

Jingwen Gu

+1 (607) 379-0535 | jgu17@illinois.edu | [Linkedin Profile](#) | [Home Page](#)

EDUCATION

University of Illinois, Urbana-Champaign, Master of Science in Computer Science Expected Graduation: May 2028

Cornell University, Cumulative GPA: 4.057/4.3

Aug. 2021 - May 2026

Selected Honors and Awards: Dean's List (All semesters since Fall 2021); AIA Medal for Academic Excellence (2026, highest GPA in graduating class); Oppenheim Ecological Thesis Prize (2026)

Relevant Coursework: CS6789 Foundations of RL (A+); CS4756 Robot Learning (A+); CS4780 Intro to ML (A+)

PUBLICATIONS & CONFERENCE PRESENTATIONS

1. **Jingwen Gu***, Yiting He*, Zhishuai Liu*, Pan Xu. (2025). Policy Regularized Distributionally Robust Markov Decision Processes with Linear Function Approximation. *arXiv preprint*. ([arXiv:2510.14246](#))
2. Jiaru Zou*, Ling Yang*, **Jingwen Gu***(equal contribution), Jiahao Qiu, Ke Shen, Jingrui He, Mengdi Wang. (2025). ReasonFlux-PRM: Trajectory-Aware PRMs for Long Chain-of-Thought Reasoning in LLMs. *NeurIPS 2025*. ([arXiv:2506.18896](#))
3. Bradley Guo, **Jingwen Gu**, Jin Peng Zhou, Wen Sun. (2025). Learning to Self-Correct through Chain-of-Thought Verification. *ICML 2025, 2nd Workshop on Test-Time Adaptation: Putting Updates to the Test (PUT)*
4. Jin Peng Zhou, Katie Z Luo, **Jingwen Gu**, Jason Yuan, Kilian Q. Weinberger, Wen Sun. (2024). Orchestrating LLMs with Different Personalizations. *arXiv preprint* ([arXiv:2407.04181](#)).
5. **Jingwen Gu**, Timur Dogan. (2025). Virtual Horizon Method: Fast Shading Calculations for UBEH using Lidar Data Rasterization. *Oral Presentation at IBPSA Building Simulation 2025*.

RESEARCH EXPERIENCE

Zanette's Lab, Carnegie Mellon University (Advisor: Prof. Andrea Zanette)

Jun. 2026 – present

- Researching on the potential of RLVR-generated language model rollouts to be reused as synthetic data for offline learning;
- Fine-tuning 0.6B, 1.7B, and 4B models that retain online learning pass@1024 performance in AIME and reasoning gym;
- Exploring methods to counter entropy collapse in reinforcement learning-based post-training for language models.

Cornell Computing and Information Science (Advisor: Prof. Wen Sun)

Jan. 2024 – May 2025

- Worked on multi-objective reinforcement learning (MORL) for language model personalization;
- Fine-tuned and aligned 7B open-source language models with PPO that simultaneously optimize 3 objectives;
- Evaluated and compared 5 different MORL methods, identifying trade-offs in reward balancing;
- Obtained 62% average win-rate against existing MORL methods on personalization.

Environmental Systems Lab, Cornell University (Advisor: Prof. Timur Dogan)

Aug. 2023 – May 2026

- Developed LiDAR-based ray tracing methods for urban topographies.
- Built a C# application interfacing with Vulkan and OpenGL, enabling fast urban-scale shading simulations on 1.7km tiles.
- Reduced computation time by **6x** compared to traditional methods while retaining 95% accuracy.

PROFESSIONAL EXPERIENCE

Research Intern, **WEIRD Lab, University of Washington (Advisor: Prof. Abhishek Gupta)**

Jun. 2025 – Aug. 2025

- Researched on in-context reinforcement learning methods for simulated robotic construction and assembly;
- Implemented meta-RL algorithms in various simulated robotic construction environments (e.g., IsaacLab, Robosuite);
- Achieved 31% improvement in trajectory returns upon deployment to out-of-distribution environments with different simulation parameters.

AI Researcher/Engineer Intern, **Dereka AI, Seattle, WA (Supervisor: Tairan Qi)**

Jun. 2024 – Aug. 2024

- Implemented quantization, pruning, and knowledge distillation for large model compression;
- Reduced model size and runtime memory by 10x while maintaining performance and inference latency;
- Built an MVP that enhances K-12 history education with AI-generated videos.

PROFESSIONAL SKILLS

Programming Languages: Python (Proficient), Java (Proficient), C# (Proficient), C/C++ (Intermediate)

Frameworks/Tools: PyTorch, NumPy, Pandas, Transformers, OpenGL, Vulkan, LaTeX

Languages: Proficient in English, Native in Mandarin Chinese