Jiaxiang Li

5-151 Keller Hall
200 Union Street Se
Minneapolis, MN 55455

Email: li003755@umn.edu
Phone: +1 (530)-220-3138
Homepage: https://jasonjiaxiangli.github.io/

Academic employment

University of Minnesota — Minneapolis, MN
Postdoctoral Associate, Department of Electrical and Computer Engineering — 2024–Now
Mentor: Mingyi Hong

Education

University of California, Davis — Davis, CA
Ph.D. in Applied Mathematics, GPA: 3.97/4.00 — 2018–2023
Advisors: Shiqian Ma and Krishnakumar Balasubramanian

Zhejiang University — Hangzhou, China
B.S. in Mathematics, GPA: 3.88/4.00 — 2014–2018

Research Interest

My research interests lie at the optimization problems arising in machine learning, operations research and other applications. Specifically, I am working on the following topics:

- Computational complexities for gradient-based, gradient-free (zeroth-order) and primal-dual algorithms for solving nonconvex optimizations and optimization on Riemannian manifolds.
- Convergence theory for deterministic and stochastic minimax/bilevel problems with applications in operations research and machine learning.
- Distributed optimization algorithms, including decentralized and federated learning.
- Theories for reinforcement learning algorithms, especially for multi-armed bandits and convergence of policy-based methods, and their applications in large language model (LLM) alignments.
- Theory and computation of supervised and unsupervised learning in general.

Fundings

- I am a **Co-PI** of NSF grant ECCS-2426064. Prof. Mingyi Hong is the PI.

Honors and Awards

- INFORMS Computing Society Prize — 2024
- INFORMS Optimization Conference Travel Award — 2024
- Graduate Student Fellowship at UC Davis — 2018–2019, 2022-2023

Publications and Preprints (* marks co-first author)

[1] X. Zhang, C. Li, S. Zeng, **J. Li**, Z. Wang, S. Lu, A. Garcia, and M. Hong, "Reinforcement Learning in Inference Time: A Perspective from Successive Policy Iterations", in *ICLR 2025 Workshop on Reasoning and Planning for Large Language Models*.

[2] A. Glentis*, **J. Li***, Q. Shang, A. Han, I. Tsaknakis, Q. Wei, and M. Hong, "Scalable parameter and memory efficient pretraining for llm: Recent algorithmic advances and benchmarking", *arXiv preprint arXiv:2505.22922*, 2025.

[3] C. Li, S. Zeng, Z. Liao, **J. Li**, D. Kang, A. Garcia, and M. Hong, "Joint reward and policy learning with demonstrations and human feedback improves alignment", in *The Thirteenth International Conference on Learning Representations (ICLR, Spotlight)*, 2025.

[4] **J. Li** and S. Ma, "Riemannian bilevel optimization", *Journal of Machine Learning Research*, vol. 26, no. 18, pp. 1–44, 2025.

[5] S. Zeng, L. Viano, C. Li, **J. Li**, V. Cevher, M. Wulfmeier, S. Ermon, A. Garcia, and M. Hong, "Aligning large language models with human feedback: Mathematical foundations and algorithm design", *Authorea Preprints*, 2025.

[6] A. Han, **J. Li**, W. Huang, M. Hong, A. Takeda, P. Jawanpuria, and B. Mishra, "SLTrain: A sparse plus low-rank approach for parameter and memory efficient pretraining", in *Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[7] X. Jiang, **J. Li**, M. Hong, and S. Zhang, "Barrier function for bilevel optimization with coupled lower-level constraints: Formulation, approximation and algorithms", *arXiv preprint arXiv:2410.10670 (Submitted to Mathematical Programming)*, 2024.

[8] **J. Li**, K. Balasubramanian, and S. Ma, "Zeroth-order Riemannian averaging stochastic approximation algorithms", *SIAM Journal on Optimization*, vol. 34, no. 4, pp. 3314–3341, 2024.

[9] **J. Li**, X. Chen, S. Ma, and M. Hong, "Problem-Parameter-Free Decentralized Nonconvex Stochastic Optimization", *Pacific Journal on Optimization (Accepted), arXiv:2402.08821*, 2024.

