

# Zhuoran Yang | Curriculum Vitae

Version: Nov 2024 – Department of Statistics and Data Science, Yale University  
✉ zhuoran.yang@yale.edu; zhuoranyang.work@gmail.com • 🌐 <https://zhuoranyang.github.io/>

## Research Interest

- Reinforcement Learning
- Nonconvex Optimization
- Artificial Intelligence
- High-Dimensional Statistics

My research interest broadly lies in the interface between machine learning, statistics, and optimization. In particular, I am interested in understanding the theoretical underpinnings of foundation models and designing a new generation of artificial intelligence algorithms for large-scale and multi-agent decision-making problems. My current research mainly focuses on the following two aspects:

- **Theoretical foundations of large language models and their usage:** Unveiling the inner mechanisms of large language models, including their training and prompting processes., and leveraging these insights for downstream applications, especially decision-making problems.
- **Multi-agent learning with strategic behaviors:** Designing multi-agent learning algorithms that can be deployed in the presence of strategic agents, particularly in contexts where agents possess their own interests and are required to interact with humans.

## Professional Experience

### Yale University

Assistant Professor of Statistics and Data Science and Computer Science (Secondary)

New Haven, CT, USA.

July. 2022 –

### UC Berkeley

Postdoctoral Researcher

Berkeley, CA, USA.

July. 2021 – June. 2022

### Simons Institute

VMware Research Fellow

Berkeley, CA, USA.

Aug. 2020 – Dec. 2020

### Tencent AI Lab

Machine Learning Research Intern

Bellevue, WA, USA.

Jun. 2018 – Sept. 2018

### Tencent AI Lab

Machine Learning Research Intern

Shenzhen, China.

