

Jiawei Zhang

Department of Computer Sciences
University of Wisconsin–Madison

Email: jzhang2924@wisc.edu
<https://jwzhang.netlify.app>

Academic Appointments

University of Wisconsin–Madison , Madison, WI <i>Assistant Professor</i>	2025 – present
Massachusetts Institute of Technology , Cambridge, MA <i>MIT Postdoctoral Fellow for Engineering Excellence</i> Advisors: Prof. Asuman Ozdaglar and Prof. Saurabh Amin	2023 – 2025
<i>Postdoctoral Associate</i> Advisor: Prof. Asuman Ozdaglar	2021 – 2023

Education

The Chinese University of Hong Kong, Shenzhen , Shenzhen, China <i>Ph.D. in Computer and Information Engineering</i> Advisor: Prof. Zhi-Quan (Tom) Luo	August 2021
University of Science and Technology of China , Hefei, China <i>B.Sc. in Mathematics (Hua Loo-Keng Talent Program)</i>	June 2016

Research Interests

- Optimization theory and algorithms
- Efficient, reliable, and generalizable training and inference for LLMs and generative models
- Optimization and AI for science and engineering

Honors and Awards

MIT Postdoctoral Fellowship for Engineering Excellence, Massachusetts Institute of Technology, 2023

Presidential Award for Outstanding Doctoral Students, The Chinese University of Hong Kong, Shenzhen, 2021

Outstanding Teaching Assistant Award, The Chinese University of Hong Kong, Shenzhen, 2020

Professional Service

Action Editor, Transactions on Machine Learning Research (TMLR).

Area Chair, International Conference on Machine Learning (ICML 2026).

Area Chair, Conference on Neural Information Processing Systems (NeurIPS).

Area Chair, International Conference on Learning Representations (ICLR).

Academic Publications

Journal Articles

“On the Convergence Analysis of Muon,” Wei Shen, Ruichuan Huang, Minhui Huang, Cong Shen, and **Jiawei Zhang**. *Transactions on Machine Learning Research (TMLR)*, 2026. **Selected for presentation at NeurIPS 2026** (Journal-to-Conference Track).

“Beamforming Optimization for Robust Sensing and Communication in Dynamic mmWave MIMO Networks,” Lei Li, **Jiawei Zhang**, and Tsung-Hui Chang. *IEEE Journal on Selected Areas in Communications*, 43(4):1354–1370, 2025.

“Decentralized Non-Convex Learning With Linearly Coupled Constraints: Algorithm Designs and Application to Vertical Learning Problem,” **Jiawei Zhang**, Songyang Ge, Tsung-Hui Chang, and Zhi-Quan Luo. *IEEE Transactions on Signal Processing*, 2022.

“A Global Dual Error Bound and Its Application to the Analysis of Linearly Constrained Nonconvex Optimization,” **Jiawei Zhang** and Zhi-Quan Luo. *SIAM Journal on Optimization*, 2022.

“Distributed Stochastic Consensus Optimization With Momentum for Nonconvex Nonsmooth Problems,” Zhiguo Wang, **Jiawei Zhang**, Tsung-Hui Chang, Jian Li, and Zhi-Quan Luo. *IEEE Transactions on Signal Processing*, 2021.

“A Proximal Alternating Direction Method of Multiplier for Linearly Constrained Nonconvex Minimization,” **Jiawei Zhang** and Zhi-Quan Luo. *SIAM Journal on Optimization*, 2020.

Refereed Conference Proceedings

“Free Draft-and-Verification: Toward Lossless Parallel Decoding for Diffusion Large Language Models,” Shutong Wu and **Jiawei Zhang**. *NeurIPS 2026*, accepted.

“SAGE: Mitigating Long-Horizon Reasoning Biases via Topological Guidance,” Xinyue Zeng, **Jiawei Zhang**, Yujun Yan, and Dawei Zhou. *NeurIPS 2026*, accepted.

“From Likelihood Convergence to Parameter Convergence in POMDPs,” Jack Zhang, Saurabh Amin, **Jiawei Zhang**, and Patrick Jaillet. *NeurIPS 2026*, accepted.

“Understanding Why Language Models Hallucinate: Testing Reasoning Against Priors,” Yangfan Hu, Xuhan Tong, Haoyue Bai, Xi Ding, Shashank Muralidhar Bharadwaj, Siyang Cao, Robert Nowak, and **Jiawei Zhang**. *Findings of EMNLP 2026*, accepted.

