

Amal Mehta

SF Bay Area · amal.mehta@gmail.com · amalmehta.com · LinkedIn · GitHub

Research engineer with six years of applied deep learning, from generative models and vision to deep reinforcement learning, taking work from data and training to policies running on real robots. Earlier neuroimaging research at UCSF, and a long-standing interest in how the brain learns.

Experience

Independent ML & Robotics Research

Self-directed projects; details at amalmehta.com

Sept 2025 – Present

San Francisco Bay Area

- **Reward hacking in code agents:** GRPO on visible tests taught an agent to delete 98% of untested branches (true correctness 31% → 2% as pass rate rose); random 25% hidden-check audits stopped it in the pilot, but 50 pre-registered follow-ups found they only delay it.
- **Brain-computer interfaces:** four projects on public neural data; label-free decoding averaged R^2 0.72 to day 138 (Wiener 0.42).
- **Post-training:** SFT then GRPO on a small transformer policy in a key-and-door gridworld maze; over 3 seeds, GRPO lifted held-out success from 70% to 80% (DPO: 81%), where RL from scratch failed.
- **Meta-RL and self-play:** quadrotor adapting to payload, wind and motor wear with seven meta-learning methods; 9×9 Go by self-play.
- **Tools:** LLM Understander, an interactive course on how LLMs are trained; Researcher Pal, Claude Code skills for research papers.
- Member, Sensorimotor AI journal club, Redwood Center (UC Berkeley); mentored undergraduate scholars from Berkeley.

Deep Learning Scientist

PhysicsAI (DARPA-funded)

Sept 2021 – Aug 2025

Pacifica, CA (remote)

- Formulated synthetic-data generation as RL: policies (TD3, PPO, DDPG) set image-generator parameters, rewarded by detector performance on real drone imagery, closing the sim-to-real gap without hand-tuned rendering.
- Built a PEARL meta-RL agent (probabilistic task embeddings over TD3) for fast few-shot adaptation across generator tasks.
- Implemented domain adaptation from scratch to hold detector performance in the low-real-label, high-synthetic-data regime.
- Built GAN-based synthesis pipelines (DCGAN, InfoGAN, FineGAN) producing training data for downstream detection models.

Deep RL Researcher

Hybrid Robotics Lab, UC Berkeley — adv. Prof. Koushil Sreenath

Aug 2023 – Mar 2024

Berkeley, CA

- Trained a single multi-stage PPO policy in IsaacGym for a Unitree Go1 quadruped to track, approach, and shoot a soccer ball at a target, replacing the multi-policy hierarchical framework used in prior work; ~60% kicking accuracy on hardware.
- Designed the three-stage reward and transition machinery and CPG-based gait tracking; prototyped RLPD (SAC with prior data) to fine-tune the sim-trained policy on real robot data.

Research Assistant

Video & Image Processing Lab, BAIR — adv. Prof. Avidesh Zakhor

June 2020 – Dec 2022

UC Berkeley

- Developed a deep CNN segmentation pipeline for invasive melanoma in large-scale histopathological images (F1 0.44 vs. 0.14 for the prior best); published at SPIE Medical Imaging 2023; Stowell-Orbison Award Poster Session at USCAP 2022.

Researcher

Abbasi-Asl Lab, UCSF — adv. Prof. Reza Abbasi-Asl

June 2020 – Aug 2020

San Francisco, CA

- Built empirical mode decomposition (EMD) pipelines for 2D and 3D brain image registration.

Education

University of California, Berkeley

B.A. Computer Science — Regents Scholar

Aug 2017 – Aug 2021

Berkeley, CA

- Focused on robotics and linear system theory; Regents, merit, and leadership awards.
- CS285 (Deep RL) project: iterative retraining of PPO agents against adversarial policies.

Stanford University

Graduate coursework (non-degree) — Autonomous Systems & Artificial Intelligence

Aug 2021 – May 2023

- CS231N (Deep Learning for Vision), CS330 (Meta-Learning), AA271A/B (Autonomy I & II), AA273 (State Estimation).
- CS330 project: few-shot imitation learning; implemented RL^2 and inverse soft Q-learning in Meta-World.

Technical Skills

Languages

Python, C++, Swift, JavaScript / TypeScript

ML / RL

PyTorch, TensorFlow, stable-baselines3, IsaacGym, OpenCV, CTC decoding

Product

SwiftUI, Three.js

Robotics

ROS, Gazebo, Unitree Go1, CPG gait control, Kalman filtering, Shared autonomy

Tools

Docker, GitHub Actions CI, CUDA/multi-GPU, W&B

Publications

1. Shah, A., **Mehta, A.**, Wang, M., Neumann, N., McCalmont, T., Zakhor, A. “Deep Learning Segmentation for Invasive Melanoma.” *SPIE Medical Imaging*, 2023.
2. Neumann, N., Wang, M., **Mehta, A.**, et al. “Quantifying Invasive Melanoma Volume by Deep Learning Segmentation at the Pixel-Level.” *USCAP Annual Meeting*, 2022.