

Mengxiao Zhang

Email: mengxiao.zhang@usc.edu; Website: <https://mengxiao.github.io/>

Research Interest

Designing practical and adaptive machine learning algorithms with strong theoretical guarantees, with a focus on general sequential learning problems, including online learning, bandit problems, reinforcement learning and game theory.

Education

Graduate study:

Department of Computer Science, University of Southern California.

Advisor: Haipeng Luo

Aug 2018 - Present

Fifth-year Ph.D. student

Undergraduate study:

Key Laboratory of Machine Perception, MOE, School of EECS, Peking University

Sep 2014 - July 2018

Advisor: Liwei Wang

Apr 2016 - Aug 2018

Bachelor of Science in Machine Intelligence, Peking University. Double major in Mathematics and Applied Mathematics

Conference Papers

(* indicates alphabetical ordering or equal contribution)

Improved High-Probability Regret for Adversarial Bandits with Time-Varying Feedback Graphs

Haipeng Luo*, Hanghang Tong*, Mengxiao Zhang*, Yuheng Zhang*, ALT 2023

No-Regret Learning in Two-Echelon Supply Chain with Unknown Demand Distribution

Mengxiao Zhang, Shi Chen, Haipeng Luo, Yingfei Wang, AISTATS 2023

Corralling a Larger Band of Bandits: A Case Study on Switching Regret for Linear Bandits

Haipeng Luo*, Mengxiao Zhang*, Peng Zhao*, Zhi-Hua Zhou*, COLT 2022

Adaptive Bandit Convex Optimization with Heterogeneous Curvature

Haipeng Luo*, Mengxiao Zhang*, Peng Zhao*, COLT 2022

No-Regret Learning in Time-Varying Zero-Sum Games

Mengxiao Zhang*, Peng Zhao*, Haipeng Luo, Zhi-Hua Zhou, ICML 2022

Last-iterate Convergence of Decentralized Optimistic Gradient Descent/Ascent in Infinite-horizon Competitive Markov Games

Chen-Yu Wei, Chung-Wei Lee*, Mengxiao Zhang*, Haipeng Luo, COLT 2021

Achieving Near Instance-Optimality and Minimax-Optimality in Stochastic and Adversarial Linear Bandits Simultaneously

Chung-Wei Lee*, Haipeng Luo*, Chen-Yu Wei*, Mengxiao Zhang*, Xiaojin Zhang*, ICML 2021

Linear Last-iterate Convergence in Constrained Saddle-point Optimization

Chen-Yu Wei, Chung-Wei Lee, Mengxiao Zhang, Haipeng Luo, ICLR 2021

Bias no more: high-probability data-dependent regret bounds for adversarial bandits and MDPs

Chung-Wei Lee*, Haipeng Luo*, Chen-Yu Wei*, Mengxiao Zhang*, NeurIPS 2020, Oral

A Closer Look at Small-loss Bounds for Bandits with Graph Feedback

Chung-Wei Lee*, Haipeng Luo*, Mengxiao Zhang*, COLT 2020

Workshop Papers

Advancing Query Rewriting in E-Commerce via Shopping Intent Learning

Mengxiao Zhang, Yongning Wu, Raif Rustamov, Hongyu Zhu, Haoran Shi, Yuqi Wu, Lei Tang, Zuohua Zhang, Chu Wang, SIGIR eCom 2022, Spotlight Talk

Preprint

Autobidders with Budget and ROI Constraints: Efficiency, Regret, and Pacing Dynamics

Brendan Lucier*, Sarath Pattathil*, Aleksandrs Slivkins*, Mengxiao Zhang*, arxiv 2301.13306.

Defective Convolutional Layers Learn Robust CNNs

Tiange Luo, Tianle Cai*, Mengxiao Zhang, Siyu Chen, Di He, Liwei Wang, arxiv 1911.08432.*

Randomness in Deconvolutional Networks for Visual Representation

Kun He, Jingbo Wang, Haochuan Li, Yao Shu, Mengxiao Zhang, Man Zhu, Liwei Wang, John E. Hopcroft, arxiv 1704.00330.

