

Yu Bai

Website: <https://yubai.org> E-mail: yubai.pku@gmail.com Twitter: @yubai01

RESEARCH INTERESTS

Foundation Models/ Large Language Models, Learning Theory, Deep Learning, Reinforcement Learning, Game Theory, Uncertainty Quantification, Non-Convex Optimization.

EMPLOYMENT

Senior Research Scientist, *Salesforce AI Research* Palo Alto, CA
Fundamental AI research on machine learning and its theory, including foundation models/LLMs, reinforcement learning, game theory, and deep learning. Oct 2019 - current

EDUCATION

Ph.D. in Statistics, Stanford University Stanford, CA
Specialization in Machine Learning. Sep 2014 - Sep 2019
Advisor: Prof. John C. Duchi.
Thesis: *When do gradient methods work well in non-convex learning problems?*

B.S. in Mathematics, Peking University Beijing, China
Honored Graduate Sep 2010 - July 2014

PUBLICATIONS

(I) Conference and Journal Publications

- Transformers as Statisticians: Provable In-Context Learning with In-Context Algorithm Selection.
Yu Bai, Fan Chen, Huan Wang, Caiming Xiong, Song Mei.
Neural Information Processing Systems (NeurIPS) 2023. Oral presentation.
- What can a Single Attention Layer Learn? A Study Through the Random Features Lens.
Hengyu Fu, Tianyu Guo, Yu Bai, Song Mei.
Neural Information Processing Systems (NeurIPS) 2023.
- Efficient RL with Impaired Observability: Learning to Act with Delayed and Missing State Observations.
Minshuo Chen, Yu Bai, H. Vincent Poor, Mengdi Wang.
Neural Information Processing Systems (NeurIPS) 2023.
- Breaking the Curse of Multiagency: Provably Efficient Decentralized Multi-Agent RL with Function Approximation.
Yuanhao Wang, Qinghua Liu, Yu Bai, Chi Jin.
Conference on Learning Theory (COLT) 2023.
- Lower Bounds for Learning in Revealing POMDPs.
Fan Chen, Huan Wang, Caiming Xiong, Song Mei, Yu Bai.
International Conference on Machine Learning (ICML) 2023.
- Improved Online Conformal Prediction via Strongly Adaptive Online Learning.
Aadyot Bhatnagar, Huan Wang, Caiming Xiong, Yu Bai.
International Conference on Machine Learning (ICML) 2023.

7. Offline Learning in Markov Games with General Function Approximation.
Yuheng Zhang, Yu Bai, Nan Jiang.
International Conference on Machine Learning (ICML) 2023.
8. Partially Observable RL with B-Stability: Unified Structural Condition and Sharp Sample-Efficient Algorithms.
Fan Chen, Yu Bai, Song Mei.
International Conference on Learning Representations (ICLR) 2023.
Notable-top-25% (Spotlight presentation).
9. The Role of Coverage in Online Reinforcement Learning.
Tengyang Xie, Dylan J. Foster, Yu Bai, Nan Jiang, Sham M. Kakade.
International Conference on Learning Representations (ICLR) 2023.
Notable-top-5% (Oral presentation).
10. Learning Rationalizable Equilibria in Multiplayer Games.
Yuanhao Wang, Dingwen Kong, Yu Bai, Chi Jin.
International Conference on Learning Representations (ICLR) 2023.
11. Identifying Good Directions to Escape the NTK Regime and Efficiently Learn Low-Degree Plus Sparse Polynomials.
Eshaan Nichani, Yu Bai, Jason D. Lee.
Neural Information Processing Systems (NeurIPS) 2022.
12. Policy Optimization for Markov Games: Unified Framework and Faster Convergence.
Runyu Zhang, Qinghua Liu, Huan Wang, Caiming Xiong, Na Li, Yu Bai.
Neural Information Processing Systems (NeurIPS) 2022.
13. Efficient Phi-Regret Minimization in Extensive-Form Games via Online Mirror Descent.
Yu Bai, Chi Jin, Song Mei, Ziang Song, Tiancheng Yu.
Neural Information Processing Systems (NeurIPS) 2022. **Oral presentation.**
14. Sample-Efficient Learning of Correlated Equilibria in Extensive-Form Games.
Ziang Song, Song Mei, Yu Bai.
Neural Information Processing Systems (NeurIPS) 2022.
15. Conformal Predictor for Improving Zero-Shot Text Classification Efficiency.
Prafulla Kumar Choubey, Yu Bai, Chien-Sheng Wu, Wenhao Liu, Nazneen Rajani.
Empirical Methods in Natural Language Processing (EMNLP) 2022.
16. Local Calibration: Metrics and Recalibration.
Rachel Luo, Aadyot Bhatnagar, Yu Bai, Shengjia Zhao, Huan Wang, Caiming Xiong, Silvio Savarese, Edward Schmerling, Marco Pavone.
Conference on Uncertainty in Artificial Intelligence (UAI) 2022.
17. Near-Optimal Learning of Extensive-Form Games with Imperfect Information.
Yu Bai, Chi Jin, Song Mei, Tiancheng Yu.
International Conference on Machine Learning (ICML) 2022.
18. When Can We Solve General-Sum Markov Games with a Large Number of Players Sample-Efficiently?
Ziang Song, Song Mei, Yu Bai.
International Conference on Learning Representations (ICLR) 2022.
19. Efficient and Differentiable Conformal Prediction with General Function Classes.
Yu Bai, Song Mei, Huan Wang, Yingbo Zhou, Caiming Xiong.
International Conference on Learning Representations (ICLR) 2022.

