

Xuzhong Wang

✉ xw3193@columbia.edu 📞 (804) 791-8290 🌐 GitHub 🔗 LinkedIn

Education

Columbia University, Fu Foundation School of Engineering *May 2028 (expected)*
B.S. in Computer Science

College of William & Mary *May 2026*
B.S. in Mathematics and Computer Science, Honors in Mathematics *GPA: 3.99/4.0*

Publications & Preprints

[†] *Authors listed alphabetically (mathematics convention).* * *Equal contribution.*

1. **X. Wang**, M. Jiang, T. Nair, G. Bhusal, Y. Zhang, H. Chen. “Towards Reliable, Generalizable, and Specific In-Context Knowledge Editing via Multi-Objective Reinforcement Learning.” *Findings of EMNLP 2026*; [arXiv:2608.25100](#).
2. S. Al Jannat, Y. Li, M. He, **X. Wang**, H. Zhao, J. Sun, D. Xu, E. Xu, Y. Sun. “Akita: A High Usability Simulation Framework for Computer Architecture.” *Under submission to HPCA 2027*; [arXiv:2604.28073](#).

Research Experience

Visiting Researcher, Sky Computing Lab, UC Berkeley Aug. 2026 – Present
Advisor: Joseph E. Gonzalez

- Developing **reinforcement learning for programmatic tool calling (PTC)**, where the policy emits executable code over tool APIs rather than one call per turn, yielding better-defined states and substantial token savings relative to conventional tool-use RL.
- Designing a **process reward model** that uses the stateful execution cells of interactive PTC as natural step boundaries, replacing final-answer-only verification that gives too sparse a signal for long-horizon tasks.

Undergraduate Researcher, D3i Lab, College of William & Mary May 2025 – Present
Advisors: Haipeng Chen, Yanfu Zhang

- Designed and implemented a **multi-objective RL framework** to train an auxiliary retriever policy for in-context knowledge editing, formulating demonstration selection as a constrained MDP solved with multi-objective GRPO
- Mitigated reward hacking observed in prior work via constrained reward formulation, improving editing specificity by **~10%** over existing baselines.
- Currently developing an **evaluation framework for LLM unlearning** that separates genuine knowledge erasure from input-side suppression, using invariance under content-preserving query transformations as the measurement axis.
- Built RL training infrastructure in a **verl** fork (custom advantage estimators, per-objective diagnostics) and ran multi-node experiments on the NAIRR/Anvil HPC cluster with BFCL evaluation.

Undergraduate Honors Thesis, Graph Theory March 2025 – May 2026
Advisors: Daniel W. Cranston, Gexin Yu

- Introduced the **Generalized Extension Lemma**, strengthening the Extension Lemma for bounding recoloring sequences from a single vertex to arbitrary subgraphs.
- Developed a **global discharging argument** that exploits global graph structure, extending beyond classical local discharging to obtain sharper density bounds.
- **Proved** that 5-list recoloring has linear diameter for $\text{mad}(G) < 3 - \varepsilon$, sharp relative to Cereceda's conjecture; **improved** the linear bound for 6-list recoloring from 3.4 to $3.75 - \varepsilon$.

Undergraduate Researcher, Scalable Architecture Lab

May 2024 – November 2025

Advisor: Yifan Sun

- Implemented a **data recording tool** in Go and SQLite for simulation data tracking, enabling efficient read/write support for experiments.
- Maintained Akita, a computer architecture simulator engine, improving developer experience through **event-driven design** and **smart ticking** for simulation.

Professional Experience

Machine Learning Engineer Intern, AnytimeAI

June 2026 – Aug. 2026

White Plains, NY

- Designed **LLM agent pipelines** for cross-document correlation analysis of legal documents.
- Built an agent routing/advisor tool that reduced model inference costs by **~60%**.
- Evaluated a multimodal ensemble for extracting and analyzing tabular data from plain text, improving task performance by **~8%**.

Machine Learning Engineer Intern, Bestlink Technology

May 2024 – July 2024

Nanjing, Jiangsu

- Designed a **data retrieval pipeline** with OpenCV to extract QA pairs from images.
- Partnered with senior engineers to architect a **LangChain-based LLM agent workflow** powering a question-answering tool for consulting clients.

Talks & Presentations

- Poster: “List-Recoloring of Graphs with Small Maximum Average Degree.” Joint Mathematics Meetings (JMM), Washington D.C., January 2026.

Grants & Awards

- Charles Center Summer Research Grant, College of William & Mary, 2025, 2026.
- Semester Research Grant, College of William & Mary, 2025.
- Cissy Patterson Grant for Research, College of William & Mary, 2025.
- JMM Travel Grant, Department of Mathematics, College of William & Mary, 2026.