Surya Murthy

The University of Texas at Austin, 2515 Speedway, Austin, TX, 78712 |(215) 964-0920 | surya.murthy@utexas.edu

Summary

PhD researcher at The University of Texas at Austin specializing in multi-agent systems and multi-task optimization. Experienced in developing algorithms that balance competing objectives, integrate human feedback, and can be deployed to physics-based simulation environments for robotic manipulation, as well as preference-based optimization settings. Proficient in Python, C++, PyTorch, CUDA, ROS, and MuJoCo, with publications at NeurIPS and AIAA SciTech.

Education

The University of Texas at Austin — Expected Graduation: May 2028
Ph.D in Computer Engineering — GPA: 3.85/4.00

The University of Illinois at Urbana-Champaign — May 2023
Bachelor of Science in Computer Engineering, James Scholar, Dean's List — GPA: 3.88/4.00

Experience

Autonomous Systems Group — Austin, TX
Graduate Research Assistant — September 2023- Present

Multi-Agent Reinforcement Learning for Air Traffic

- Built large-scale simulation environments in the BlueSky simulator to model multi-agent air traffic scenarios.
- Applied reinforcement learning methods to induce real-time coordination among agents and explore trade-offs among agent-level objectives, focusing on scalability and robustness.
- Collaborated with NASA and MIT Lincoln Labs on demonstrations and data collection. Presented results at AIAA SciTech.

Multi-Objective Optimization with Minimal Feedback

- Developed a cooperative bargaining framework to find fair solutions in multi-objective settings, with applications in multitask reinforcement learning.
- Applied DiBS to the MetaWorld MT10 benchmark of ten robotic manipulation tasks using a shared Soft Actor–Critic policy, achieving strong performance and robustness to non-affine reward transformations compared to leading baselines.
- Co-authored a NeurIPS 2025 publication on convergence and transformation invariance, and contributed to an extension on deep learning and reinforcement learning applications submitted to ICLR 2026.

Estimating Preferences in Human–Autonomy Interaction

- Developed a comparison-based optimization algorithm to estimate user preferences and guide decision-making under multiple objectives. Integrating user feedback to iteratively refine preference estimates and update optimization outputs.
- Proved theoretical convergence of the algorithm to Pareto-optimal solutions.
- Validated the approach in a user study, outperforming automated negotiation baselines in interactions with real participants.

Micron Technologies Inc. — Boise, ID
Intern - Process Control Engineer — May 2022- August 2022

- Analyzed large-scale chip fabrication datasets with SQL and machine learning to identify high-failure process steps.
- Built a predictive ML pipeline to estimate process checkpoint measurements from sensor data, improving efficiency.

Selected Publications - Scholar Link

A Reinforcement Learning Approach to Safe and Quiet Urban Air Mobility Traffic Management — AIAA SciTech 2025

Cooperative Bargaining Games Without Utilities: Mediated Solutions from Direction Oracles — NeurIPS 2025

DiBS-MTL: Transformation-Invariant Multitask Learning with Direction Oracles — ICLR 2026 (under review)

Skills

Programming: Python, PyTorch, ROS, Mujoco Simulator, Docker, SQL, C, C++, CUDA, GitHub, MATLAB, SystemVerilog