Yuchen Fang

Department of Mathematics, University of California, Berkeley
Email:yuchenfang@berkeley.edu | Website: ychenfang.github.io

Education

University of California, Berkeley — California, US
Ph.D. in Applied Mathematics — 2023 – 2028
Advisors: Prof. James Demmel and Prof. Javad Lavaei.

The University of Chicago — Illinois, US
M.S. in Computational and Applied Mathematics — 2021 – 2022
Thesis: *Fully Stochastic Trust-Region Sequential Quadratic Programming for Equality-Constrained Optimization*
Thesis Advisors: Prof. Mladen Kolar and Prof. Mihai Anitescu.

Xiamen University — Fujian, CN
B.S. in Statistics — 2017 – 2021
Thesis Advisor: Prof. Yingxing Li.

Highlight Projects (*-equal contribution)

[2] **Yuchen Fang**, Roger Hsiao, James Demmel, and Javad Lavaei. Frozen Eyes: Fast Latent Planning with Pretrained Encoders
Submitted to *CVPR 2027*, Paper.

- Developed Frozen Eyes, a compact world model that maps frozen pretrained visual features into a 192-dimensional planning state, training only an adapter and action-conditioned predictor. Reduced downstream training to 1.8 hours on a single A6000 GPU, including feature extraction, versus 9.4 hours for LeWM.
- Achieved 90.0% success on Push-T with 0.63-second planning, exceeding DINO-WM by 16.5 percentage points while planning approximately 200× faster.
- Achieved 97.8% and 67.7% success on near- and distant-goal Two-Room navigation using V-JEPA 2 features, outperforming LeWM by 11.5 and 49.9 percentage points, respectively.
- Conducted controlled ablations across DINOv2, V-JEPA 2, and SigLIP 2, demonstrating the benefits of coarse spatial pooling and identifying task-dependent adapter architectures for manipulation and navigation.

[1] **Yuchen Fang***, Gechen Qu*, Jason Choi, Claire Tomlin, James Demmel, and Javad Lavaei. TRIP-PO: Trust-Region Interior-Point Policy Optimization
In preparation for *RSS 2027*, Paper Draft

- Developed TRIP-PO, a trust-region interior-point extension of PPO for constrained RL. Each update first takes a normal step toward feasibility, then a PPO step on a log-barrier objective whose cost weights grow as the remaining slack shrinks.
- Trained a Unitree G1 humanoid to climb stairs under 7 hardware constraints. TRIP-PO matches the reward of PPO-Lagrangian, but reduce the proportion of iterations in which constraints are violated from 97% to 1%.
- Benchmarked against 5 constrained-RL baselines on Safety-Gymnasium. TRIP-PO stays within limits on all tasks, achieving the highest return and the lowest cumulative violation.

Open-Sourced Environment (*-equal contribution)

[1] Chenyu Li*, Zeyu Sun*, **Yuchen Fang***, Yuhan Cao, and Guangting Yu. AER: An Open Environment & Benchmark For LLM Agents Making Real-World Decisions
Website, technical paper in preparation.

Publications

[5] Roger Hsiao, **Yuchen Fang**, Xiangru Huang, Ruilong Li, Hesam Rabeti, Zan Gojcic, Javad Lavaei, James Demmel, and Sophia Shao. 3DGS^2-TR: A Scalable Second-Order Trust-Region Method for 3D Gaussian Splatting
International Conference on Machine Learning (ICML) 2026, Paper Website

[4] Jihun Kim, **Yuchen Fang**, and Javad Lavaei. On the Sharp Input-Output Analysis of Nonlinear Systems under Adversarial Attacks
International Conference on Machine Learning (ICML) 2026 Spotlight, Paper

[3] Santiago Gonzalez, et al. (including **Yuchen Fang**). QEDBench: Quantifying the Alignment Gap in Automated Evaluation of University-Level Mathematical Proofs
International Conference on Machine Learning (ICML) 2026, Paper Website

[2] **Yuchen Fang**, Javad Lavaei, and Sen Na. High-Probability Complexity Bounds of Trust-Region Stochastic Sequential Quadratic Programming with Heavy-Tailed Noise
Mathematical Programming (2026), DOI: 10.1007/s10107-026-02357-x

[1] **Yuchen Fang**, Sen Na, Michael W. Mahoney, and Mladen Kolar. Fully Stochastic Trust-Region Sequential Quadratic Programming for Equality-Constrained Optimization Problems
SIAM Journal on Optimization (2024), DOI: 10.1137/22M1537862
Short version accepted by *NeurIPS Workshop on Higher-Order Optimization in Machine Learning (HOOML) 2022*

