Anthony Roumi

+1 (310) 850-1195 | tonyroumi@gmail.com | Linkedin | https://github.com/tonyroumi

Personal Statement: Software Engineer and Algorithm Developer with research engineering experience in distributed systems, embodied AI, and multi-agent control. Led technical development on a Phase III SBIR data-fusion engine and directed a student robotics project that trained a reinforcement-learning locomotion policy for a humanoid robot. Independently conducting machine learning research for physical systems across simulation engines, with the goal of building the probabilistic, statistical, and practical intuition needed to pursue meaningful research projects from first principles. Seeking to apply this blend of algorithmic, systems, and ML expertise to advance autonomous control solutions.

Experience

Software Engineer/Algorithm Developer| Lyntris | *San Diego, CA* July 2025 - Present

- Lead technical and software engineering efforts on a Phase III SBIR program, defining architecture, developing features and algorithms, and delivering the project on schedule.
- Contributing to a distributed data fusion engine in C++, integrating tracking capabilities into customer facing SDKs across Linux and containerized environments.
- Integrate and debug EKF pipelines for air-vehicle state estimation from azimuth and elevation sensors, using multi-hypothesis tracking to associate measurements, which increases estimation reliability.
- Designed and implemented a distributed consensus and task-reassignment algorithm that lets multi-platform teams keep shared tasking and operational state even under real radio-communication conditions.
- Developed a portable containerized demo, including configurable scenarios, live data collection pipelines, and custom AFSIM plugins for customer demonstrations and internal collaboration.

UCSD Triton Droids - *Software / Controls Team Lead* Nov. 2024 - June 2025

- Led the robotics team to develop reinforcement learning based pipelines for locomotion of a humanoid, coordinating simulation (MJX, Python) and experimental testing across subteams, which produced a functional walking controller used in the organization's demo.
- Trained a joystick controlled locomotion policy in MJX for a simplified humanoid model using PyTorch achieving stable simulated walking that served as the baseline controller for subsequent research.
- Developed onboarding curriculum covering simulation, RL, and embodied control to ramp new members onto active research and development.

Machine Learning Research Intern | ObjectSecurity | *San Diego, CA* July 2024 - Oct. 2024

- Designed and led a research project investigating neuron interpretability in transformer models by analyzing activation patterns across layers for large labeled datasets of "high confidence" and "low confidence" predictions.
- Analyzed correlated activity between neurons, input features, and model outputs to identify specific neurons strongly associated with model confidence.
- Performed neuron ablation studies to measure contribution of identified neurons, demonstrating their influence on model outputs.

Projects

Motion Retargeting and Reinforcement Learning for Embodied Character Control | *Python, Jax, Pytorch, MuJoCo, MJX*

- Independently reproduced Adversarial Motion Priors (AMP) in MuJoCo from the original paper to train physically plausible stylistic locomotion policies satisfying task objectives.
- Implemented a skeleton aware motion retargeting network, enabling motion retargeting between topologically equivalent skeletons without paired supervision.
- Conducting ongoing R&D on a motion retargeting architecture leveraging large scale online mocap datasets, extending skeleton aware representation learning toward scalable cross embodiment transfer of mocap data.

Multi-Agent Reinforcement Learning Benchmark (Robosuite) | *Python, Pytorch, MuJoCo*

- Designed a MARL benchmark in robosuite spanning cooperative manipulation tasks such as two arm lifting, peg in hole insertion, and object handover to evaluate algorithm performance.
- Implemented and trained MAPPO and MASAC baselines, comparing on policy vs. off policy performance, coordination, stability, and sample efficiency across tasks.
- Built training and evaluation pipelines with standardized metrics: task success rate, learning stability, and agent coordination.

Skills

ML/RL: PPO, SAC, AMP, Imitation Learning, Representation Learning, World Models, Stochastic Policy Optimization
Controls/Estimation: Kalman filters, Bayesian filtering, state space models, PID control, system dynamics
Frameworks Programming & Tools: MuJoCo, MJX, IsaacSim, C++, Python, Pytorch, Scipy, Scikit-learn, Pandas, Numpy, Jax, Docker

Education

University of California – San Diego
B.S. in Computer Engineering (Provost Honors Student)