Jun. 2017 – Sept. 2017

## Education

### Princeton University

Doctor of Philosophy

Department of Operations Research and Financial Engineering

Princeton, NJ, USA

Sept. 2015 – May. 2021

### Tsinghua University

Bachelor of Mathematics

Specialized in Probability and Statistics

Beijing, China

Sept. 2011 – Jun. 2015

### École Normale Supérieure

Exchange Student

Paris, France

Jan. 2015 – May. 2015

### The University of Wisconsin-Madison

Exchange Student

Madison, WI, USA

Jan. 2014 – May. 2014

## Awards

2020: VMware Research Fellow, Simons Institute

2018: Tencent Ph.D Fellowship

2015: Francis Robbins Upton Fellowship

2012: Tsinghua Xuetao Fellowship

## Publications and Preprints

---

### Journal Publications and Submissions

- [J1] Luofeng Liao, Zuyue Fu, Zhuoran Yang, Yixin Wang, Dingli Ma, Mladen Kolar, Zhaoran Wang. **Instrumental Variable Value Iteration for Causal Offline Reinforcement Learning**. Journal of Machine Learning Research, 2024.
- [J2] Zihao Li, Boyi Liu, Zhaoran Wang, Zhuoran Yang, Mengdi Wang. **Double Duality: Variational Primal-Dual Policy Optimization for Constrained Reinforcement Learning**. Journal of Machine Learning Research, 2023.
- [J3] Hong, Mingyi, Hoi-To Wai, Zhaoran Wang, and Zhuoran Yang. **A two-timescale stochastic algorithm framework for bilevel optimization: Complexity analysis and application to actor-critic**. SIAM Journal on Optimization, 2023.
- [J4] Qi Cai, Zhuoran Yang, Jason D. Lee, Zhaoran Wang. **Neural Temporal Difference and Q Learning Provably Converge to Global Optima**. Mathematics of Operations Research, 2023.
- [J5] Jianqing Fan, Zhuoran Yang, Mengxin Yu. **Understanding Implicit Regularization in Over-Parameterized Nonlinear Statistical Model**. Journal of the American Statistical Association, 2023.
- [J6] Han Zhong, Zhuoran Yang, Zhaoran Wang, Michael I. Jordan. **Can Reinforcement Learning Find Stackelberg-Nash Equilibria in General-Sum Markov Games with Myopically Rational Followers?** Journal of Machine Learning Research, 2023.
- [J7] Chi Jin, Zhuoran Yang, Zhaoran Wang, Michael I. Jordan. **Provably Efficient Reinforcement Learning with Linear Function Approximation**. Mathematics of Operations Research, 2023.
- [J8] Pratik Ramprasad, Yuantong Li, Zhuoran Yang, Zhaoran Wang, Will Wei Sun, and Guang Cheng. **Online Bootstrap Inference for Policy Evaluation in Reinforcement Learning**. Journal of the American Statistical Association, 2023.
- [J9] Xie, Qiaomin, Yudong Chen, Zhaoran Wang, and Zhuoran Yang. **Learning Zero-Sum Simultaneous-Move Markov Games Using Function Approximation and Correlated Equilibrium**. Mathematics of Operations Research, 2023.
- [J10] Zhuoran Yang, Lin F. Yang, Ethan X. Fang, Tuo Zhao, Zhaoran Wang, and Matey Neykov. **Misspecified Nonconvex Statistical Optimization for Phase Retrieval**. Mathematical Programming, 2019.
- [J11] Zhuoran Yang, Yang Ning, Han Liu. **On Semiparametric Exponential Family Graphical Models**. Journal of Machine Learning Research, 2018.
- [J12] Sen Na, Zhuoran Yang, Zhaoran Wang, Mladen Kolar. **High-dimensional Varying Index Coefficient Models via Stein's Identity**. Journal of Machine Learning Research, 2019.
- [J13] Jianqing Fan, Han Liu, Zhaoran Wang, Zhuoran Yang. **Curse of Heterogeneity: Computational Barriers in Sparse Mixture Models and Phase Retrieval**. (2018) Submitted to Annals of Statistics, under Revision.
- [J14] Jianqing Fan, Zhaoran Wang, Yuchen Xie, Zhuoran Yang. **A Theoretical Analysis of Deep Q-Learning**. (2020) Submitted to Annals of Statistics, under Revision.
- [J15] Krishna Balasubramanian, Jianqing Fan, Zhuoran Yang. **Tensor Methods for Additive Index Models under Discordance and Heterogeneity**, (2018) Submitted to Annals of Statistics, under Revision.

### Conference Publications

- [C1] Siyu Chen, Heejune Sheen, Tianhao Wang, Zhuoran Yang. **Unveiling Induction Heads: Provable Training Dynamics and Feature Learning in Transformers**. Conference on Neural Information Processing Systems (NeurIPS), 2024.
- [C2] Awni Altabaa, Zhuoran Yang. **On the Role of Information Structure in Reinforcement Learning for Partially-Observable Sequential Teams and Games**. Conference on Neural Information Processing Systems (NeurIPS), 2024.
- [C3] Siyu Chen, Heejune Sheen, Tianhao Wang, and Zhuoran Yang. **Training Dynamics of Multi-Head Softmax Attention for In-Context Learning: Emergence, Convergence, and Optimality**. Conference on Learning Theory, (2024).
- [C4] Nuoya Xiong, Zhihan Liu, Zhaoran Wang, Zhuoran Yang. **Sample-Efficient Multi-Agent RL: An Optimization Perspective**. International Conference on Learning Representations (ICLR), 2024.