20. Understanding the Under-Coverage Bias in Uncertainty Estimation.
Yu Bai, Song Mei, Huan Wang, Caiming Xiong.
*Neural Information Processing Systems (NeurIPS) 2021. **Spotlight presentation.***
21. Policy Finetuning: Bridging Sample-Efficient Online and Offline Reinforcement Learning.
Tengyang Xie, Nan Jiang, Huan Wang, Caiming Xiong, Yu Bai.
Neural Information Processing Systems (NeurIPS) 2021.
22. Sample-Efficient Learning of Stackelberg Equilibria in General-Sum Games.
Yu Bai, Chi Jin, Huan Wang, Caiming Xiong.
Neural Information Processing Systems (NeurIPS) 2021.
23. Near-Optimal Offline Reinforcement Learning via Double Variance Reduction.
Ming Yin, Yu Bai, Yu-Xiang Wang.
Neural Information Processing Systems (NeurIPS) 2021.
24. Don't Just Blame Over-parameterization for Over-confidence: Theoretical Analysis of Calibration in Binary Classification.
Yu Bai, Song Mei, Huan Wang, Caiming Xiong.
International Conference on Machine Learning (ICML) 2021.
25. Exact Gap between Generalization Error and Uniform Convergence in Random Feature Models.
Zitong Yang, Yu Bai, Song Mei.
International Conference on Machine Learning (ICML) 2021.
26. How Important is the Train-Validation Split in Meta-Learning?
Yu Bai, Minshuo Chen, Pan Zhou, Tuo Zhao, Jason D. Lee, Sham Kakade, Huan Wang, Caiming Xiong.
International Conference on Machine Learning (ICML) 2021.
27. A Sharp Analysis of Model-based Reinforcement Learning with Self-Play.
Qinghua Liu, Tiancheng Yu, Yu Bai, Chi Jin.
International Conference on Machine Learning (ICML) 2021.
28. Near Optimal Provable Uniform Convergence in Off-Policy Evaluation for Reinforcement Learning.
Ming Yin, Yu Bai, Yu-Xiang Wang.
*Artificial Intelligence and Statistics (AISTATS) 2021. **Oral presentation.***
29. Towards Understanding Hierarchical Learning: Benefits of Neural Representations.
Minshuo Chen, Yu Bai, Jason D. Lee, Tuo Zhao, Huan Wang, Caiming Xiong, Richard Socher.
Neural Information Processing Systems (NeurIPS) 2020.
30. Near-Optimal Reinforcement Learning via Self-Play.
Yu Bai, Chi Jin, Tiancheng Yu.
Neural Information Processing Systems (NeurIPS) 2020.
31. Provable Self-Play Algorithms for Competitive Reinforcement Learning.
Yu Bai, Chi Jin.
International Conference on Machine Learning (ICML) 2020.
32. Beyond Linearization: On Quadratic and Higher-Order Approximation of Wide Neural Networks.
Yu Bai, Jason D. Lee.
International Conference on Learning Representations (ICLR) 2020.

