neuro, NLP

- Developed a neural translator model that predicts neural activity in one brain region from activity patterns in other regions using deep learning and sequence-modeling techniques.

Neural Spike Vectorization with Attention — RL, NLP, Comp Neuro

- Developed an attention-based model for vectorizing brain spike sequences.
- Compared latent representations of RL-trained RNNs with vectorized neural activity from mice performing decision-making tasks using optimal-transport metrics.

Technologies

Languages: Python, PyTorch, JAX, C++, Java, TypeScript, LaTeX

Databases: MySQL, NoSQL, PostgreSQL, SQLite, Azure

Topics: Algorithm Design, Reinforcement Learning, Machine Learning, NLP, Computational Neuroscience

Languages: English, Chinese (Mandarin)