- [C5] Jianliang He, Han Zhong, Zhuoran Yang. **Sample-efficient Learning of Infinite-horizon Average-reward MDPs with General Function Approximation.** International Conference on Learning Representations (ICLR), 2024.
- [C6] Juno Kim, Kakei Yamamoto, Kazusato Oko, Zhuoran Yang, Taiji Suzuki. **Symmetric Mean-field Langevin Dynamics for Distributional Minimax Problems.** International Conference on Learning Representations (ICLR), 2024.
- [C7] Jianliang He, Siyu Chen, Fengzhuo Zhang, Zhuoran Yang. **From Words to Actions: Unveiling the Theoretical Underpinnings of LLM-Driven Autonomous Systems.** International Conference on Machine Learning (ICML), 2024.
- [C8] Zehao Dou, Minshuo Chen, Mengdi Wang, Zhuoran Yang. **Theory of Consistency Diffusion Models: Distribution Estimation Meets Fast Sampling.** International Conference on Machine Learning (ICML), 2024.
- [C9] Han Shen, Zhuoran Yang, Tianyi Chen. **Principled Penalty-based Methods for Bilevel Reinforcement Learning and RLHF.** International Conference on Machine Learning (ICML), 2024.
- [C10] Nuoya Xiong, Zhaoran Wang, Zhuoran Yang. **A General Framework for Sequential Decision-Making under Adaptivity Constraints.** International Conference on Machine Learning (ICML), 2024.
- [C11] Sirui Zheng, Chenjia Bai, Zhuoran Yang, Zhaoran Wang. **How Does Goal Relabeling Improve Sample Efficiency?** International Conference on Machine Learning (ICML), 2024.
- [C12] Kakei Yamamoto, Kazusato Oko, Zhuoran Yang, Taiji Suzuki. **Mean Field Langevin Actor-Critic: Faster Convergence and Global Optimality beyond Lazy Learning.** International Conference on Machine Learning (ICML), 2024.
- [C13] Zihan Zhu, Ethan X. Fang, Zhuoran Yang. **Online Performative Gradient Descent for Learning Nash Equilibria in Decision-Dependent Games.** Conference on Neural Information Processing Systems (NeurIPS), 2023.
- [C14] Shuang Qiu, Ziyu Dai, Han Zhong, Zhaoran Wang, Zhuoran Yang, Tong Zhang. **Posterior Sampling for Competitive RL: Function Approximation and Partial Observation.** Conference on Neural Information Processing Systems (NeurIPS), 2023.
- [C15] Zhihan Liu, Miao Lu, Wei Xiong, Han Zhong, Hao Hu, Shenao Zhang, Sirui Zheng, Zhuoran Yang, and Zhaoran Wang. **Maximize to Explore: One Objective Function Fusing Estimation, Planning, and Exploration.** Conference on Neural Information Processing Systems (NeurIPS), 2023.
- [C16] Fengzhuo Zhang, Vincent Tan, Zhaoran Wang, Zhuoran Yang. **Learning Regularized Monotone Graphon Mean-Field Games.** Conference on Neural Information Processing Systems (NeurIPS), 2023.
- [C17] Haoran Xu, Li Jiang, Jianxiong Li, Zhuoran Yang, Zhaoran Wang, Victor Wai Kin Chan, Xianyuan Zhan. **Offline RL with No OOD Actions: In-Sample Learning via Implicit Value Regularization.** International Conference on Learning Representations (ICLR), 2023.
- [C18] Miao Lu, Yifei Min, Zhaoran Wang, Zhuoran Yang. **Pessimism in the Face of Confounders: Provably Efficient Offline Reinforcement Learning in Partially Observable Markov Decision Processes.** International Conference on Learning Representations (ICLR), 2023.
- [C19] Zhuoqing Song, Jason D. Lee, Zhuoran Yang. **Can We Find Nash Equilibria at a Linear Rate in Markov Games?** International Conference on Learning Representations (ICLR), 2023.
- [C20] Wenhao Zhan, Jason D. Lee, Zhuoran Yang. **Decentralized Optimistic Hyperpolicy Mirror Descent: Provably No-Regret Learning in Markov Games.** International Conference on Learning Representations (ICLR), 2023.
- [C21] Siyu Chen, Jibang Wu, Yifan Wu, Zhuoran Yang. **Learning to Incentivize Information Acquisition: Proper Scoring Rules Meet Principal-Agent Model.** International Conference on Machine Learning (ICML), 2023.
- [C22] Yulai Zhao, Zhuoran Yang, Zhaoran Wang, Jason Lee. **Local Optimization Achieves Global Optimality in Multi-Agent Reinforcement Learning.** International Conference on Machine Learning (ICML), 2023.

