

JINGWEN LIU

(+1) 858-319-5689
jingwenliu@cs.columbia.edu
<https://ljwo105.github.io/>

EDUCATION	Columbia University <i>Ph.D. in Computer Science</i> • Advisors: Daniel Hsu and Alexandr Andoni • Research area: Foundations of deep learning and RL post-training University of California, San Diego <i>B.S. in Mathematics and Computer Science</i> • Summa Cum Laude	September 2023 - October 2027 (Expected) New York, NY September 2019 - June 2023 La Jolla, CA
RESEARCH INTEREST	Theoretical and Empirical Science of Deep Learning and Reinforcement Learning <i>Understanding the benefits of data repetition in neural network optimization; Understanding the emergent capabilities in RL post-training.</i>	
INDUSTRY EXPERIENCE	Microsoft Research <i>Research Intern, hosted by Akshay Krishnamurthy</i> • Post-trained language models with up to 1.7B parameters using supervised fine-tuning (SFT) and RLVR (GRPO variants) on reasoning game tasks. • Designed controlled curricula and synthetic training environments for recursive self-improvement (RSI) to understand the underlying mechanisms such as emergent backtracking and feature improvement.	May 2026 - August 2026 New York, NY
SELECTED PUBLICATIONS	Science of Deep Learning 1. Jingwen Liu, Ezra Edelman, Surbhi Goel, Bingbin Liu. “Less Data, Faster Training: repeating smaller datasets speeds up learning via sampling biases.” <i>ICML</i> , 2026. Also presented at the <i>ICLR 2026 Workshop on Scientific Methods for Understanding Deep Learning</i> (contributed talk). Transformer Architecture Design 1. Jingwen Liu, Alexandr Andoni, Daniel Hsu. “Fixed Universal Transformers.” <i>NeurIPS</i> , 2026. 2. Jingwen Liu, Hantao Yu, Clayton Sanford, Alexandr Andoni, Daniel Hsu. “Fast attention mechanisms: a tale of parallelism.” <i>NeurIPS</i> , 2025.	
WORKING PROJECTS	Recursive self-improvement through RL curriculum <i>with Nived Rajaraman, Bingbin Liu, Eshaan Nichani, Dylan Foster and Akshay Krishnamurthy at Microsoft Research</i>	

- PUBLICATIONS
1. Rashida Hakim, Rachel Cummings, Jingwen Liu. “Optimal Data Acquisition with Correlated Privacy Costs.” *In submission*, 2026.
 2. Jingwen Liu, Alexandr Andoni, Daniel Hsu. “Fixed Universal Transformers.” *Advances in Neural Information Processing Systems (NeurIPS)*, 2026.
 3. Jingwen Liu, Ezra Edelman, Surbhi Goel, Bingbin Liu. “Less Data, Faster Training: repeating smaller datasets speeds up learning via sampling biases.” *Workshop on Scientific Methods for Understanding Deep Learning, ICLR 2026*, **Contributed talk**. *The International Conference on Machine Learning (ICML)*, 2026.
 4. Navid Ardeshtir*, Samuel Deng*, Daniel Hsu*, Jingwen Liu*. “Group-realizable multi-group learning by minimizing empirical risk.” *The 37th International Conference on Algorithmic Learning Theory (ALT)*, 2026.
 5. Jingwen Liu, Hantao Yu, Clayton Sanford, Alexandr Andoni, Daniel Hsu. “Fast attention mechanisms: a tale of parallelism.” *Advances in Neural Information Processing Systems (NeurIPS)*, 2025.
 6. Samuel Deng*, Daniel Hsu*, Jingwen Liu*. “Group-wise oracle-efficient algorithms for online multi-group learning.” *Advances in Neural Information Processing Systems (NeurIPS)*, 2024.
 7. Zeyuan Chen, Yinbo Chen, Jingwen Liu, Xingqian Xu, Vidit Goel, Zhangyang Wang, Humphrey Shi, Xiaolong Wang. “VideoINR: Learning Video Implicit Neural Representation for Continuous Space-Time Super-Resolution.” *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.

- ACADEMIC EXPERIENCE
- | | |
|---|-----------------------------|
| University of Pennsylvania | May 2025 - July 2025 |
| <i>Hosted by Surbhi Goel</i> | Philadelphia, PA |
| <ul style="list-style-type: none"> • Developed data repetition strategies for optimization speedup during pre-training. | |
| Simons Institute for the Theory of Computing | August 2024 - December 2024 |
| <i>UC Berkeley</i> | Berkeley, CA |
| <ul style="list-style-type: none"> • Participant in Modern Paradigms in Generalization and Special Year on Large Language Models and Transformers, Part 1. | |
| University of California, San Diego | April 2022 - August 2022 |
| <i>Hosted by Sanjoy Dasgupta</i> | La Jolla, CA |
| <ul style="list-style-type: none"> • Proved the inconsistency of the distance-based method for nearest neighbor-like algorithms with abstain rules. • Designed and performed experiments to investigate the behavior of both adaptive and nonadaptive nearest neighbor-like algorithms and calibrated parametric methods. | |

- AWARDS AND HONORS
- **Mudd Fellowship**, Columbia University 2023
Awarded for a small number of admitted PhD students at Columbia Engineering.
 - **Provost Honors**, UCSD 2019-2023