Ruqi Zhang ruqiz@purdue.edu · ruqizhang.github.io

Education

Ph.D., Statistics; MS, Computer Science, Cornell University, 2016-2021

Committee: Christopher De Sa, Thorsten Joachims, Giles Hooker

B.S., Mathematics, Renmin University of China, 2012-2016

Professional Experience

2022–now	*Assistant Professor*, Computer Science, Purdue University
2021–2022	*Postdoctoral Fellow*, Institute for Foundations of Machine Learning, UT Austin
2020	*Research Intern*, Microsoft Research New England
2019	*Research Intern*, Microsoft Research Cambridge, UK

Publications

1. Training bayesian neural networks with sparse subspace variational inference.
   J. Li, Z. Miao, Q. Qiu, and R. Zhang.
   International Conference on Learning Representations, 2024

2. Entropy-mcmc: Sampling from flat basins with ease.
   B. Li and R. Zhang.
   International Conference on Learning Representations, 2024

3. Balance is essence: Accelerating sparse training via adaptive gradient correction.
   B. Lei, D. Xu, R. Zhang, S. He, and B. K. Mallick.
   Conference on Parsimony and Learning, 2024

4. Rethinking data distillation: Do not overlook calibration.
   D. Zhu, B. Lei, J. Zhang, Y. Fang, Y. Xie, R. Zhang, and D. Xu.
   In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4935–4945, 2023

5. Discs: A benchmark for discrete sampling.
   K. Goshvadi, H. Sun, X. Liu, A. Nova, R. Zhang, W. S. Grathwohl, D. Schuurmans, and H. Dai.
   In *Thirty-seventh Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2023

6. Analysis of climate campaigns on social media using bayesian model averaging.
   T. Islam, R. Zhang, and D. Goldwasser.
   AAAI/ACM Conference on AI, Ethics, and Society, 2023.
   Oral Presentation, top 11%

7. Dp-fast mh: Private, fast, and accurate metropolis-hastings for large-scale bayesian inference.
   W. Zhang and R. Zhang.
   International Conference on Machine Learning, 2023

8. Calibrating the rigged lottery: Making all tickets reliable.
   B. Lei, R. Zhang, D. Xu, and B. Mallick.
   International Conference on Learning Representations, 2023

9. Efficient informed proposals for discrete distributions via newton's series approximation.
   Y. Xiang, D. Zhu, B. Lei, D. Xu, and R. Zhang.
   In *International Conference on Artificial Intelligence and Statistics*, pages 7288–7310. PMLR, 2023

10. Sampling in constrained domains with orthogonal-space variational gradient descent.
    R. Zhang, Q. Liu, and X. Tong.
    Advances in Neural Information Processing Systems, 35, 37108–37120, 2022

11. A langevin-like sampler for discrete distributions.
R. Zhang, X. Liu, and Q. Liu.
In *International Conference on Machine Learning*, pages 26375–26396, 2022

12. Low-precision stochastic gradient langevin dynamics.
R. Zhang, A. G. Wilson, and C. De Sa.
In *International Conference on Machine Learning*, pages 26624–26644. PMLR, 2022

13. Meta-learning divergences for variational inference.
R. Zhang, Y. Li, C. De Sa, S. Devlin, and C. Zhang.
In *International Conference on Artificial Intelligence and Statistics*, pages 4024–4032. PMLR, 2021

14. Asymptotically optimal exact minibatch metropolis-hastings.
R. Zhang, A. F. Cooper, and C. M. De Sa.
Advances in Neural Information Processing Systems, 33, 19500–19510, 2020.
Spotlight Presentation, top 3%

15. Amagold: Amortized metropolis adjustment for efficient stochastic gradient mcmc.
R. Zhang, A. F. Cooper, and C. De Sa.
In *International Conference on Artificial Intelligence and Statistics*, pages 2142–2152. PMLR, 2020

16. Cyclical stochastic gradient mcmc for bayesian deep learning.
R. Zhang, C. Li, J. Zhang, C. Chen, and A. G. Wilson.
In *International Conference on Learning Representations*, 2020.
Oral Presentation, top 2%

17. Poisson-minibatching for gibbs sampling with convergence rate guarantees.
R. Zhang and C. M. De Sa.
In *Advances in Neural Information Processing Systems*, pages 4922–4931, 2019.
Spotlight Presentation, top 2.5%