- [C23] Jiacheng Guo, Zihao Li, Huazheng Wang, Mengdi Wang, Zhuoran Yang, Xuezhou Zhang. **Provably Efficient Representation Learning with Tractable Planning in Low-Rank POMDP**. International Conference on Machine Learning (ICML), 2023.
- [C24] Yixuan Wang, Simon Zhan, Ruochen Jiao, Zhilu Wang, Wanxin Jin, Zhuoran Yang, Zhaoran Wang, Chao Huang, Qi Zhu. **Enforcing Hard Constraints with Soft Barriers: Safe Reinforcement Learning in Unknown Stochastic Environments**. International Conference on Machine Learning (ICML), 2023.
- [C25] Yifei Min, Tianhao Wang, Ruitu Xu, Zhaoran Wang, Michael Jordan, Zhuoran Yang. **Learn to Match with No Regret: Reinforcement Learning in Markov Matching Markets**. Conference on Neural Information Processing Systems (NeurIPS), 2022.
- [C26] Gene Li, Junbo Li, Anmol Kabra, Nathan Srebro, Zhaoran Wang, Zhuoran Yang. **Exponential Family Model-Based Reinforcement Learning via Score Matching**. Conference on Neural Information Processing Systems (NeurIPS), 2022.
- [C27] Grigoris Velegkas, Zhuoran Yang, Amin Karbasi. **Reinforcement Learning with Logarithmic Regret and Policy Switches**. Conference on Neural Information Processing Systems (NeurIPS), 2022.
- [C28] Fengzhuo Zhang, Boyi Liu, Kaixin Wang, Vincent Tan, Zhuoran Yang, Zhaoran Wang. **Relational Reasoning via Set Transformers: Provable Efficiency and Applications to MARL**. Conference on Neural Information Processing Systems (NeurIPS), 2022.
- [C29] Baihe Huang, Jason D. Lee, Zhaoran Wang, Zhuoran Yang. **Towards General Function Approximation in Zero-Sum Markov Games**. International Conference on Learning Representations (ICLR), 2022.
- [C30] Chenjia Bai, Lingxiao Wang, Zhuoran Yang, Zhi-Hong Deng, Animesh Garg, Peng Liu, Zhaoran Wang. **Pessimistic Bootstrapping for Uncertainty-Driven Offline Reinforcement Learning**. International Conference on Learning Representations (ICLR), 2022.
- [C31] Qi Cai, Zhuoran Yang, Zhaoran Wang. **Reinforcement Learning from Partial Observation: Linear Function Approximation with Provable Sample Efficiency**. International Conference on Machine Learning (ICML), 2022.
- [C32] Zhihan Liu, Yufeng Zhang, Zuyue Fu, Zhuoran Yang, Zhaoran Wang. **Learning from Demonstration: Provably Efficient Adversarial Policy Imitation with Linear Function Approximation**. International Conference on Machine Learning (ICML), 2022.
- [C33] Siyu Chen, Donglin Yang, Jiayang Li, Senmiao Wang, Zhuoran Yang, Zhaoran Wang. **Adaptive Model Design for Markov Decision Process**. International Conference on Machine Learning (ICML), 2022.
- [C34] Han Zhong, Wei Xiong, Jiyuan Tan, Liwei Wang, Tong Zhang, Zhaoran Wang, Zhuoran Yang. **Pessimistic Minimax Value Iteration: Provably Efficient Equilibrium Learning from Offline Datasets**. International Conference on Machine Learning (ICML), 2022.
- [C35] Boxiang Lyu, Zhaoran Wang, Mladen Kolar, Zhuoran Yang. **Pessimism meets VCG: Learning Dynamic Mechanism Design via Offline Reinforcement Learning**. International Conference on Machine Learning (ICML), 2022.
- [C36] Zhihan Liu, Lu Miao, Zhaoran Wang, Michael Jordan, Zhuoran Yang. **Welfare Maximization in Competitive Equilibrium: Reinforcement Learning for Markov Exchange Economy**. International Conference on Machine Learning (ICML), 2022.
- [C37] Xiaoyu Chen, Han Zhong, Zhuoran Yang, Zhaoran Wang, Liwei Wang. **Human-in-the-loop: Provably Efficient Preference-based Reinforcement Learning with General Function Approximation**. International Conference on Machine Learning (ICML), 2022.
- [C38] Yufeng Zhang, Siyu Chen, Zhuoran Yang, Michael Jordan, Zhaoran Wang. **Wasserstein Flow Meets Replicator Dynamics: A Mean-Field Analysis of Representation Learning in Actor-Critic**. Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [C39] Minshuo Chen, Yan Li, Ethan Wang, Zhuoran Yang, Zhaoran Wang, Tuo. Zhao **Pessimism Meets Invariance: Provably Efficient Offline Mean-Field Multi-Agent RL**. Conference on Neural Information Processing Systems (NeurIPS), 2021.

