

MIAO LU

Tel: 650-250-9790 | E-mail: miaolu@stanford.edu | Web: miaolu3.github.io

Google Scholar | LinkedIn | Last update: Nov.2024

EDUCATION

Stanford University

Ph.D. student in Operations Research, Department of Management Science & Engineering
Advisor: Jose Blanchet

Stanford, USA
Sep.2023 - present

University of Science and Technology of China

B.S. in Mathematics & Applied Mathematics, with *summa cum laude*

Hefei, China
Sep.2018 - Jun.2022

INDUSTRIAL EXPERIENCES

Ubiquant Investment

Quantitative research intern

Shanghai, China
Jun.2022 - Sep.2022

RESEARCH INTERESTS

My ultimate goal of research is to: (i) develop the *foundations* of the next generation artificial intelligence (AI) techniques and (ii) make the *best use of* AI to effectively address challenging managerial and societal problems. Towards such a goal, my current research interests concentrate on the following perspectives and directions.

- **Theoretical foundations of AI.** I work on the theory of reinforcement learning, deep learning, and large foundation models through the lens of statistics and optimization. The goal is to push forward the understanding and to improve the *training* and *inference mechanism* of AI models in a principled manner. I also hold the goal of understanding the robustness and safety of AI systems through the tool of distributional robustness.
- **Large language models capabilities and alignment.** I work to push the boundary of planning/reasoning capabilities while enhancing the alignment of LLMs through the way of reinforcement learning (RL).

PUBLICATIONS

- [10] **Can Neural Networks Achieve Optimal Computational-statistical Tradeoff? An Analysis on Single-Index Model**
Siyu Chen*, Beining Wu*, **Miao Lu**, Zhuoran Yang, Tianhao Wang
NeurIPS Workshop on Mathematics of Modern Machine Learning (M3L) 2024 **Oral**
- [9] **Provably Mitigating Overoptimization in RLHF: Your SFT Loss is Implicitly an Adversarial Regularizer**
Zhihan Liu*, **Miao Lu***, Shenao Zhang, Boyi Liu, Hongyi Guo, Yingxiang Yang, Jose Blanchet, Zhaoran Wang
Neural Information Processing Systems (NeurIPS) 2024
ICML Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ARLET) 2024
- [8] **Distributionally Robust Reinforcement Learning with Interactive Data Collection: Fundamental Hardness and Near-Optimal Algorithm**
Miao Lu*, Han Zhong*, Tong Zhang, Jose Blanchet
Neural Information Processing Systems (NeurIPS) 2024
ICML Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ARLET) 2024
- [7] **Benign Oscillation of Stochastic Gradient Descent with Large Learning Rates**
Miao Lu*, Beining Wu*, Xiaodong Yang, Difan Zou
International Conference on Learning Representations (ICLR) 2024
NeurIPS Workshop on Mathematics of Modern Machine Learning (M3L) 2023
- [6] **Double Pessimism is Provably Efficient for Distributionally Robust Offline Reinforcement Learning: Generic Algorithm and Robust Partial Coverage**
Jose Blanchet[†], **Miao Lu**[†], Tong Zhang[†], Han Zhong[†]
Neural Information Processing Systems (NeurIPS) 2023
Extended version under review at *Mathematics of Operations Research (MOR)*
- [5] **Maximize to Explore: One Objective Function Fusing Estimation, Planning, and Exploration**
Zhihan Liu*, **Miao Lu***, Wei Xiong*, Han Zhong, Hao Hu, Shenao Zhang, Sirui Zheng, Zhuoran Yang, Zhaoran Wang
Neural Information Processing Systems (NeurIPS) 2023 **Spotlight**

- [4] **Pessimism in the Face of Confounders: Provably Efficient Offline Reinforcement Learning in Partially Observable Markov Decision Processes**
Miao Lu, Yifei Min, Zhaoran Wang, Zhuoran Yang
International Conference on Learning Representations (ICLR) 2023
- [3] **Welfare Maximization in Competitive Equilibrium: Reinforcement Learning for Markov Exchange Economy**
 Zhihan Liu*, **Miao Lu***, Zhaoran Wang, Michael I. Jordan, Zhuoran Yang
International Conference on Machine Learning (ICML) 2022
- [2] **Learning Pruning-Friendly Networks via Frank-Wolfe: One-Shot, Any-Sparsity, and No Retraining**
Miao Lu*, Xiaolong Luo*, Tianlong Chen, Wuyang Chen, Dong Liu, Zhangyang Wang
*International Conference on Learning Representations (ICLR) 2022 **Spotlight***
- [1] **Learning Robust Policy against Disturbance in Transition Dynamics via State-Conservative Policy Optimization**
 Yufei Kuang, **Miao Lu**, Jie Wang, Qi Zhou, Bin Li, Houqiang Li
Association for Advancement of Artificial Intelligence (AAAI) 2022

(Note: authors with * contributed equally to the work, and [†] represents alphabetical order.)

RESEARCH VISITING EXPERIENCES

Toyota Technological Institute at Chicago Student visitor hosted by Tianhao Wang and Zhiyuan Li	<i>Chicago, USA</i> <i>July.2024 - Aug.2024</i>
The University of Hong Kong Research assistant hosted by Difan Zou, Dept. of Computer Science & Institute of Data Science	<i>Hong Kong, China</i> <i>Apr.2023 - Aug.2023</i>

AWARDS AND HONORS

Xinhe Scholarship (outstanding undergraduate researchers, School of the Gifted Young, USTC)	<i>Mar.2023</i>
Yuanqing Yang Scholarship (top scholarship, School of Mathematical Sciences, USTC)	<i>Jan.2022</i>
The 41st Guo Moruo Scholarship (highest honor, USTC)	<i>Dec.2021</i>
Chinese National Scholarship (top scholarship, Ministry of Education of China)	<i>Nov.2019, 2020</i>

INVITED TALKS

Computational-statistical Trade-off of Learning Single-index Models via Neural Networks [10]	
◦ 2nd Mathematics of Modern Machine Learning Workshop (M3L), Vancouver, BC, Canada	<i>Dec.2024</i>
Theoretical Foundations of Distributionally Robust Reinforcement Learning [6, 8]	
◦ 2024 INFORMS annual meeting, Seattle, WA, USA [8]	<i>Oct.2024</i>
◦ 58th Annual Conference on Information Sciences and Systems (CISS), Princeton, NJ, USA [6]	<i>Mar.2024</i>
◦ 2023 INFORMS annual meeting, Phoenix, AZ, USA [6]	<i>Oct.2023</i>

TEACHING ASSISTANT

Differential Equations (2020 fall, USTC) (PI: Wuqing Ning, Dept. of Applied Math., USTC)	<i>Sep.2020 - Jan.2021</i>
---	----------------------------

ACADEMIC SERVICES

Journal Reviewer

Annals of Applied Probability (AOAP)
 Operations Research (OR)

Conference Reviewer

Neural Information Processing Systems (NeurIPS; 2023, 2024)
 International Conference on Machine Learning (ICML; 2024)
 International Conference on Learning Representations (ICLR; 2024, 2025)
 International Conference on Artificial Intelligence and Statistics (AISTATS; 2025)
 ICML Workshop on Aligning Reinforcement Learning Experimentalists and Theorists (ARLET; 2024)
 NeurIPS Workshop on Mathematics of Modern Machine Learning (M3L; 2024)
 Association for the Advancement of Artificial Intelligence (AAAI; 2025)