18. Large scale sparse clustering.
R. Zhang and Z. Lu.
In *International Joint Conference on Artificial Intelligence*, pages 2336–2342, 2016

Awards and Honors

Ross-Lynn Research Scholar Fund · 2023

Conference on Uncertainty in Artificial Intelligence (UAI) Top Reviewer · 2023

ICML Best Reviewers (Top 10%) · 2021

Spotlight Rising Star in Data Science · University of Chicago · 2020

NeurIPS Top 10% Reviewers Award · 2020

NeurIPS Travel Grant · 2019

Academic Outstanding Scholarship · Renmin University of China · 2013-2015

Exchange Students Scholarship · University of Helsinki · 2015

Talks

Low-precision Sampling for Probabilistic Deep Learning, Invited talk at NeurIPS Workshop on Machine Learning with New Compute Paradigms, NeurIPS, December 2023

Sampling in Discrete and Constrained Domains, Invited talk at ICML Workshop on Structured Probabilistic Inference & Generative Modeling, ICML, July 2023

Scalable and Reliable Inference for Probabilistic Modeling, Invited talk at Center for Data Science and Machine Learning, National University of Singapore, October 2022

A Langevin-like Sampler for Discrete Distributions, Spotlight presentation at ICML, July 2022

Low-Precision Stochastic Gradient Langevin Dynamics, Spotlight presentation at ICML, July 2022

Scalable and Reliable Inference for Probabilistic Modeling, Invited talk at Simons Institute, UC Berkeley, November 2021

Asymptotically Optimal Exact Minibatch Metropolis-Hastings, Spotlight talk in Rising Stars in Data Science Workshop at University of Chicago, January 2021

Asymptotically Optimal Exact Minibatch Metropolis-Hastings, Spotlight presentation at NeurIPS, December 2020

Cyclical Stochastic Gradient MCMC for Bayesian Deep Learning, Oral presentation at ICLR, April 2020

Poisson-Minibatching for Gibbs Sampling with Convergence Rate Guarantees, Spotlight presentation at NeurIPS, December 2019

Teaching

Purdue · CS 37300 · Data Mining And Machine Learning · Spring 2024
Graduate lecture course

Purdue · CS 57800 · Statistical Machine Learning · Fall 2023
Graduate lecture course

Purdue · CS 57800 · Statistical Machine Learning · Spring 2023
Graduate lecture course

Purdue · CS 59200 · Probabilistic Machine Learning · Fall 2022
Graduate seminar course

Purdue · CS 57100 · Artificial Intelligence · Fall 2022, Fall 2023
Guest Lecturer

Cornell · CS 4820, Introduction to Analysis of Algorithms · Spring 2021
Head Teaching Assistant

Cornell · ILRST 5050 · Statistics at Work · Fall 2018
Teaching Assistant

Cornell · STSCI 2100 · Introductory Statistics · Spring 2018
Teaching Assistant

Cornell · MATH 3110 · Introduction to Analysis · Spring 2017
Teaching Assistant,

Cornell · STSCI 3110 · Probability Models and Inference for the Social Sciences · Fall 2016, Fall 2017, Fall 2020
Teaching Assistant

Service

Organizer

4th, 5th, 6th Symposium on Advances in Approximate Bayesian Inference (AABI) · 2022&2023&2024

ICML Workshop on Sampling and Optimization in Discrete Space · 2023

ICML Women in Machine Learning (WiML) Un-Workshop on Safely Navigating Scalability-reliability Trade-offs in ML Methods · 2021

Reviewer

Transactions on Machine Learning Research

Statistics and Computing

NeurIPS 2018, 2019, 2020, 2021, 2022, 2023

NeurIPS Workshop Proposal 2024

ICML 2019, 2020, 2021 (expert reviewer), 2022, 2023

ICLR 2019, 2020, 2021, 2022, 2023

AISTATS 2020, 2021, 2022, 2023

AAAI 2020

UAI 2019

MLSys 2024

Symposium on Advances in Approximate Bayesian Inference 2019, 2020

I Can't Believe It's Not Better@NeurIPS 2020

Neural Compression Workshop@ICLR 2021

Panelist

NSF Panel 2022, 2023

Mentor

AISTATS 2022 Submission Mentoring Program

Department

Purdue CS PhD Admission Committee 2023, 2024