- [C40] Lingxiao Wang, Zhuoran Yang, Zhaoran Wang. **Provably Efficient Causal Reinforcement Learning with Confounded Observational Data.** Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [C41] Yingjie Fei, Zhuoran Yang, Yudong Chen, Zhaoran Wang. **Exponential Bellman Equation and Improved Regret Bounds for Risk-Sensitive Reinforcement Learning.** Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [C42] Runzhe Wu, Yufeng Zhang, Zhuoran Yang, Zhaoran Wang. **Offline Constrained Multi-Objective Reinforcement Learning via Pessimistic Dual Value Iteration.** Conference on Neural Information Processing Systems (NeurIPS), 2021.
- [C43] Jin, Ying, Zhuoran Yang, and Zhaoran Wang. **Is Pessimism Provably Efficient for Offline RL?** International Conference on Machine Learning (ICML), 2021.
- [C44] Shuang Qiu, Xiaohan Wei, Jieping Ye, Zhaoran Wang, Zhuoran Yang. **Provably Efficient Fictitious Play Policy Optimization for Zero-Sum Markov Games with Structured Transitions.** International Conference on Machine Learning (ICML), 2021.
- [C45] Tengyu Xu, Zhuoran Yang, Zhaoran Wang, Yingbin Liang. **Doubly Robust Off-Policy Actor-Critic: Convergence and Optimality.** International Conference on Machine Learning (ICML), 2021.
- [C46] Shuang Qiu, Jieping Ye, Zhaoran Wang, Zhuoran Yang. **On Reward-Free RL with Kernel and Neural Function Approximations: Single-Agent MDP and Markov Game.** International Conference on Machine Learning (ICML), 2021.
- [C47] Yingjie Fei, Zhuoran Yang, Zhaoran Wang. **Risk-Sensitive Reinforcement Learning with Function Approximation: A Debiasing Approach.** International Conference on Machine Learning (ICML), 2021.
- [C48] Zhuoran Yang, Chi Jin, Zhaoran Wang, Mengdi Wang, Michael I. Jordan. **On Function Approximation in Reinforcement Learning: Optimism in the Face of Large State Spaces** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C49] Yingjie Fei, Zhuoran Yang, Yudong Chen, Zhaoran Wang, Qiaomin Xie. **Risk-Sensitive Reinforcement Learning: Near-Optimal Risk-Sample Tradeoff in Regret,** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C50] Yingjie Fei, Zhuoran Yang, Zhaoran Wang, Qiaomin Xie. **Dynamic Regret of Policy Optimization in Non-stationary Environments,** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C51] Shuang Qiu, Zhuoran Yang, Xiaohan Wei, Jieping Ye, Zhaoran Wang. **Upper Confidence Primal-Dual Reinforcement Learning for CMDP with Adversarial Loss.** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C52] Luofeng Liao, You-Lin Chen, Zhuoran Yang, Bo Dai, Mladen Kolar, and Zhaoran Wang **Provably Efficient Neural Estimation of Structural Equation Models: An Adversarial Approach.** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C53] Wai, Hoi-To, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong. **Provably Efficient Neural GTD for Off-Policy Learning,** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C54] Yufeng Zhang, Qi Cai, Zhuoran Yang, Yongxin Chen, Zhaoran Wang. **Can Temporal-Difference and Q-Learning Learn Representation? A Mean-Field Theory.** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C55] Wanxin Jin, Zhaoran Wang, Zhuoran Yang, Shaoshuai Mou. **Pontryagin Differentiable Programming: An End-to-End Learning and Control Framework.** Conference on Neural Information Processing Systems (NeurIPS), 2020.
- [C56] Shuang Qiu, Xiaohan Wei, Zhuoran Yang. **Robust One-Bit Recovery via ReLU Generative Networks: Improved Statistical Rates and Global Landscape Analysis.** International Conference on Machine Learning (ICML), 2020.
- [C57] Qi Cai, Zhuoran Yang, Chi Jin, Zhaoran Wang. **Provably Efficient Exploration in Policy Optimization.** International Conference on Machine Learning (ICML), 2020.