33. Provably Efficient Q-Learning with Low Switching Cost.
Yu Bai, Tengyang Xie, Nan Jiang, Yu-Xiang Wang.
Neural Information Processing Systems (NeurIPS) 2019.
34. ProxQuant: Quantized Neural Networks via Proximal Operators
Yu Bai, Edo Liberty, Yu-Xiang Wang.
International Conference on Learning Representations (ICLR) 2019.
35. Subgradient Descent Learns Orthogonal Dictionaries.
Yu Bai, Qijia Jiang, Ju Sun.
International Conference on Learning Representations (ICLR) 2019.
36. Approximability of Discriminators Implies Diversity in GANs.
Yu Bai, Tengyu Ma, Andrej Risteski.
International Conference on Learning Representations (ICLR) 2019.
37. The Landscape of Empirical Risk for Non-convex Losses.
Song Mei, Yu Bai, Andrea Montanari.
The Annals of Statistics 46 (6A), 2747-2774, 2018.

(II) Preprints

1. Transformers as Decision Makers: Provable In-Context Reinforcement Learning via Supervised Pretraining.
Licong Lin, Yu Bai, Song Mei.
Submitted, 2023. arXiv preprint arXiv:2310.08566.
2. How Do Transformers Learn In-Context Beyond Simple Functions? A Case Study on Learning with Representations.
Tianyu Guo, Wei Hu, Song Mei, Huan Wang, Caiming Xiong, Silvio Savarese, Yu Bai.
Submitted, 2023.
3. Text2Data: Low-Resource Data Generation with Textual Control.
Shiyu Wang, Yihao Feng, Tian Lan, Ning Yu, Yu Bai, Ran Xu, Huan Wang, Caiming Xiong, Silvio Savarese.
Submitted, 2023.
4. Is Inverse Reinforcement Learning Harder than Standard Reinforcement Learning?
Lei Zhao, Mengdi Wang, Yu Bai.
Submitted, 2023.
5. Sample-Efficient Learning of POMDPs with Multiple Observations In Hindsight.
Jiacheng Guo, Minshuo Chen, Huan Wang, Caiming Xiong, Mengdi Wang, Yu Bai.
Submitted, 2023. arXiv preprint arXiv:2307.02884.
6. Unified Algorithms for RL with Decision-Estimation Coefficients: No-Regret, PAC, and Reward-Free Learning.
Fan Chen, Song Mei, Yu Bai.
Submitted, 2022. arXiv preprint arXiv:2209.11745.

INVITED TALKS

Transformers as Statisticians: Provable In-Context Learning with In-Context Algorithm Selection.

Invited talk at Airbnb, July 2023.

Partially Observable RL with B-Stability: Unified Structural Condition and Sharp Sample-Efficient Algorithms.

Invited talk at SIAM OP23, Seattle, June 2023.

Recent Progresses on the Theory of Multi-Agent Reinforcement Learning and Games.

Guest Lecture at CS332, Stanford University, Oct 2022.

When Can We Learn General-Sum Markov Games Sample-Efficiently with A Large Number of Players?

RL Theory Virtual Seminars, May 2022.

Near-Optimal Learning of Extensive-Form Games with Imperfect Information

CISS Conference, Princeton University, March 2022.

Learning and Games Program, Simons Institute on the Theory of Computing, April 2022.

Understanding the Under-Coverage Bias in Uncertainty Estimation

Statistics Department Seminar, Rutgers University, October 2021.

Spotlight presentation at ICML 2021 Workshop on Distribution-free Uncertainty Quantification, July 2021.

Sample-Efficient Learning of Stackelberg Equilibria in General-Sum Games

Spotlight presentation at ICML 2021 Workshop on Reinforcement Learning Theory, July 2021.

How Important is the Train-Validation Split in Meta-Learning?

One World Seminar on the Mathematics of Machine Learning, October 2020.

Provable Self-Play Algorithms for Competitive Reinforcement Learning.

Facebook AI Research, March 2020.

Beyond Linearization: On Quadratic and Higher-Order Approximation of Wide Neural Networks.

Simons Institute on the Theory of Computing, August 2020.

ProxQuant: Quantizing Neural Networks via Proximal Operators

Bytedance AI Lab, December 2018.

Amazon AI, September 2018.

