Jiamin Xu

jx389@cornell.edu — +1 607-697-4379

Education

Cornell Tech — New York, NY
Ph.D. in Operations Research and Information Engineering — Sep 2024 – Present
GPA: 4.175/4.3

Fudan University — Shanghai, China
B.S. in Applied Mathematics — Sep 2019 – 2024
GPA: 3.84/4.0 (Rank: 3/154), National Scholarship

Publications

- **Integrating Causal DAGs in Deep RL: Activating Minimal Markovian States with Multi-Order Exposure** link
  ICML 2026 SPIGM Workshop, **Xu**, Maasch, Gan
- **Fast Non-Episodic Finite-Horizon RL with K-Step Lookahead Thresholding** link
  ICML 2026, **Xu**,Gan
- **From Restless to Contextual: A Thresholding Bandit Reformulation for Finite-horizon Improvement** link
  AISTATS 2026, **Xu**, Nazarov, Rastogi, Perianez, Gan

Working Papers

- **On Experimentation With Heterogeneous Subgroups: An Asymptotic Optimal δ-Weighted-PAC Design** link
  Xu, Simchi-Levi, Wang

Research & Experience

State Construction for Deep RL — Cornell Tech
Advisor: Kyra Gan — Feburary 2026 – May 2026

- Proposed methods to build **minimal Markovian state** from raw observations given DAG.
- Proposed **MOSE** which lets the neural network to automatically build state from history.
- Empirically validated the proposed algorithm outperforms stacking frames as state.
- Empirically validated with the minimal Markovian state, MOSE outperform the version that does not have this information.

Heart Failure AI Initiative — NewYork-Presbyterian
Advisor: Kyra Gan — Feb 2026 – Present

- Using LLM to build prediction models to help decision making.
- Prompting LLM to extract data.
- Building simulator to simulate patient trajectories.

K-Step Lookahead For Finite Horizon RL — Cornell Tech
Advisor: Kyra Gan — Octorber 2025 – Feburary 2026

- Proposed learning **K-step lookahead thresholding** policy rather than optimal policy online to improve the convergence rate in non-episodic finite horizon RL.
- Developed **minimax optimal online algorithms** of learning the **K-step lookahead thresholding** policy.
- Empirically validate that the propsoed policies outperform a wide range of tabular RL algorithms.

Thresholding Bandit Approach to Improve Finite-Horizon Restless Bandit — Cornell Tech
Advisor: Kyra Gan — August 2024 – Octorber 2025

- Designed **thresholding contextual bandit framework** reducing RB problem complexity while preserving optimality guarantees.
- Developed **fast converging online algorithm** with provable finite-horizon performance improvements.
- Built simulation pipelines to validate theoretical gains vs. baselines.

Contextual Bandit Experimental Design — MIT
Advisors: David Simchi-Levi, Chonghuan Wang — May 2023 – Jan 2024

- Designed algorithms balancing **regret vs. estimation error** under heterogeneous treatment effects.
- Derived **minimax lower bounds** and constructed matching algorithms.
- Developed δ-weighted PAC framework for context-dependent decision accuracy.

Technical Skills

Programming: Python, C
Areas: Reinforcement Learning, Bandits

Leadership & Activities

- Invited Speaker, Amazon Scout RL Lunch Talk 2026/06/12
- Invited Speaker, INFORMS Healthcare Conference 2026
- Invited Speaker, INFORMS Annual Meeting 2025
- Co-organizer, INFORMS Annual Meeting Session 2025