- [C58] Lingxiao Wang, Qi Cai, Zhuoran Yang, Zhaoran Wang. **On the Global Optimality of Model-Agnostic Meta-Learning**. International Conference on Machine Learning (ICML), 2020.
- [C59] Lingxiao Wang, Zhuoran Yang, Zhaoran Wang. **Breaking the Curse of Many Agents: Provable Mean Embedding Q-Iteration for Mean-Field Reinforcement Learning**. International Conference on Machine Learning (ICML), 2020.
- [C60] Sen Na, Yuwei Luo, Zhuoran Yang, Zhaoran Wang, and Mladen Kolar. **Semiparametric Nonlinear Bipartite Graph Representation Learning with Provable Guarantees**. International Conference on Machine Learning (ICML), 2020.
- [C61] Yufeng Zhang, Qi Cai, Zhuoran Yang, and Zhaoran Wang. **Generative Adversarial Imitation Learning with Neural Network Parameterization: Global Optimality and Convergence Rate**. International Conference on Machine Learning (ICML), 2020.
- [C62] Zuyue Fu, Zhuoran Yang, Yongxin Chen, Zhuoran Yang, **Actor-Critic Provably Finds Nash Equilibria of Linear-Quadratic Mean-Field Games**, International Conference on Learning Representations (ICLR), 2020.
- [C63] Lingxiao Wang, Qi Cai, Zhuoran Yang, Zhaoran Wang. **Neural Policy Gradient Methods: Global Optimality and Rates of Convergence**. International Conference on Learning Representations (ICLR), 2020.
- [C64] Chen, Minshuo, Yizhou Wang, Tianyi Liu, Zhuoran Yang, Xingguo Li, Zhaoran Wang, and Tuo Zhao. **On Computation and Generalization of Generative Adversarial Imitation Learning**. International Conference on Learning Representations (ICLR), 2020.
- [C65] Zhuoran Yang, Yongxin Chen, Mingyi Hong, Zhaoran Wang. **Provably Global Convergence of Actor-Critic: A Case for Linear Quadratic Regulator with Ergodic Cost**, Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [C66] Boyi Liu, Qi Cai, Zhuoran Yang, Zhaoran Wang. **Neural Proximal Policy Optimization Attains Optimal Policy**, Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [C67] Kaiqing Zhang, Zhuoran Yang, Tamer Başar. **Policy Optimization Provably Converges to Nash Equilibria in Zero-Sum Linear Quadratic Games**, Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [C68] Lingxiao Wang, Zhuoran Yang, Zhaoran Wang. **Statistical-Computational Tradeoff in Single Index Models**, Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [C69] Wai, Hoi-To, Zhuoran Yang, Zhaoran Wang, and Mingyi Hong, Kexin Tang. **Variance Reduced Policy Evaluation with Smooth Function Approximation**. (with H. Wai, M. Hong, Z. Wang, and K. Tang). (2019). Conference on Neural Information Processing Systems (NeurIPS).
- [C70] Ming Yin, Zhuoran Yang, Zhaoran Wang, Mladen Kolar. **Convergent Policy Optimization for Safe Reinforcement Learning**. Conference on Neural Information Processing Systems (NeurIPS), 2019.
- [C71] Xiaohan Wei, Zhuoran Yang, Zhaoran Wang. **On the Statistical Rate of Nonlinear Recovery in Generative Models with Heavy-Tailed Data**, International Conference on Machine Learning (ICML), 2019.
- [C72] **Multi-Agent Reinforcement Learning via Double Averaging Primal-Dual Optimization**, (with H. Wai, Z. Wang, and M. Hong). Conference on Neural Information Processing Systems (NeurIPS), 2018
- [C73] Xinyang Yi, Zhaoran Wang, Zhuoran Wang, Constantine Caramanis, Han Liu. **More Supervision, Less Computation: Statistical-Computational Tradeoffs in Weakly Supervised Learning**, Conference on Neural Information Processing Systems (NeurIPS), 2016.
- [C74] Zhuoran Yang, Zhaoran Wang, Han Liu, Yonina Eldar, Tong Zhang. **Sparse Nonlinear Regression: Parameter Estimation and Asymptotic Inference**, (with Z. Wang, H. Liu, Y. Eldar, and T. Zhang). International Conference on Machine Learning (ICML), 2016.

## Book Chapters.....

- [B1] Kaiqing Zhang, Zhuoran Yang, and Tamer Başar. **Multi-Agent Reinforcement Learning: A Selective Overview of Theories and Algorithms**. Handbook of Reinforcement Learning and Control, 2021.

## Manuscripts and Preprints

---

- [P1] Xinyang Hu, Fengzhuo Zhang, Siyu Chen, and Zhuoran Yang. **Unveiling the statistical foundations of chain-of-thought prompting methods**. arXiv:2408.14511 (2024).
- [P2] Siyu Chen, Mengdi Wang, and Zhuoran Yang. **Actions Speak What You Want: Provably Sample-Efficient Reinforcement Learning of the Quantal Stackelberg Equilibrium from Strategic Feedbacks**. arXiv:2307.14085 (2023).
- [P3] Siyu Chen, Yitan Wang, Zhaoran Wang, and Zhuoran Yang. **A Unified Framework of Policy Learning for Contextual Bandit with Confounding Bias and Missing Observations**. arXiv:2303.11187 (2023).
- [P4] Yufeng Zhang, Fengzhuo Zhang, Zhuoran Yang, and Zhaoran Wang. **What and How does In-Context Learning Learn? Bayesian Model Averaging, Parameterization, and Generalization**. arXiv:2305.19420 (2023).
- [P5] Zihao Li, Zhuoran Yang, and Mengdi Wang. **Reinforcement Learning with Human Feedback: Learning Dynamic Choices via Pessimism**. arXiv:2305.18438 (2023).