Graduate Courses:

2018 Fall: CSCI 670 Advanced Analysis of Algorithm. Instructor: David Matthias Kempe, GPA:4.0
2018 Fall: CSCI 671 Randomized Algorithm. Instructor: David Matthias Kempe, GPA:4.0
2019 Spring: DSO 699 Statistic Maching Learning. Instructor: Jason D.Lee, GPA: 4.0
2019 Fall: CSCI 675 Convex and Combinatorial Optimization. Instructor: Shaddin Dughmi, GPA: 4.0
2020 Spring: CSCI 672 Approximation Algorithm. Instructor: Shaddin Dughmi, GPA: 4.0
2020 Fall: CSCI 556 Introduction to Cryptography. Instructor: Ming-deh Huang, GPA: 4.0

Mentorship

Game AI design for Rock Paper Scissors. *Jun. 2019 - Aug. 2019*

Dept.of Computer Science, University of Southern California

SHINE program: Mentoring Justin Jang from West Ranch High school on a project aiming at designing an Rock Paper Scissors AI by using online learning algorithms such as AdaNormalHedge.

Internship

Carnegie Mellon University *June 2017 - Sept 2017*

Mentor: Maria Florina Balcan and Ellen Vitercik

Project: Comparison Based Choices in Economics

Amazon.com *May 2020 - Aug 2020*

Applied Science Intern at Sponsor Brand Response Model Team

Mentor: Chu Wang

Project: Contextual Query Rewriting and Beyond

Amazon.com *May 2021 - Aug 2021*

Applied Science Intern at Sponsor Brand Response Model Team

Mentor: Chu Wang

Project: Improved Seq2Seq Query Rewriting in E-commerce via Shopper Intent Learning

Microsoft Research *May 2022 - Aug 2022*

Research Intern

Mentor: Alex Slivkins and Brendan Lucier

Project: Multi-agent auction under resource and return-on-investment constraints

Teaching Assistant

CSCI 570: Analysis of Algorithms, Instructor: Victor Adamchik *2020 Spring*

CSCI 567: Machine Learning, Instructor: Haipeng Luo *2020 Fall*

CSCI 567: Machine Learning, Instructor: Haipeng Luo *2021 Fall*

CSCI 670: Advanced Analysis of Algorithms, Instructor: Ming-Deh Huang *2022 Spring*

CSCI 670: Advanced Analysis of Algorithms, Instructor: Ming-Deh Huang *2023 Spring*

Awards:

Recipient of National Scholarship in 2016 (Top 1 in department);

Recipient of National Scholarship in 2017 (Top 1 in department);

Recipient of Peking University Foundation Scholarship in 2015;

Recipient of Peking University Merit Student in 2015;

Recipient of Peking University Merit Student in 2016;

Recipient of Peking University Pacemaker of Merit Student in 2017 (Top 1 in department);

Recipient of Peking University Outstanding Undergraduates, Beijing Outstanding Undergraduates in 2018;

Service:

Reviewer:

- Advances in Neural Information Processing Systems (NeurIPS): 2022, 2021, 2020

- International Conference on Machine Learning (ICML): 2023, 2022, 2021, 2020
- Conference on Learning Theory (COLT): 2023
- Annual ACM Symposium on Theory of Computing (STOC): 2023
- International Joint Conference on Artificial Intelligence (IJCAI): 2023, 2022, 2021
- ACM International Conference on Knowledge Discovery and Data Mining (KDD): 2021
- International Conference on Artificial Intelligence and Statistics (AISTATS): 2023
- Conference on Algorithmic Learning Theory (ALT): 2021

Talk:

- Corraling a Larger Band of Bandits: A Case Study on Switching Regret for Linear Bandits, Conference on Learning Theory (COLT), London (online), July, 2022
- Adaptive Bandit Convex Optimization with Heterogeneous Curvature, Conference on Learning Theory (COLT), London (online), July, 2022
- No-Regret Learning in Time-Varying Zero-Sum Games, AI-TIME invited talk (online), August, 2022
- Achieving Near Instance-Optimality and Minimax-Optimality in Stochastic and Adversarial Linear Bandits Simultaneously, International Conference on Machine Learning (ICML), virtual, July, 2021
- Linear Last-iterate Convergence in Constrained Saddle-point Optimization, RPI invited talk (online), April, 2021
- Bias-no-more: High-probability Data-dependent Regret Bounds for Adversarial Bandits and MDPs, Advances in Neural Information Processing Systems (NeurIPS), virtual, December, 2020
- A Closer Look at Small-loss Bounds for Bandits with Graph Feedback, Conference on Learning Theory (COLT), virtual, July, 2020