On the Generalization and Approximation in GANs

Google Brain, November 2018.

Salesforce Research, November 2018.

Stanford ML Seminar, October 2018.

Optimization Landscape of Some Non-convex Learning Problems

Stanford Theory Seminar, April 2018.

Stanford ML Seminar, April 2017.

SERVICE

Area Chair / Senior Program Committee

- International Conference on Learning Representations (ICLR) 2024
- Neural Information Processing Systems (NeurIPS) 2023
- Artificial Intelligence and Statistics (AISTats) 2023-2024

Associate Editor at Journals

- Transactions of Machine Learning Research (TMLR)

Reviewing

- Conference:
 - Neural Information Processing Systems (NeurIPS) 2018-2022
 - International Conference on Machine Learning (ICML) 2019-2021, 2023
 - International Conference on Learning Representations (ICLR) 2019-2023
 - Conference on Computational Learning Theory (COLT) 2019-2020, 2022-2023
 - Algorithmic Learning Theory (ALT) 2024
 - Artificial Intelligence and Statistics (AISTats) 2020
 - IEEE International Symposium on Information Theory (IEEE-ISIT) 2018
- Journal:
 - The Annals of Statistics (AoS)
 - Journal of the American Statistical Association (JASA)
 - Journal of the Royal Statistical Society, Series B (JRSS-B)
 - Journal of Machine Learning Research (JMLR)
 - Transactions of Machine Learning Research (TMLR)
 - IEEE Transactions on Signal Processing (IEEE-TSP)
 - SIAM Journal on Control and Optimization (SICON)

MENTORING EXPERIENCE

Interns

- Shiyu Wang (co-mentored) 2023
Ph.D. at Emory University. Summer Intern at Salesforce.
- Runyu Zhang 2022
Ph.D. at Harvard SEAS. Summer Intern at Salesforce.
- Michael Curry (co-mentored) 2021
Ph.D. at Maryland CS. Summer Intern at Salesforce. → Postdoc at University of Zurich.
- Tolga Ergen 2021
Ph.D. at Stanford EE. Summer Intern at Salesforce. → Samsung AI Research.
- Minshuo Chen 2020-2021
Ph.D. at Georgia Tech ISyE. Summer Intern at Salesforce. → Postdoc at Princeton ECE.
- Darshan Thaker 2020
Masters at Columbia CS. Spring Intern at Salesforce. → Ph.D. at JHU CS.

Undergraduate Student Mentoring

- Lei Zhao 2023
Undergraduate at USTC.
- Hengyu Fu 2023
Undergraduate at Peking University.
- Jiacheng Guo 2023
Undergraduate at Fudan University. → Ph.D. at Princeton ECE.
- Fan Chen 2022-2023
Undergraduate at Peking University. → Ph.D. at MIT EECS.
- Dingwen Kong 2022
Undergraduate at Peking University. → Ph.D. at MIT EECS.

- Ziang Song
Undergraduate at Peking University → Ph.D. at Stanford Statistics.

2021-2022

TEACHING EXPERIENCE

As Instructor:

Guest Lecturer, Nonparametric Statistics (Stats205), Fall 2019.
Guest Lecturer, Theory of Statistics (Stats300B), Spring 2018.
Session Instructor, Theory of Probability (Stats310A), Fall 2017.

As Teaching Assistant (selected):

Statistical Learning Theory (CS229T), as head TA.
Modern Markov Chains (Stats 318).
Theory of Probability (Stats310A/B/C).
Theory of Statistics (Stats300A/B).
Statistical Inference (Stats200).
Introduction to Stochastic Processes (Stats217).

INTERNSHIPS

Research Intern, *Amazon AI*

Host: Edo Liberty & Yu-Xiang Wang

Palo Alto, CA

June 2018 - Sep 2018

Proposed ProxQuant, a prox-gradient method with quantization-inducing regularizers for training quantized neural networks. Paper published in ICLR 2019.

Research Intern, *Google Research*

Host: Li Zhang

Mountain View, CA

June 2016 - Sep 2016

Proposed adaptive sampling strategies for softmax in deep networks for extreme classification which achieved state-of-the-art accuracy on a large-scale Youtube benchmark dataset. Algorithm implemented in Tensorflow (`tf.contrib.nn.rank_sampled_softmax_loss`).