Preprints

[7] Junyu Guo, **Yuchen Fang**, Shangding Gu, James Demmel, Javad Lavaei. Remembered but Not Used: Understanding In-Context Update Failures in LLMs
Submitted to *ICLR 2027*, Paper

[6] **Yuchen Fang**, Sen Na, James Demmel, and Javad Lavaei. Trust-Region Interior-Point Stochastic Sequential Quadratic Programming for Expectation Constrained Optimization with Heavy-Tailed Noise
Submitted to *Mathematical Programming*, Paper

[5] Haoxuan Wang, **Yuchen Fang**, and Sen Na. TR-SSQP: A Trust-Region Method for Constrained Stochastic Optimization under Heavy-Tailed Noise
Submitted to *SIAM Journal on Optimization*, arXiv

[4] **Yuchen Fang**, James Demmel, and Javad Lavaei. Why is Normalization Preferred? A Worst-Case Complexity Theory for Stochastically Preconditioned SGD under Heavy-Tailed Noise
Submitted to *Journal of Machine Learning Research*, arXiv

[3] **Yuchen Fang**, Xinshou Zheng, and Javad Lavaei. TRSVR: An Adaptive Stochastic Trust-Region Method with Variance Reduction
Submitted to *ICLR 2027*, arXiv

[2] **Yuchen Fang**, Jihun Kim, Sen Na, James Demmel, and Javad Lavaei. A Trust-Region Interior-Point Stochastic Sequential Quadratic Programming Method
Minor revision at *SIAM Journal on Optimization*, arXiv.

[1] **Yuchen Fang**, Sen Na, Michael W. Mahoney, and Mladen Kolar. Trust-Region Sequential Quadratic Programming for Stochastic Optimization with Random Models
Revision at *Mathematical Programming*, arXiv
Short version accepted by *NeurIPS Workshop on Higher-Order Optimization in Machine Learning (HOOML) 2022*

Presentations

[6] Nov. 2026, "Trust-Region Interior-Point Stochastic Sequential Quadratic Programming for Expectation Constrained Optimization with Heavy-Tailed Noise", *INFORMS Annual Meeting*

[5] Oct. 2025, "High-Probability Complexity Bounds of Trust-Region Stochastic Sequential Quadratic Programming with Heavy-Tailed Noise", *INFORMS Annual Meeting*

[4] Dec. 2022, "Trust-Region Sequential Quadratic Programming for Stochastic Optimization with Random Models: First-Order Stationarity", *NeurIPS Workshop on Higher-Order Optimization in Machine Learning (HOOML) 2022*

[3] Dec. 2022, "Fully Stochastic Trust-Region Sequential Quadratic Programming for Equality-Constrained Optimization Problems", *NeurIPS Workshop on Higher-Order Optimization in Machine Learning (HOOML) 2022*

[2] Nov. 2022, "Fully Stochastic Trust-Region Sequential Quadratic Programming for Equality-Constrained Optimization Problems", *CAM Student Seminar, The University of Chicago*

[1] May 2022, "Trust-Region Stochastic Sequential Quadratic Programming for Constrained Optimization", *ACMNTW Workshop on Optimization and Machine Learning 2022, Northwestern University*

Awards

Simons Dissertation Fellowship in Mathematics (research support, 40 recipients nationwide annually) 2026
Outstanding Undergraduate, Xiamen University 2021

Professional Service

Journal Reviewer: Mathematical Programming, PNAS

Conference Reviewer:

Association for the Advancement of Artificial Intelligence (AAAI)	2027
International Conference on Machine Learning (ICML)	2026
• Golden reviewer 2026	
International Conference on Learning Representations (ICLR)	2024–2027
Conference on Neural Information Processing Systems (NeurIPS)	2023–2026
International Conference on Artificial Intelligence and Statistics (AISTATS)	2023, 2024
Midwest Machine Learning Symposium	2023

Teaching:

GSI of Math 32 Pre-calculus, University of California, Berkeley	Spring 2025
GSI of Math 1A Calculus, University of California, Berkeley	Fall 2023 - Fall 2024
Grader of STAT 37710 Machine Learning, The University of Chicago	Fall 2022

Skills

Programming Languages: R, Python, CUDA, MATLAB, Julia, C, LaTeX