

# JIAXIANG LI

---

1 Meta Way,  
Menlo Park, CA 94025

Email: [jasonljx96@gmail.com](mailto:jasonljx96@gmail.com)  
Phone: +1 (530)-220-3138  
Homepage: <https://jasonjiaxiangli.github.io/>

## APPOINTMENTS

---

**Meta Platforms, Inc.**  
AI Research Scientist

Sunnyvale, CA  
Aug 2025–Now

- Large-scale model training and optimization: design and implement novel optimizer ideas (e.g. Shampoo, Muon and DiLoCo algorithms) to accelerate and stabilize AI training

**University of Minnesota**  
Postdoctoral Associate, Department of Electrical and Computer Engineering  
Mentor: Mingyi Hong and Shuzhong Zhang

Minneapolis, MN  
Jan 2024–Aug 2025

## EDUCATION

---

**University of California, Davis**  
Ph.D. in Applied Mathematics, GPA: 3.97/4.00  
Advisor: Shiqian Ma and Krishnakumar Balasubramanian

Davis, CA  
Sept 2018–Dec 2023

**Zhejiang University**  
B.S. in Mathematics, GPA: 3.88/4.00

Hangzhou, China  
Sept 2014–July 2018

## RESEARCH INTEREST

---

My research lies at the intersection of optimization/learning theory and artificial intelligence, with a focus on developing principled algorithms in modern machine learning. Specifically, I work on:

- Algorithm design for gradient-based, gradient-free (zeroth-order), and primal–dual methods for large-scale nonconvex optimization, especially foundation model pre-training and fine-tuning.
- Distributed, decentralized and federated optimization algorithms for training large-scale AI systems.
- Parameter-efficient pre-training and fine-tuning methods for large language models (LLMs).
- Theoretical foundations of reinforcement learning, especially policy-based methods, and their role in aligning LLMs.

## 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\***, A. Han, and M. Hong, “A minimalist optimizer design for llm pretraining”, *Forty-third International Conference on Machine Learning (Accepted)*, 2026.
- [3] S. Wang\*, F. Zhang\*, **J. Li\***, C. Du, C. Du, T. Pang, Z. Yang, M. Hong, and V. Y. Tan, “Muon outperforms adam in tail-end associative memory learning”, in *The Fourteenth International Conference on Learning Representations (ICLR) (Accepted)*, 2026.
- [4] Z. Zhang, M. S. Oh, Z. Qin, **J. Li**, X. Zhang, and J. Liu, “A tale of two problems: Multi-task bilevel learning meets equality constrained multi-objective optimization”, in *Forty-third International Conference on Machine Learning (Accepted)*, 2026.
- [5] 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.
- [6] C. He, **J. Li**, B. Jiang, S. Ma, and S. Zhang, “On relatively smooth optimization over riemannian manifolds”, (*Major Revision in SIAM Journal on Optimization*) *arXiv preprint arXiv:2508.03048*, 2025.
- [7] X. Jiang, **J. Li**, M. Hong, and S. Zhang, “A correspondence-driven approach for bilevel decision-making with nonconvex lower-level problems”, (*Submitted to Mathematical Programming*) *arXiv preprint arXiv:2509.01148*, 2025.
- [8] C. Li, J. Leng, **J. Li**, Y. Sun, S. Chen, S. Shahrampour, and A. Garcia, “ADARL: Adaptive Low-Rank Structures for Robust Policy Learning under Uncertainty”, *arXiv preprint arXiv:2510.11899*, 2025.
- [9] 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.
- [10] **J. Li** and S. Ma, “Riemannian bilevel optimization”, *Journal of Machine Learning Research*, vol. 26, no. 18, pp. 1–44, 2025.
- [11] **J. Li**, S. Ma, and T. Srivastava, “A riemannian alternating direction method of multipliers”, *Mathematics of Operations Research*, vol. 50, no. 4, pp. 3222–3242, 2025.
- [12] H. Nie, **J. Li**, and Z. Wen, “An augmented lagrangian value function method for lower-level constrained stochastic bilevel optimization”, (*Submitted to Journal of Scientific Computing*) *arXiv preprint arXiv:2509.24249*, 2025.
- [13] B. Song, **J. Li**, R. Wang, S. Lu, and M. Hong, “A framework for quantifying how pre-training and context benefit in-context learning”, *arXiv preprint arXiv:2510.22594*, 2025.
- [14] H. Wang, Z. Pan, C. He, **J. Li**, and B. Jiang, “Federated learning on riemannian manifolds: A gradient-free projection-based approach”, (*Submitted to IEEE Transactions on Automatic Control*) *arXiv preprint arXiv:2507.22855*, 2025.
- [15] 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”, *IEEE Signal Processing Magazine (Accepted)*, 2025.
- [16] X. Zhang, C. Li, S. Zeng, **J. Li**, Z. Wang, K. Lin, S. Lu, A. Garcia, and M. Hong, “Aligning frozen llms by reinforcement learning: An iterative reweight-then-optimize approach”, *arXiv preprint arXiv:2506.17828*, 2025.
- [17] 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.
- [18] X. Jiang, **J. Li**, M. Hong, and S. Zhang, “Barrier function for bilevel optimization with coupled lower-level constraints: Formulation, approximation and algorithms”, (*Major revision in Mathematical Programming*) *arXiv preprint arXiv:2410.10670*, 2024.
- [19] **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.
- [20] **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.
- [21] **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.

- [22] 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.
- [23] 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.
- [24] **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.
- [25] **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

---

- ICCOPT, session speaker July 2025  
Topic: A Riemannian ADMM
- 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, 2025), NeurIPS (2023, 2024, 2025), ICLR (2023, 2024), AAAI (2024)

## INTERNSHIP EXPERIENCE

---

**Amazon.com, Inc.**  
Applied Scientist Intern

Seattle, WA  
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.**

Software Engineer Intern, Machine Learning

Menlo Park, CA

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