“Smoothing the Score Function to Enhance Generalization in Diffusion Models,” Xinyu Zhou, **Jiawei Zhang** (co-advising), and Stephen J. Wright. *CVPR 2026*.

“MLorc: Momentum Low-rank Compression for Memory Efficient Large Language Model Adaptation,” Wei Shen, Zhang Yaxiang, Minhui Huang, Mengfan Xu, **Jiawei Zhang**, and Cong Shen. *AISTATS 2026*.

“Beyond RLHF and NLHF: Population-Proportional Alignment under an Axiomatic Framework,” Kihyun Kim, **Jiawei Zhang** (co-first author; corresponding author), Asuman E. Ozdaglar, and Pablo A. Parrilo. *ICLR 2026*.

“Reducing Contextual Stochastic Bilevel Optimization via Structured Function Approximation,” Maxime Bouscary, **Jiawei Zhang**, and Saurabh Amin. *ICLR 2026*.

“A Single-Loop First-Order Algorithm for Linearly Constrained Bilevel Optimization,” Wei Shen, **Jiawei Zhang**, Minhui Huang, and Cong Shen. *NeurIPS 2025*.

“Contextual Optimization Under Model Misspecification: A Tractable and Generalizable Approach,” Omar Bennouna, **Jiawei Zhang** (co-first author; corresponding author), Saurabh Amin, and Asuman E. Ozdaglar. *ICML 2025*.

“Stochastic Smoothed Primal-Dual Algorithms for Nonconvex Optimization with Linear Inequality Constraints,” Ruichuan Huang, **Jiawei Zhang** (co-last author; corresponding author), and Ahmet Alacaoglu. *ICML 2025*.

“Uniformly Stable Algorithms for Adversarial Training and Beyond,” Jiancong Xiao, **Jiawei Zhang** (co-first author; corresponding author), Zhi-Quan Luo, and Asuman Ozdaglar. *ICML 2024*.

“A Unified Linear Programming Framework for Offline Reward Learning from Human Demonstrations and Feedback,” Kihyun Kim, **Jiawei Zhang** (co-first author), Asuman E. Ozdaglar, and Pablo Parrilo. *ICML 2024*.

“Stochastic Smoothed Gradient Descent Ascent for Federated Minimax Optimization,” Wei Shen, Minhui Huang, **Jiawei Zhang**, and Cong Shen. *AISTATS 2024*.

“A Smoothed Bregman Proximal Gradient Algorithm for Decentralized Nonconvex Optimization,” Wenqiang Pu, **Jiawei Zhang**, Rui Zhou, Xiao Fu, and Mingyi Hong. *ICASSP 2024*.

“Linearly Constrained Bilevel Optimization: A Smoothed Implicit Gradient Approach,” Prashant Khanduri, Ioannis Tsaknakis, Yihua Zhang, Jia Liu, Sijia Liu, **Jiawei Zhang**, and Mingyi Hong. *ICML 2023*.

“Revisiting the Linear-Programming Framework for Offline RL with General Function Approximation,” Asuman Ozdaglar, Sarath Pattathil, **Jiawei Zhang** (*corresponding author*), and Kaiqing Zhang. *ICML 2023*. (*Alphabetical author order*.)

“What is a Good Metric to Study Generalization of Minimax Learners?” Asuman Ozdaglar, Sarath Pattathil, **Jiawei Zhang** (*corresponding author*), and Kaiqing Zhang. *NeurIPS 2022*. (*Alphabetical author order*.)

“When Expressivity Meets Trainability: Fewer than n Neurons Can Work,” **Jiawei Zhang**, Yushun Zhang, Mingyi Hong, Ruoyu Sun, and Zhi-Quan Luo. *NeurIPS 2021*.

“Communication Efficient Primal-Dual Algorithm for Nonconvex Nonsmooth Distributed Optimization,” Congliang Chen, **Jiawei Zhang**, Li Shen, Peilin Zhao, and Zhi-Quan Luo. *AISTATS 2021*.

“A Single-Loop Smoothed Gradient Descent-Ascent Algorithm for Nonconvex-Concave Min-Max Problems,” **Jiawei Zhang**, Peijun Xiao, Ruoyu Sun, and Zhi-Quan Luo. *NeurIPS 2020*.

“A Proximal Dual Consensus Method for Linearly Coupled Multi-Agent Non-Convex Optimization,” **Jiawei Zhang**, Songyang Ge, Tsung-Hui Chang, and Zhi-Quan Luo. *ICASSP 2020*.

