Yang "William" Xu

Email | Web | GitHub | LinkedIn

Goal: Seeking a full-time position in machine learning research.

Education

Purdue University, West Lafayette, USA — Aug 2022–Present
Ph.D. in Industrial Engineering (Operations Research Track)

Purdue University, West Lafayette, USA — Aug 2022–May 2025
M.S. in Electrical and Computer Engineering

The Chinese University of Hong Kong, Shenzhen, China — Sept 2018–May 2022
B.E. in Electronic and Information Engineering

Experience

- **Purdue University** — *August 2022–Present*
  Graduate Research and Teaching Assistant — West Lafayette, USA
  - Conduct research in machine learning on the theory and applications of decision-making and performance evaluation in stochastic systems, with publications at NeurIPS, ICML, AAAI, and TMLR.
  - Lead recitation sessions, hold in-person and online office hours, and grade exams in Math and IE departments.
- **Amazon.com Services LLC** — *May 2026–Present*
  Applied Scientist Intern — Seattle, USA
  - Conduct research on LLM-based content generation and evaluation to improve business impact.

Projects

- **Decision-Making** — *August 2022–Present*
  - Conducted research on reinforcement learning and stochastic approximation for decision-making, including robust, constrained, and quantum average-reward problems, with provable finite-sample guarantees.
  - Developed learning-based robotic navigation algorithms for dense crowds, emphasizing safe operation and strong generalization in multi-agent environments.
- **Performance Evaluation** — *August 2022–Present*
  - Developed theory for policy evaluation in robust average-reward reinforcement learning and for Markov chain Poisson equations, targeting principled long-run performance analysis.
  - Contributed to RLLTE (470+ stars), a modular RL toolkit; and SIREN, a reliable and statistically principled evaluation method of large language models.

Selected Publications and Preprints

- **Y. Xu***, J. Zhang*, H. Sun, Z. Zhou, T. Cao, V. Aggarwal, "Towards Reliable LLM Evaluation: Correcting the Winner's Curse in Adaptive Benchmarking," *https://arxiv.org/abs/2605.05973*, 2026.
- J. Zhang, **Y. Xu**, V. Aggarwal, "Don't Freeze, Don't Crash: Extending the Safe Operating Range of Neural Navigation in Dense Crowds," *https://arxiv.org/abs/2603.06729*, 2026.
- **Y. Xu**, V. Aggarwal, "Persistent-Transient Policy Evaluation for Markov Chains via Minimal Peripheral Quotients," *https://arxiv.org/abs/2602.00474*, 2026.
- **Y. Xu**, S. Ganesh, V. Aggarwal, "Efficient Q-Learning and Actor-Critic Methods for Robust Average Reward Reinforcement Learning," *https://arxiv.org/abs/2506.07040*, 2025.
- **Y. Xu**, W. Mondal, V. Aggarwal, "Finite-Sample Analysis of Policy Evaluation for Robust Average Reward Reinforcement Learning," *NeurIPS*, 2025.
- **Y. Xu***, S. Ganesh*, W. Mondal, Q. Bai, V. Aggarwal, "Global Convergence for Average Reward Constrained MDPs with Primal-Dual Natural Actor Critic Algorithm," *NeurIPS*, 2025.
- B. Ganguly*, **Y. Xu***, V. Aggarwal, "Quantum Speedups in Regret Analysis of Infinite Horizon Average-Reward Markov Decision Processes," *ICML*, 2025.
- **Y. Xu**, V. Aggarwal, "Accelerating Quantum Reinforcement Learning with a Quantum Natural Policy Gradient Based Approach," *ICML*, 2025.
- M. Yuan, Z. Zhang, **Y. Xu**, S. Luo, B. Li, X. Jin, W. Zeng, "RLLTE: Long-Term Evolution Project of Reinforcement Learning," *AAAI*, 2025.

Awards

- NeurIPS Financial Aid Award — *2025*
- Dean's List Award, The Chinese University of Hong Kong, Shenzhen — *2020, 2021*

Skills

- **Programming Languages and Frameworks**: Python, C/C++, MATLAB, R, LaTeX, Linux, Git