

Banghua Zhu

banghua@uw.edu \diamond <http://banghua.me/>

CURRENT EMPLOYMENT

Assistant Professor at University of Washington, ECE, Seattle, WA 2025 -
Research Area: Large Language Models, Human-Centered Machine Learning, Reinforcement Learning,
Game Theory, Statistics, Information Theory

Cofounder at Nexusflow AI, Palo Alto, CA 2023 - Now
Lead the Post-training of Large Language Models for Enterprise AI Agents

EDUCATION

University of California, Berkeley, California, USA 2018 - 2024
Ph.D. in Department of EECS, advised by Prof. Jiantao Jiao and Prof. Michael I. Jordan.
Thesis: Towards Efficient Fine-tuning and Serving of Large Language Models
GPA: 4.0/4.0

Tsinghua University, Beijing, China 2014 - 2018
B.E. in Department of Electronic Engineering
B.S. in Department of Mathematics
GPA: 95.2/100 Ranking 1st/262 in Department of Electronic Engineering

PUBLICATIONS AND MANUSCRIPTS

Journal Paper

- [1] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. Generalized resilience and robust statistics. *Annals of Statistics*.
- [2] Paria Rashidinejad, **Banghua Zhu**, Cong Ma, Jiantao Jiao, Stuart Russell. Bridging Offline Reinforcement Learning and Imitation Learning: A Tale of Pessimism. *IEEE Transactions on Information Theory*.
- [3] Cong Ma, **Banghua Zhu**, Jiantao Jiao, Martin J Wainwright. Minimax Off-Policy Evaluation for Multi-Armed Bandits. *IEEE Transactions on Information Theory*.
- [4] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. Robust estimation via generalized quasi-gradients *Information and Inference: A Journal of the IMA*.
- [5] **Banghua Zhu**, Jiantao Jiao, David Tse. Deconstructing Generative Adversarial Networks. *IEEE Transactions on Information Theory*
- [6] **Banghua Zhu**^{*}, Ziao Wang^{*}, Nadim Ghaddar^{*}, Jiantao Jiao, Lele Wang. Noisy Computing of the OR and MAX Functions. *IEEE Journal on Selected Areas in Information Theory*

Conference Paper

- [1] **Banghua Zhu**, Michael I. Jordan, Jiantao Jiao. Data Refinement: Mitigating Reward Overfitting and Overoptimization in RLHF. *ICML 2024*.
- [2] Ying Sheng, Shiyi Cao, Dacheng Li, Coleman Richard Charles Hooper, Nicholas Lee, Shuo Yang, Christopher Chou, **Banghua Zhu**, Lianmin Zheng, Kurt Keutzer, Joseph E. Gonzalez, Ion Stoica. SLoRA: Scalable Serving of Thousands of LoRA Adapters. *MLSys 2024*.
- [3] Qingyue Zhao, **Banghua Zhu**. Towards the Fundamental Limits of Knowledge Transfer over Finite Domains. *ICLR 2024*.
- [4] Jinning Li, Xinyi Liu, **Banghua Zhu**, Jiantao Jiao, Masayoshi Tomizuka, Chen Tang, Wei Zhan. Guided Online Distillation: Promoting Safe Reinforcement Learning by Offline Demonstration. *ICRA 2024*.

