Banghua Zhu

264 Cory Hall ◇ University of California, Berkeley◇ California 94709, USA
banghua@berkeley.edu ◇ http://people.eecs.berkeley.edu/~banghua/

EDUCATION

University of California, Berkeley, California, USA *2018 - 2024 (expected)*
Ph.D. candidate in Department of EECS, UC Berkeley, advised by Prof. Jiantao Jiao and Prof. Michael I. Jordan.
Research Interest: Theoretical Statistics, Information Theory, Foundation Models, Human-AI Interaction, Bandit and Reinforcement Learning, Economics and Game Theory
GPA: 4.0/4.0

Tsinghua University, Beijing, China *2014 - 2018*
B.E. in Department of Electronic Engineering
B.S. in Department of Mathematics (double major)
GPA: 95.2/100 Ranking 1^{st}/262 in Department of Electronic Engineering

The University of Texas at Austin, TX, USA *2016.8 - 2016.12*
Exchange student at the Department of Electrical and Computer Engineering
GPA: 4.0/4.0

PUBLICATIONS AND MANUSCRIPTS

Published

[1] **Banghua Zhu**, Jiantao Jiao, Michael I. Jordan. Principled Reinforcement Learning with Human Feedback from Pairwise or K-wise Comparisons. *ICML 2023.*

[2] Geng Zhao*, **Banghua Zhu***, Jiantao Jiao, Michael I. Jordan. Online Learning in Stackelberg Games with an Omniscient Follower. *ICML 2023.*

[3] Ikechukwu Uchendu, Ted Xiao, Yao Lu, **Banghua Zhu**, Mengyuan Yan, Josphine Simon, Matthew Bennice, Chuyuan Fu, Cong Ma, Jiantao Jiao, Sergey Levine, Karol Hausman. Jump-Start Reinforcement Learning. *ICML 2023.*

[4] **Banghua Zhu**, Ziao Wang, Nadim Ghaddar, Jiantao Jiao, Lele Wang. On the Optimal Bounds for Noisy Computing. *ISIT 2023.*

[5] Ziao Wang, Nadim Ghaddar, **Banghua Zhu**, Lele Wang. Noisy Sorting Capacity. *In submission.*

[6] **Banghua Zhu**, Stephen Bates, Zhuoran Yang, Yixin Wang, Jiantao Jiao, Michael I. Jordan. The Sample Complexity of Online Contract Design. *EC 2023.*

[7] **Banghua Zhu***, Lun Wang*, Qi Pang*, Shuai Wang, Jiantao Jiao, Dawn Song, Michael I. Jordan. Byzantine-Robust Federated Learning with Optimal Rates and Privacy Guarantee. *AISTATS 2023.*

[8] **Banghua Zhu**, Jiantao Jiao, Michael I. Jordan. Robust Estimation for Nonparametric Families via Generative Adversarial Networks. *ISIT 2022.*

[9] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. Generalized resilience and robust statistics. *Annals of Statistics.*

[10] Paria Rashidinejad, **Banghua Zhu**, Cong Ma, Jiantao Jiao, Stuart Russell. Bridging Offline Reinforcement Learning and Imitation Learning: A Tale of Pessimism. *NeurIPS 2021.*

[11] Cong Ma, **Banghua Zhu**, Jiantao Jiao, Martin J Wainwright. Minimax Off-Policy Evaluation for Multi-Armed Bandits. *IEEE Transactions on Information Theory.*

[12] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. Robust estimation via generalized quasi-gradients *Information and Inference: A Journal of the IMA.*

[13] **Banghua Zhu**, Jiantao Jiao, David Tse. Deconstructing Generative Adversarial Networks. *Published on IEEE Transactions on Information Theory*

[14] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. When does the Tukey median work? *ISIT 2020.*

[15] **Banghua Zhu**, Longzhuang He, Jintao Wang. Joint optimization of physical layer using neural network. *IEEE Journal on Selected Areas in Communications.*

[16] Abolfazl Hashemi, **Banghua Zhu**, Haris Vikalo. Sparse Factorization Decomposition for Haplotype Assembly of Diploids and Polyploids. *BMC Genomics*

[17] **Banghua Zhu**, Longzhuang He, Jintao Wang, Jian Song. Fast Sparse Bayesian Learning based Symbol Detection for Massive Spatial Modulation. *IEEE Broadcast Symposium 2016.*

[18] Abolfazl Hashemi, **Banghua Zhu**, Haris Vikalo. Tensor Factorization Framework for Haplotype Assembly of Diploids and Polyploids. *RECOMB-seq 2016*

Manuscripts

[1] **Banghua Zhu**, Hiteshi Sharma, Felipe Vieira Frujeri, Shi Dong, Chenguang Zhu, Michael I. Jordan, Jiantao Jiao. Fine-Tuning Language Models with Advantage-Induced Policy Alignment. *In submission.*

[2] **Banghua Zhu**, Ying Sheng, Lianmin Zheng, Clark Barrett, Michael I. Jordan, Jiantao Jiao. On Optimal Caching and Model Switching for Large Model Inference. *In submission.*

[3] **Banghua Zhu**, Mingyu Ding, Philip Jacobson, Ming Wu, Wei Zhan, Michael I. Jordan, Jiantao Jiao. Doubly Robust Self-Training. *In submission.*

[4] **Banghua Zhu**, Sai Praneeth Karimireddy, Jiantao Jiao, Michael I. Jordan. Online Learning in a Creator Economy. *In submission.*

WORKING EXPERIENCE

Research Intern at Microsoft Research, Mountain View *2022.1 - 2022.11*
Mentor: **Hiteshi Sharma, Shi Dong and Felipe Vieira Frujeri**

- Developed sample-efficient and stable algorithm for policy optimization in reinforcement learning with human feedback.
- Developed RL-based algorithms for hallucination reduction.

Student Researcher at Google Robotics, Mountain View *2022.1 - 2022.11*
Mentor: **Yao Lu**

- Developed reinforcement learning algorithms for online and offline training in real-world robotics.
- Developed principled theory for off-policy learning and off-policy evaluation.

Research Intern at Electrical Engineering Department, Stanford University *2017.7 - 2017.10*
Advisor: **Prof. David Tse**

- Cleaned, processed and analyzed Illumina, Pacbio and Nanopore reads for genome sequencing.
- Designed pipeline to assemble tandem repeats from hybrid reads and successfully applied that to Mucin 2 gene assembly.
- Developing integrated tools for analyzing and assembling tandem repeats.

AWARDS

2023	David J. Sakrison Memorial Prize
2019-2021	Berkeley EECS Department Award
2018	Beijing Excellent Undergrad Award
2017	Tsinghua Science and Innovation Scholarship
2017	Bao Gang Excellent Student Scholarship
2016	Qualcomm Scholarship
2016	National Scholarship
2015	National Scholarship
2015	Singapore Technologies Engineering China Scholarship