Yinglun Xu

+1(404)200-3909
yinglun6@illinois.edu
www.xuyinglun.com

Education

2021 - **University of Illinois Urbana-Champaign**, Illinois, United States
Ph.D. in Computer Science, Advisor: Prof. Singh Gagandeep, expected May 2025

2019 - 2021 **Georgia Institute of Technology**, Georgia, United States
M.S. in Electrical and Computer Engineering, Advisor: Prof. Jacob Abernethy

2015 - 2019 **Peking University**, Beijing, China
B.S. in Physics, Advisor: Prof. Yun-Feng Xiao

Research Interests

My research interests lie in **Machine Learning and Reinforcement Learning**, including

- Reinforcement learning from human feedback (**RLHF**), offline preference-based reinforcement learning
- Trustworthy Reinforcement Learning: Adversarial Attack, Provably Efficient Exploration, Provably Robust Exploration, and Verification on Deep Reinforcement Learning (**DRL**)
- Multi-arm Bandit (**MAB**) Learning Theories

Internship

Summer 2023 **Amazon**, CA, United States
Applied Scientist Intern on Search Experience Science team

Summer 2022 **Amazon**, WA, United States
Applied Scientist Intern on Core Machine Learning Science team

Research Experience

June 2021 - Present **University of Illinois Urbana-Champaign**, Illinois, United State
Graduate Research Assistant, Advisor: Prof. Singh Gagandeep

- Study efficient data poisoning attack against deep reinforcement learning algorithms in black box setting [In submission]
- Study provably efficient deep reinforcement learning and its robust variants
- Study offline reinforcement learning that uses transformers for function approximation
- Study offline preference-based reinforcement learning and design an efficient learning algorithm for the setting [In submission]. The next step is to extend to the setting where the preference feedback provided by humans, which is also known as reinforcement learning from human feedback (RLHF)

Dec. 2019 - June 2021 **Machine Learning Theory Group, Georgia Institute of Technology**, Georgia, United State
Graduate Research Assistant, Advisor: Prof. Jacob Abernethy

- Design a truthful and robust bandit mechanism for Pay-Per-Click advertising auction [In submission]
- Study adversarial attack against randomized bandit algorithm and discover a fundamental reason why some bandit algorithms are not robust [NeurIPS 2021]

Oct. 2018 - Dec. 2018 **Nonlinear Photonics Laboratory, California Institute of Technology**, California, United State
Undergraduate Research Assistant, Advisor: Prof. Alireza Marandi

- Design an on-chip circuit to simulate an Ising model which could solve NP-hard problems. [US Patent 2020]

Oct. 2016 - June 2019 **Microcavity Photonics Group, Peking University**, Beijing, China
Undergraduate Research Assistant, Advisor: Prof. Yun-Feng Xiao

- Develop theories for efficiently characterizing nano-particles through their signals collected by an on-chip micro-circuit [PRA 2018]

Publications

* indicates equal contribution. [Google Scholar Profile]

arXiv **Reward Poisoning Attack Against Offline Reinforcement Learning**
Yinglun Xu*, Rohan Gumaste*, Gagandeep Singh

arXiv **Efficient Two-Phase Offline Deep Reinforcement Learning from Preference Feedback**
Yinglun Xu, Gagandeep Singh

arXiv **Black-Box Targeted Reward Poisoning Attack Against Online Deep Reinforcement Learning**
Yinglun Xu, Gagandeep Singh

arXiv **On the robustness of epsilon greedy in multi-agent contextual bandit mechanism**
Yinglun Xu, Bhuvesh Kumar, Jacob Abernethy

TMLR 2023 **Efficient Reward Poisoning Attacks on Online Deep Reinforcement Learning** (Featured Certification)
Yinglun Xu, Qi Zeng, Gagandeep Singh

PNAS 2022 **Single-molecule optofluidic microsensor with interface whispering gallery modes**
Xiao-Chong Yu, Shui-Jing Tang, Wenjing Liu, Yinglun Xu, Qihuang Gong, You-Ling Chen, Yun-Feng Xiao

US Patent **Thin-film optical parametric oscillators**
Alireza Marandi, Luis Ledezma, Yinglun Xu, Ryan Briggs

NeurIPS 2021 **Observation-Free Attacks on Stochastic Bandits**
Yinglun Xu, Bhuvesh Kumar, Jacob Abernethy.

M.S. Thesis **Adversarial Attack and Robust Learning in Multi-Arm Bandit Problems**
Yinglun Xu

ICML 2020 **Bridging Truthfulness and Corruption-Robustness in Multi-Armed Bandit Mechanisms** (Incentives in Machine Learning Workshop)
Jacob Abernethy, Bhuvesh Kumar, Thodoris Lykouris, Yinglun Xu (Alphabetically ordered)

PRA 2018 **Mode splitting induced by an arbitrarily shaped Rayleigh scatterer in a whispering-gallery microcavity**
Yinglun Xu, Shui-jing Tang, Xiaochong Yu, Yi-Lin Chen, Daquan Yang, Qihuang Gong, Yun-Feng Xiao.

Skills

Programming Languages: Python, C++
Mathematics: Ordinary and partial differential equations, Probability Theory