“Scalable Gaussian Process Using Inexact ADMM for Big Data,” Yue Xu, Feng Yin, **Jiawei Zhang**, Wenjun Xu, Shuguang Cui, and Zhi-Quan Luo. *ICASSP 2019*.

“Learning Word Vectors with Linear Constraints: A Matrix Factorization Approach,” Wenye Li, **Jiawei Zhang**, Jianjun Zhou, and Laizhong Cui. *IJCAI 2018*.

Workshop Papers

“The Power of Duality Principle in Offline Average-Reward Reinforcement Learning,” Asuman Ozdaglar, Sarath Pattathil, **Jiawei Zhang** (*corresponding author*), and Kaiqing Zhang. *ICML 2023 Workshop on Duality for Modern Machine Learning*. (*Alphabetical author order*.)

“Smoothed-SGDmax: A Stability-Inspired Algorithm to Improve Adversarial Generalization,” Jiancong Xiao, **Jiawei Zhang** (*co-first author*), Zhi-Quan Luo, and Asuman Ozdaglar. *NeurIPS 2022 ML Safety Workshop*.

Preprints

“Interactive-Policy Distillation with Bidirectional Propose-and-Verify,” Shutong Wu, Xiwen Chen, Brendan Rappazzo, Daiheng Zhang, Anderson Schneider, Yuriy Nevmyvaka, and **Jiawei Zhang**. *arXiv preprint arXiv:2609.36546*, 2026.

“MemFold: Learning Compact Soft Memory for Long-Context Personalization via On-Policy Optimization,” Jingxuan Wu, Yuzhe Yang, Yiqiao Huang, Chengzhi Liu, Qingni Wang, Chengxuan Qian, Shutong Wu, **Jiawei Zhang**, and Xin Eric Wang. *arXiv preprint arXiv:2609.36435*, 2026.

“Obtaining Game-Stationary Points for Smooth Nonconvex-Nonconcave Minimax Problems via First-Order Methods,” Mengxuan Dong, Yao Yao, and **Jiawei Zhang**. *arXiv preprint arXiv:2609.32970*, 2026.

“When the Wrong Key Wins: Understanding and Detecting Hallucinations in LLMs,” Xuhan Tong, Haoyue Bai, Dawei Zhou, Naichen Shi, and **Jiawei Zhang**. *arXiv preprint arXiv:2609.15106*, 2026.

“Denoising as Projection: Constrained Optimization with Gradient-Guided Diffusion,” Runyu Zhang, **Jiawei Zhang** (*corresponding author*), Gioele Zardini, Saurabh Amin, and Asuman Ozdaglar. *arXiv preprint arXiv:2608.29507*, 2026.

“Off-Manifold Collapse in Guided Protein Language Models,” Shuibai Zhang, Xinchu Liu, Fred Zhangzhi Peng, Zhihan Yang, Shutong Wu, Yingzi Ma, and **Jiawei Zhang**. *arXiv preprint arXiv:2608.18597*, 2026.

“MatrAIx: Simulating the World with 8.3 Billion Persona Agents,” Xiaomin Li, Yuexing Hao, Jianheng Hou, ..., **Jiawei Zhang**, ..., Dawn Song. *arXiv preprint arXiv:2608.04205*, 2026. (Author list abbreviated.)

“Inverse Reinforcement Learning without an Optimal Demonstrator: A Feasible Reward Set Approach,” Kihyun Kim, Shripad Deshmukh, Nikos Vlassis, and **Jiawei Zhang**. *arXiv preprint arXiv:2605.30903*, 2026.

“Demonstrations, CoT, and Prompting: A Theoretical Analysis of ICL,” Xuhan Tong, Yuchen Zeng, and **Jiawei Zhang**. *arXiv preprint arXiv:2603.19611*, 2026.

“Why GRPO Needs Normalization: A Local-Curvature Perspective on Adaptive Gradients,” Cheng Ge, Caitlyn Heqi Yin, Hao Liang, and **Jiawei Zhang**. *arXiv preprint arXiv:2601.23135*, 2026.

“Smoothed Proximal Lagrangian Method for Nonlinear Constrained Programs,” Wenqiang Pu, Kaizhao Sun, and **Jiawei Zhang**. *arXiv preprint arXiv:2408.15047*, 2024.

“On the Iteration Complexity of Smoothed Proximal ALM for Nonconvex Optimization Problem with Convex Constraints,” **Jiawei Zhang**, Wenqiang Pu, and Zhi-Quan Luo. *arXiv preprint arXiv:2207.06304*, 2022.