Fatemeh Nourzad

nourzad.1@osu.edu | atinourzad.github.io | fatemeh-nourzad | atinourzad | +1 (213) 354-3760

Education

2022 – present Ph.D. in Electrical and Computer Engineering (Expected graduation: May 2028)
The Ohio State University, Columbus, OH
Advisor: Prof. Ness B. Shroff

2019 – 2022 M.S. in Electrical and Computer Engineering
University of Tehran, Tehran, Iran
Thesis: Developed a robust **time-series** pipeline for seismic waveforms, automating **earthquake detection/classification** and **feature extraction** under distribution shift (varying noise/recording conditions).

2014 – 2018 B.S. in Electrical and Computer Engineering,
University of Tehran, Tehran, Iran

Selected Publications

- **Selective Forgetting in Tool-Augmented Language-Model Agents.**
  Fatemeh Nourzad[*], A. Roknilamouki[*], Q. Wu[*], E. Ekici, Y. Liang, N. B. Shroff.
  ***NeurIPS** 2026 TTCL Workshop; **CoLLAs** 2026 Work-In-Progress Track.* [paper]

- **Barriers to Pareto Steerability in Preference-Conditioned LLM Alignment.**
  Fatemeh Nourzad, D. Sow, Y. Liang, M. Shi, M. Zhang, Y. Li, E. Ekici, N. B. Shroff.
  ***ICLR** 2026 I Can't Believe It's Not Better Workshop.* [paper]

- **One Model, Many Goals: Meta-Learning Preference-Conditioned Alignment for Lifelong LLM Agents.**
  Fatemeh Nourzad, D. Sow, Y. Liang, M. Shi, M. Zhang, Y. Li, E. Ekici, N. B. Shroff.
  ***ICLR** 2026 Lifelong Agents Workshop.* [paper]

- **Democratizing LLM Training with the Edge: Challenges, Systems, and the Road Ahead.**
  Z. Luo, J. Cai[*], C. Le Denmat[*], S. Nair[*], **Fatemeh Nourzad**[*], et al.
  ***IEEE/ACM Transactions on Networking** (ToN)*, invited paper; publication forthcoming.

- **Democratizing LLM Training Across Geo-Distributed Heterogeneous Compute: A Systems Position and Roadmap.**
  Z. Luo, J. Cai[*], C. Le Denmat[*], S. Nair[*], **Fatemeh Nourzad**[*], et al.
  ***NeurIPS** 2026 CODE-FM Workshop.* [paper]
  Related extended abstract: **Toward WAN-Aware LLM Training Across Heterogeneous, Geo-Distributed Sites.**
  ***ACM SIGCOMM** 2026 NAIC Workshop*, **Oral Presentation (only extended abstract selected)**. [paper]

- **Provably Efficient RL for Linear MDPs under Instantaneous Safety Constraints in Non-Convex Feature Spaces.**
  A. Roknilamouki, A. Ghosh, M. Shi, **Fatemeh Nourzad**, E. Ekici, N. B. Shroff.
  *International Conference on Machine Learning (**ICML**)*, 2025. [paper]

- **Rethinking Multi-objective Resolution in Federated LLM Alignment.**
  Fatemeh Nourzad, A. Roknilamouki, H. Yang, E. Ekici, J. Liu, N. B. Shroff.
  Manuscript under review, 2026.

- **Thinking Beyond Tools: How Tool-Use RL Reshapes Deliberation.**
  Fatemeh Nourzad[*], A. Roknilamouki[*], Q. Wu[*], E. Ekici, Y. Liang, N. B. Shroff.
  Manuscript under review, 2026.

- **Learning to Steer the Pareto Frontier: Meta-Learning for Multi-Objective LLM Alignment.**
  Fatemeh Nourzad, D. Sow, Y. Liang, M. Zhang, Y. Li, E. Ekici, N. B. Shroff.
  Manuscript under review, 2026.

[*]Co-primary authors with equal contribution in alphabetical order.

Research Experience

Research Assistant, The Ohio State University 2024–Present

- **Tool-Use and Reasoning Retention in LLM Agents:** Investigated how tool-use RL reshapes reasoning, revealing consistent reasoning compression despite task-dependent accuracy changes; **developed methods to help retain reasoning capabilities** while learning to use tools, and characterized trade-offs among reasoning accuracy, deliberation length, and tool-use performance.
- **Preference-Conditioned LLM Alignment:** Developed a meta-learning framework for a single controllable policy that supports **continuous inference-time steering without retraining**; demonstrated improved Pareto coverage and generalization to unseen preferences over strong baselines. *Collaboration with researchers at Google.*
- **Pareto Steerability:** Identified **optimization interference and geometric limitations** that hinder continuous preference steering; conducted systematic comparisons to characterize alignment failure modes. *Collaboration with researchers at Google.*
- **Federated Multi-Objective Alignment:** Developed a communication-efficient federated RL framework using regularized client-side optimization; mitigated cross-client objective disagreement, reduced objective-gradient communication overhead, and established **finite-time Pareto-stationarity bounds**, with validation across heterogeneous clients and LLMs.

Research Assistant, AI-EDGE 2023–2026

- **Geo-Distributed Foundation-Model Training:** Co-developed a WAN-aware framework for pretraining and federated fine-tuning across heterogeneous, cross-institution GPU clusters; integrated communication-efficient synchronization and fault-tolerant recovery. Achieved up to **4× smaller update payloads** and reached target training loss **15.2% faster** than the best evaluated baseline in cross-institution pretraining.

Technical Skills

Programming: Python, C/C++

ML/LLM Frameworks: PyTorch, Hugging Face (Transformers, TRL, PEFT, Accelerate), Weights & Biases

LLM Alignment & RL: RLHF, PPO, GRPO, DPO, preference-conditioned alignment, tool-use RL

Optimization & Learning: Multi-objective optimization, Pareto optimization, meta-learning, federated learning

Distributed Training & Systems: LoRA, FSDP, DeepSpeed, SLURM, CUDA, gRPC

Infrastructure & Tools: Linux, Git, Docker, AWS, Google Cloud

Professional Service & Leadership

Peer Reviewer: NeurIPS, NeurIPS Workshops, ICLR, AISTATS, NAACL, IEEE/ACM ToN, ICDCS 2022–Present
Co-Organizer: Experts in AI and Networking, AI-EDGE Institute 2023–Present
Panelist: AI-EDGE Summer REU Program 2025
Teaching Assistant: University of Tehran (Machine Learning, Reinforcement Learning, Stochastic Processes) 2017–2022
Presenter: AIMACCS Workshop, The Ohio State University June 2026

Honors & Awards

Oral Presenter: ACM SIGCOMM NAIC Workshop; **only extended abstract selected for an oral presentation** 2026
Selected Presenter: Kraus Memorial ECE Graduate Poster Competition 2025
Ranked 2nd in M.S. graduating class, University of Tehran 2022
National Entrance Exam Rankings: Top 0.4% (M.S., 150k+ candidates); Top 0.2% (B.S., 500k+ candidates) 2014–2019