

Yuhao Li

✉ yuhaoli@umass.edu 🔗 chkxw

Education

University of Massachusetts Amherst

Ph.D. in Computer Science Advisor: Prof. Hao Zhang

Amherst, MA

Sept 2026 – Expected May 2031

University of Wisconsin–Madison

B.S. in Computer Science

Madison, WI

Jan 2023 – Aug 2025

University of Michigan – Shanghai Jiao Tong University Joint Institute

B.E. in Electrical and Computer Engineering

Shanghai, China

Sept 2020 – Aug 2026

Publications

Conference and Journal Papers

Adapting Execution-Time Objectives for Multi-Robot Policies via Collaborative Flow Policy Guidance

Williard Joshua Jose, [Yuhao Li](#), Zihao Deng, Hao Zhang

Robotics: Science and Systems (RSS), 2026 [\[Paper\]](#) [🔗](#)

Abstract Sim2Real through Approximate Information States

Yunfu Deng, [Yuhao Li](#), Josiah P. Hanna

IEEE Robotics and Automation Letters (RA-L), 2026; presented at *IROS 2026* [\[Paper\]](#) [🔗](#) [\[arXiv\]](#) [🔗](#)

Reinforcement Learning Within the Classical Robotics Stack: A Case Study in Robot Soccer

Adam Labiosa, Zhihan Wang, Siddhant Agarwal, William Cong, Geethika Hemkumar, Abhinav Narayan Harish, Benjamin Hong, Josh Kelle, Chen Li, [Yuhao Li](#), Zisen Shao, Peter Stone, Josiah P. Hanna

IEEE International Conference on Robotics and Automation (ICRA), 2025 [\[Paper\]](#) [🔗](#)

Under Review

SF-Push: Safe and Fair Pose-to-Pose Collaborative Transport with Multiple Quadruped Robots

[Yuhao Li](#), Williard Joshua Jose, Hao Zhang

Under review at *ICRA 2027* [\[Project Page\]](#) [🔗](#)

Workshop Papers

Safe and Fair Multi-Quadruped Collaborative Box Pushing

[Yuhao Li](#), Williard Joshua Jose, Hao Zhang

ICRA 2026 Workshop on Multi-Agent Robotic Systems [\[Paper\]](#) [🔗](#) [\[Project Page\]](#) [🔗](#)

Controlling Generative Flow Policies of Multi-Robot Systems with Gradient-Based Guidance

Williard Joshua Jose, [Yuhao Li](#), Zihao Deng, Hao Zhang

ICRA 2026 Workshop on Multi-Agent Robotic Systems; RSS 2026 Workshop on Diffusion for Robot Learning [ICRA-W] [🔗](#) [\[RSS-W\]](#) [🔗](#)

Research Experience

Human-Centered Robotics Lab, University of Massachusetts Amherst

Research Assistant (Ph.D. student since Sept 2026), Advisor: **Prof. Hao Zhang**

Amherst, MA

June 2025 – Present

- **Safe and fair multi-quadruped transport (first author, SF-Push).** Safety and fairness constraints conflict with task progress, and a single scalar reward fixes that trade-off in advance. SF-Push keeps one policy gradient per robot and per objective, takes its update direction from a scale-free minimum-norm combination and its step length from the task gradient alone, so no trade-off weights need tuning. Deployed on real quadruped robots.
- **Mixed-motive box pushing (first author, ICRA 2026 workshop).** Rewarding task progress alone makes a quadruped pair either crowd into unsafe spacing or leave one robot idle; added team-level proximity and fairness terms as potential-based reward shaping on top of **MAPPO** to fix both.

- **MARL infrastructure.** Built a vectorized training framework that connects GPU-parallel simulators (**Isaac**, **Genesis**) with RL/MARL libraries (Stable-Baselines3, RSL-RL, OpenRL, HARL), including a vectorized PettingZoo-style interface with multi-agent autoreset; used it to train all policies and baselines in the projects above.
- **Real-robot experiments.** Integrated **ROS 2** with an **OptiTrack** motion-capture system and deployed learned policies on AgileX LIMO robots for a decentralized multi-robot fire-suppression project; co-authored the collaborative flow-policy guidance work (RSS 2026).

Prediction and Action Lab, University of Wisconsin–Madison

Madison, WI

Undergraduate Researcher, Advisors: **Prof. Josiah Hanna**, **Prof. Peter Stone**

May 2024 – June 2025

- **Abstract sim-to-real (RA-L 2026).** Implemented robot-side code to collect real-world transitions and augmented the data for coverage; trained **PPO** navigation policies in the learned latent space.
- **RoboCup SPL.** Built the near-goal dribble-to-goal policy with imitation learning (**DAgger**, **BCO**, **GAIL**); built a multi-agent robot-soccer environment on the **PettingZoo** AEC API with a customized SimRobot backend for self-play. The team won the **2024 RoboCup SPL Challenge Shield Division** (ICRA 2025 paper).

Shanghai Jiao Tong University

Shanghai, China

Undergraduate Researcher, Advisor: **Prof. Paul Weng**

May 2022 – Dec 2022

- Trained **DQN** policies to accelerate Traveling Salesman Problem solving; modified the Stable-Baselines3 replay buffer to handle TSP graphs of varying size.

Honors and Awards

- RoboCup SPL Challenge Shield Division Champion (team), 2024
- UW–Madison Undergraduate Scholarship for Summer Study, Summer 2024
- Dean’s List, UW–Madison, Spring 2023 and Spring 2024

Open-Source Software

- **TypedShmem:** a heap in POSIX shared memory (atomics + locks) as a C++ library with pybind11 Python bindings, for fast C++/Python runtime communication. [\[GitHub\]](#) [🔗](#)

Skills

Reinforcement Learning: MAPPO, PPO, DQN, imitation learning (DAgger, BCO, GAIL); Stable-Baselines3, RSL-RL, OpenRL, HARL, PettingZoo, Gymnasium

Robotics and Simulation: ROS 1/2, OptiTrack motion capture, Isaac, Genesis

Programming: Python, C/C++, PyTorch, L^AT_EX