[10] **J. Li**, S. Ma, and T. Srivastava, "A Riemannian Alternating Direction Method of Multipliers", *Mathematics of Operations Research (Accepted), arXiv:2211.02163*, 2024.

[11] **J. Li**, S. Zeng, H.-T. Wai, C. Li, A. Garcia, and M. Hong, "Getting More Juice Out of the SFT Data: Reward Learning from Human Demonstration Improves SFT for LLM Alignment", in *Thirty-Eighth Annual Conference on Neural Information Processing Systems (NeurIPS)*, 2024.

[12] X. Zhang, S. Zeng, **J. Li**, K. Lin, and M. Hong, "Policy optimization can be memory-efficient: LLM Alignment Through Successive Policy Re-weighting (SPR)", in *NeurIPS 2024 Workshop on Fine-Tuning in Modern Machine Learning: Principles and Scalability*, 2024.

[13] Y. Zhang*, P. Li*, J. Hong*, **J. Li***, Y. Zhang, W. Zheng, P.-Y. Chen, J. D. Lee, W. Yin, M. Hong, Z. Wang, S. Liu, and T. Chen, "Revisiting Zeroth-Order Optimization for Memory-Efficient LLM Fine-Tuning: A Benchmark", in *Proceedings of the 41st International Conference on Machine Learning (ICML)*, vol. 235, PMLR, 2024, pp. 59 173–59 190.

[14] **J. Li**, K. Balasubramanian, and S. Ma, "Stochastic Zeroth-order Riemannian Derivative Estimation and Optimization", *Mathematics of Operations Research*, vol. 48, no. 2, pp. 1183–1211, 2023.

[15] **J. Li** and S. Ma, "Federated Learning on Riemannian Manifolds", *Appl. Set-Valued Anal. Optim. 5 (2023), 213-232*, 2023.

TEACHING

- **Lecturer** at University of Minnesota
  - EE5239 Introduction to Nonlinear Optimization (Jointly teaching with Prof. Mingyi Hong)

- **Teaching Assistant** at UC Davis
  - MAT 127 Real Analysis
  - MAT 168 Optimization
  - MAT 170 Math in Data Science
  - MAT 21 ABCD Calculus Series
  - MAT 17 ABC Calculus for Biology Series

Talks

- INFORMS Annual meeting, session speaker — October 2024
  Topic: A Riemannian ADMM
- INFORMS Optimization Society Conference, session speaker — March 2024
  Topic: Zeroth order moving average methods on manifolds
- INFORMS Optimization Society Conference, session speaker — March 2022
  Topic: Zeroth order optimization on Riemannian manifolds
- Modeling and Optimization: Theory and Applications (MOPTA), session speaker — August 2021
  Topic: Zeroth order optimization on Riemannian manifolds
- SIAM Conference on Optimization, session speaker — July 2021
  Topic: Zeroth order optimization on Riemannian manifolds

Academic services

- Journal Reviewer
  INFORMS Journal on Computing, Information and Inference: A Journal of the IMA, Journal of Machine Learning Research, Computational Optimization and Applications, Journal of Scientific Computing, Journal of Global Optimization
- Conference Reviewer
  AISTATS (2021, 2022, 2023), ICML (2022, 2024), NeurIPS (2023, 2024), ICLR (2023, 2024), AAAI (2024)

Industry Experience

Amazon.com, Inc. — Seattle, WA
Applied Scientist Intern — June 2023–Sept 2023

- Amazon Prime Science ML Team
- Built customer-prime service engagement history aggregation model with 1.2% offline evaluation gain
- Developed data pipeline for engagement history features using Pyspark and AWS services
- Built customer-prime service engagement trajectory model using seq2seq modeling and Attention networks

Meta Platforms, Inc. — Menlo Park, CA
Software Engineer Intern, Machine Learning — June 2022–Sept 2022

- Ads Core Machine Learning ENG Team
- Built reinforcement learning based Slate-Q long-term value optimization for Ads recommender in Facebook
- Designed large-scale data pipeline and managed data from industry-level databases using SQL and Hadoop Hive
- Initiated and conduct large-scale deep Q-networks training in PyTorch