

JIAXIANG LI 李佳翔

Durham Hall 213,
250 Perry St, Blacksburg, VA 24061

Email: jasonljx@vt.edu
Homepage: <https://jasonjiaxiangli.github.io/>

APPOINTMENTS

Virginia Tech Assistant Professor, Grado Department of Industrial and Systems Engineering	Blacksburg, VA Aug 2026–Now
Meta Platforms, Inc. AI Research Scientist	Sunnyvale, CA Aug 2025–Aug 2026
University of Minnesota Postdoctoral Associate, Department of Electrical and Computer Engineering	Minneapolis, MN Jan 2024–Aug 2025

EDUCATION

University of California, Davis Ph.D. in Applied Mathematics	Davis, CA Sept 2018–Dec 2023
Zhejiang University B.S. in Mathematics	Hangzhou, Zhejiang, China Sept 2014–July 2018

RESEARCH INTEREST

My research studies the two-way interaction between optimization, operations research (OR), and artificial intelligence (AI). My established work focuses on **optimization and OR for AI**, connecting mathematical principles with practical challenges in modern machine learning. I am also developing **AI for optimization and OR**, using learned models to improve algorithms and decision-making. Specifically, I work on:

- Optimization algorithms and training schedules for foundation-model pre-training, fine-tuning, and alignment.
- Memory-, parameter-, and communication-efficient methods for distributed, decentralized, and federated learning.
- Bilevel, minimax, multi-objective, Riemannian, and zeroth-order optimization, with theory and machine-learning applications.
- Reinforcement learning and sequential decision-making, especially policy-based methods for aligning AI systems.
- Learning-augmented optimization and AI-assisted modeling, algorithm selection, and search for optimization and OR.

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

- [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 (ICML) (Accepted)*, 2026.
- [3] C. Li, J. Leng, **J. Li**, Y. Sun, S. Chen, S. Shahrampour, and A. Garcia, “Adaptive Rank Control for Robust Reinforcement Learning”, *Transactions on Machine Learning Research*, 2026, ISSN: 2835-8856.
- [4] **J. Li**, Z. Bu, and S. Xu, “Towards joint scaling laws with optimal batch size schedules”, *arXiv preprint arXiv:2607.27731*, 2026.
- [5] 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.
- [6] 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)*, 2026.
- [7] 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 (ICML)(Accepted)*, 2026.
- [8] 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.
- [9] 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.
- [10] 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.
- [11] 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.
- [12] **J. Li** and S. Ma, “Riemannian Bilevel Optimization”, *Journal of Machine Learning Research*, vol. 26, no. 18, pp. 1–44, 2025.
- [13] **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.
- [14] 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.
- [15] 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.
- [16] 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.
- [17] 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.
- [18] 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.
- [19] 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.
- [20] **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.

- [21] **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.
- [22] **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.
- [23] 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.
- [24] 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.
- [25] **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.
- [26] **J. Li** and S. Ma, “Federated Learning on Riemannian Manifolds”, *Appl. Set-Valued Anal. Optim.* 5 (2023), 213–232, 2023.

TEACHING

- **Lecturer** at Virginia Tech
 - ISE 5405 Optimization I: Linear Programming
- **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)

PAST 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
- 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
- Initiated and conduct large-scale deep Q-networks training in PyTorch