Fernando Palafox

PhD Student in Robotics, UT Austin

fernandopalafox@utexas.edu

palafox.info

Summary

My research focuses on developing general agents that can efficiently adapt to new dynamics, environments, and tasks. I am specifically interested in how pretraining recipes can make adaptation faster and more parameter-efficient, and how active exploration can improve sample efficiency. I am currently working on:

- Real-world continual learning for humanoid loco-manipulation tasks by adapting world models distilled from simulation.
- Tactile sensing for contact-rich humanoid loco-manipulation tasks.
- RL fine-tuning of LLMs that keep humans in the loop of automated scientific discovery, supported by the **Cosmos Institute**.

I am seeking research scientist roles in robotics starting summer 2027.

Education

PhD in Robotics 2022 - Present

University of Texas at Austin

MS in Aerospace Engineering 2021 - 2022

University of Colorado Boulder

BS in Aerospace Engineering 2017 - 2021

University of Colorado Boulder

Experience

Graduate Research Assistant at UT Austin 2022 - Present

- Research on robotic adaptation to new and uncertain environments. Designed pretraining recipes for faster adaptation [1, 2] and information-gathering strategies for sample-efficient learning [3]. Applied these methods to world models for model-based reinforcement learning.
- Worked on game theory [5, 6, 7] and numerical optimization [8] for multi-agent planning and decision-making.

Research Intern at Honda Research Institute Summer 2024

- Research on active learning algorithms to help autonomous vehicles reduce uncertainty about driver behavior during highway merges.

Research Assistant at CU Boulder 2020 - 2022

- Research on open-source software radio receivers for GPS, Galileo, and WAAS satellite systems, and low-cost local positioning systems [9, 10].

Technical Skills

Real Robot Experience: Quadrupeds and humanoids
Languages: Python, Julia, MATLAB, C++
Learning Frameworks: JAX and PyTorch
Libraries/Tools: MuJoCo, Git, ROS
Topics: Machine Learning, Reinforcement Learning, Numerical Optimization, Statistical Estimation, Information Theory

Publications

Google Scholar

1. **Palafox, F.**, & Fridovich-Keil, D. (2026). Amortized Low-Rank Adaptation for Model-Based Reinforcement Learning. CoRL 2026 Workshop: Pretrain to Adapt.

2. Levy, J., Westenbroek, T., Huang, K., **Palafox, F.**, Yin, P., Omidshafiei, S., Kim, D.-K., Gupta, A., & Fridovich-Keil, D. (2026). Simulation Distillation: Pretraining World Models in Simulation for Rapid Real-World Adaptation. RSS 2026.

3. **Palafox, F.**, Li, J., Milzman, J., & Fridovich-Keil, D. (2026). Generalized Information Gathering Under Dynamics Uncertainty. Under review.

4. Ly, N., Tatsuoka, C., Nagaraj, J., Levy, J., **Palafox, F.**, Fridovich-Keil, D., & Lu, H. (2026). Data-Driven Modeling and Correction of Vehicle Dynamics. JMLMC 2026.

5. Qiu, T., Ouano, E., **Palafox, F.**, Ellis, C., & Fridovich-Keil, D. (2025). A Player Selection Network for Scalable Game-Theoretic Prediction and Planning. TMLR 2025.

6. **Palafox, F.**, Milzman, J., Lee, D. H., Park, R., & Fridovich-Keil, D. (2025). Smooth Information Gathering in Two-Player Noncooperative Games. AAMAS 2025.

7. **Palafox, F.**, Yu, Y., & Fridovich-Keil, D. (2024). Learning Hyperplanes for Multi-Agent Collision Avoidance in Space. AAS/AIAA 2024.

8. Li, J., Chiu, C.-Y., Peters, L., **Palafox, F.**, Karabag, M., Alonso-Mora, J., Sojoudi, S., Tomlin, C., & Fridovich-Keil, D. (2023). Scenario-Game ADMM: A Parallelized Scenario-Based Solver for Stochastic Noncooperative Games. CDC 2023.

9. Bernabeu, J., **Palafox, F.**, Li, Y., & Akos, D. M. (2022). A Collection of SDRs for Global Navigation Satellite Systems (GNSS). ION ITM 2022.

10. **Palafox, F.**, Ruane, L., Palo, S., & Akos, D. (2021). Developing a Low-Cost, High Performance, SDR-Based Local Positioning System. ION GNSS+ 2021.

Service and Outreach

Research Mentorship

Current

- Jai Nagaraj. Undergraduate at UT Austin

Past

- Nguyen Ly. Data-Driven Modelling of Nonautonomous Vehicle Dynamics. Undergraduate at UT Austin.
- Ryan Chin. Multi-Agent Collision Avoidance in Space. Undergraduate at MIT.
- Ryan J. Park. Smooth Information Gathering in Two-Player Noncooperative Games. PhD student at UC Berkeley.

Reviewer for CoRL, NeurIPS, WAFR, TAC, IROS, ICRA, CDC, RA-L, L-CSS

Honors, Awards, and Fellowships

- Cosmos Institute Research Grant, 2026. Awarded for Training LLMs for Human Insight.
- Graduate School Fellowship, University of Texas at Austin, 2022-2026
- UROP Assistantship, University of Colorado Boulder. Undergraduate research funding for work on low-cost, high-performance positioning systems. 2021
- Best Technical Leadership in Electronics, University of Colorado Boulder, Aerospace Engineering senior capstone projects 2021.
- Engineering Merit Scholarship, University of Colorado Boulder, 2018-2019
- Dean's List, University of Colorado Boulder, F 2017, S/F 2018, S/F 2019, S/F 2020, S/F 2021

Talks

- April 2026. Invited talk on Generalized Information Gathering Under Dynamics Uncertainty at the Intelligent Control (ICON) Lab, UC Berkeley, hosted by Negar Mehr.
- Feb. 2024. Presented Scenario-Game ADMM: A Parallelized Scenario-Based Solver for Stochastic Noncooperative Games at the Texas Robotics Graduate Portfolio Seminar.

Professional Society Memberships

- Tau Beta Pi, Engineering Honor Society
- Sigma Gamma Tau, Aerospace Engineering Honor Society
- Society of Hispanic Professional Engineers