- [5] Cassidy Laidlaw, **Banghua Zhu**, Stuart Russell, Anca Dragan. A Theoretical Explanation of Deep RL Performance in Stochastic Environments. *ICLR 2024*.
- [6] Tianhao Wu, **Banghua Zhu**, Ruoyu Zhang, Zhaojin Wen, Kannan Ramchandran, Jiantao Jiao. Pairwise Proximal Policy Optimization: Harnessing Relative Feedback for LLM Alignment. *In submission. Preliminary version on NeurIPS 2023 Foundation Models for Decision Making Workshop*.
- [7] Venkat Krishna Srinivasan, Zhen Dong, **Banghua Zhu**, Brian Yu, Hanzi Mao, Damon Mosk-Aoyama, Kurt Keutzer, Jiantao Jiao, Jian Zhang. NexusRaven: a commercially-permissive Language Model for Function Calling. *In submission. Preliminary version on NeurIPS 2023 Foundation Models for Decision Making Workshop*.
- [8] Baihe Huang, **Banghua Zhu**, Hanlin Zhu, Jason Lee, Jiantao Jiao and Michael I. Jordan. Towards Optimal Statistical Watermarking. *In submission. Preliminary version on NeurIPS 2023 Workshop on Socially Responsible Language Modelling Research*.
- [9] Hanlin Zhu, **Banghua Zhu**, Jiantao Jiao. Efficient Prompt Caching for Large Language Model Inference via Embedding Similarity. *In submission. Preliminary version on NeurIPS 2023 Workshop on Machine Learning for Systems*.
- [10] **Banghua Zhu**, Ying Sheng, Lianmin Zheng, Clark Barrett, Michael I. Jordan, Jiantao Jiao. On Optimal Caching and Model Multiplexing for Large Model Inference. *NeurIPS 2023*.
- [11] **Banghua Zhu**, Mingyu Ding, Philip Jacobson, Ming Wu, Wei Zhan, Michael I. Jordan, Jiantao Jiao. Doubly Robust Self-Training. *NeurIPS 2023*.
- [12] **Banghua Zhu**, Jiantao Jiao, Michael I. Jordan. Principled Reinforcement Learning with Human Feedback from Pairwise or K -wise Comparisons. *ICML 2023*.
- [13] Geng Zhao*, **Banghua Zhu***, Jiantao Jiao, Michael I. Jordan. Online Learning in Stackelberg Games with an Omniscient Follower. *ICML 2023*.
- [14] Ikechukwu Uchendu, Ted Xiao, Yao Lu, **Banghua Zhu**, Mengyuan Yan, Joséphine Simon, Matthew Bennis, Chuyuan Fu, Cong Ma, Jiantao Jiao, Sergey Levine, Karol Hausman. Jump-Start Reinforcement Learning. *ICML 2023*.
- [15] **Banghua Zhu**, Ziao Wang, Nadim Ghaddar, Jiantao Jiao, Lele Wang. On the Optimal Bounds for Noisy Computing. *ISIT 2023*.
- [16] Ziao Wang, Nadim Ghaddar, **Banghua Zhu**, Lele Wang. Variable-Length Insertion-Based Noisy Sorting. *ISIT 2023*.
- [17] Ziao Wang, Nadim Ghaddar, **Banghua Zhu**, Lele Wang. Noisy Sorting Capacity. *In submission. Preliminary version is published in ISIT 2022*.
- [18] **Banghua Zhu**, Stephen Bates, Zhuoran Yang, Yixin Wang, Jiantao Jiao, Michael I. Jordan. The Sample Complexity of Online Contract Design. *EC 2023*.
- [19] **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*.
- [20] **Banghua Zhu**, Jiantao Jiao, Michael I. Jordan. Robust Estimation for Nonparametric Families via Generative Adversarial Networks. *ISIT 2022*.
- [21] **Banghua Zhu**, Jiantao Jiao, Jacob Steinhardt. When does the Tukey median work? *ISIT 2020*.

Manuscripts

- [1] **Banghua Zhu***, Evan Frick*, Tianhao Wu*, Hanlin Zhu and Jiantao Jiao. Starling-7B: Improving Helpfulness and Harmlessness of LLM with RLAIIF. *In preparation*.
- [2] Zhikai Li, Xiaoxuan Liu, **Banghua Zhu**, Zhen Dong, Qingyi Gu, Kurt Keutzer. QFT: Quantized Full-parameter Tuning of LLMs with Affordable Resources. *In submission*.
- [3] **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*.
- [4] **Banghua Zhu**, Sai Praneeth Karimireddy, Jiantao Jiao, Michael I. Jordan. Online Learning in a Creator Economy. *In submission*.

PAST WORK EXPERIENCE

Research Intern at Microsoft Research, Mountain View *2023.3 - 2023.6*

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
2023	Berkeley EECS Award for Undergraduate Researcher Mentoring
2019-2021	Berkeley EECS Department Award
2019	Stanford Citadel Datathon First Prize
2018	Beijing Excellent Undergrad Award
2017	Tsinghua Science and Innovation Scholarship
2017	Bao Gang Excellent Student Scholarship
2016	Qualcomm Scholarship
2015-2016	National Scholarship
2015	Singapore Technologies Engineering China Scholarship