## Teaching Experiences

---

**S&DS 431/631: Optimization and Computation:** Fall 2022, Fall 2023, Fall 2024

**S&DS 685: Theory of Reinforcement Learning:** Spring 2023

**S&DS 432/632: Advanced Optimization Techniques:** Spring 2024

## Funding Support

---

**NSF DMS :** # 2413243

**Title:** A Statistical Foundation of In-Context Learning and Chain-of-Thought Prompting with Large Language Models

## Invited Talks and Seminars

---

11/2024 University of North Carolina - Chapel Hill, Statistics Colloquium  
11/2024 MIT Data Science Lab Seminar  
09/2024 UCL Computational Neuroscience Unit  
08/2024 Bernoulli-IMS World Congress in Probability and Statistics  
08/2024 Fidelity AI Seminar  
07/2024 EC 2024 Workshop, "INFORMS Workshop on Market Design"  
05/2024 NYU Math and Data (MaD) Seminar  
04/2024 Harvard Economics & Computer Science Research Seminar  
02/2024 AAAI 2024 Tutorial, "TH8: Learning with Multiple Objectives Beyond Bilevel Optimization"  
12/2023 ISAIM 2024 Special Session on Deep Reinforcement Learning  
09/2023 Duke Fuqua, Decision Science Seminar  
04/2023 Georgia Tech ISYE Statistics Seminar  
10/2022 University of Michigan, "Modern Statistical and Machine Learning Methods for Big Data" Workshop  
08/2022 Joint Statistical Meetings  
07/2022 International Conference on Continuous Optimization  
05/2022 Simons Institute, "Multi-Agent Reinforcement Learning and Bandit Learning" Workshop  
11/2021 University of Virginia, Computer Science Seminar  
11/2021 Simons Institute, Theory of Reinforcement Learning Reunion  
09/2021 Rutgers, Statistics Seminar  
04/2021 University of Maryland, ECE Seminar  
03/2021 USC, Computer Science Seminar  
03/2021 Harvard, Statistics Seminar  
03/2021 UPenn, ESE Seminar  
03/2021 MIT, EECS Seminar  
03/2021 University of Chicago, Statistics and Computer Science Seminar  
03/2021 NYU Stern, Statistics Seminar  
02/2021 Johns Hopkins, Mathematical Institute for Data Science (MINDS)  
02/2021 USC Marshall, Statistics Seminar  
02/2021 Brown Data Science Initiative, Statistics Seminar  
02/2021 Purdue, IE Seminar  
01/2021 Yale, Statistics Seminar  
10/2020 UIUC, IDS2 Seminar Series  
10/2020 (Online) Reinforcement Learning Seminar Series  
09/2020 University of Wisconsin-Madison, SILO Seminar Series

## Professional Services

---

**2024–present:** Area Chair of ICML, NeurIPS, ICLR, AAAI

**2023–present:** Associate Editor of ACM/IMS Journal of Data Science

**2015–present:** Reviewers of various journals and conferences

- Journals: Annals of Statistics (AOS), Journal of the American Statistical Association (JASA), Journal of Machine Learning Research (JMLR), Machine Learning Journal (MLJ), Operations Research (OR), IEEE Transactions on Automatic Control (TAC), IEEE Transactions on Information Theory (TIT), IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI), Proceedings of IEEE
- International Conference on Artificial Intelligence and Statistics (AISTATS), IEEE Conference on Decision and Control (CDC), Annual Conference on Learning Theory (COLT), International Conference on Learning Representations (ICLR), International Conference on Machine Learning (ICML), Conference on Neural Information Processing Systems (NeurIPS)

**2020 – 2021:** Co-organizing *Online Seminar on Mathematical Foundations of Data Science*, which is a weekly virtual seminar series on topics related to statistics, optimization, and machine learning. [link]

**2022–2024:** Co-organizing *Statistics and Data Science Seminar at Yale University*, weekly