Anthony Liang

aliang8.github.io ◇ anthony.liang@tri.global ◇ linkedin.com/in/ant-liang

RESEARCH INTERESTS

My research focuses on scalable robot learning, with an emphasis on data-efficient human feedback (gaze, demonstrations, preferences, interventions) and reward foundation models that provide generalizable reward functions across tasks, embodiments, and environments. I also work on core components for embodied intelligence, including vision-language-action models for low-level control, trace models for structured planning, and world models for scene understanding and predictive reasoning. In parallel, I am investigating RL post-training for LLM-based web agents to improve reliability in sequential UI tasks and long-horizon interactions.

EDUCATION

University of Southern California, Los Angeles, CA — Aug 2021 - May 2026
Ph.D. in Computer Science (Co-advised by Erdem Bıyık and Stephen Tu) — *GPA: 3.9/4.0*

University of Michigan Rackham Graduate School, Ann Arbor, MI — Aug 2017 - 2021
M.S. in Robotics (Advisor: Honglak Lee) — *GPA: 4.0/4.0*

University of Michigan, Ann Arbor, MI — Aug 2017 - 2021
B.S.E. (Advisor: Honglak Lee) — *GPA: 3.65/4.0*

RESEARCH EXPERIENCE

LiRA Lab and Statistical Learning Lab, USC — Aug 2021 - May 2026
Ph.D. Student, PI: Erdem Bıyık, Stephen Tu

- Robot learning with multimodal human feedback, sample-efficient RL, learning from unlabelled videos

Intelligent Robot Lab, Carnegie Mellon — May 2020 - May 2021
Visiting Researcher with Changliu Liu

- Hierarchical RL for safe control of autonomous vehicles in dynamic environments

Deep Learning Lab, University of Michigan — Jan 2019 - May 2021
Research Intern with Honglak Lee

- Sample-efficient RL for embodied task learning

PROFESSIONAL EXPERIENCE

Toyota Research Institute — May 2026 - Present
Research Scientist — *Los Altos, CA*

- Post-training and sample-efficient real-world RL, policy improvement in World-Action Models, learning from suboptimal data

Google DeepMind — May 2025 - Sept 2025
Research Intern with Kalpesh Krishna, Jacob Eisenstein — *NYC/Seattle*

- RL post-training for reasoning models

Google Research — May 2023 - Feb 2024
Research Intern with Chih-Wei Hsu, Yinlam Chow, Guy Tennenholtz, Craig Boutilier — *Remote NYC*

- Bayesian RL for Markov Decision Processes with gradually changing latent dynamics
- Data augmentation at critical states for robot imitation learning with Stephen Tu
- Generative modeling for online RL policies

Meta AI - Multimodal Learning Team — May 2022 - Aug 2022
Research Intern with Paul Crook and Andrea Madatto — *Redmond, WA*

- Fine-tuning large language models for task-oriented dialogue generation

Amazon Science — May 2021 - Aug 2021
Applied Science Intern with Thiago Mosquiero — *Seattle, WA*

- Collaborative filtering for recommending new brands and products to consumers

Invisible.ai May 2020 - Aug 2020
AI Research Intern *Remote*

- Improving computer vision models for real-time object detection and tracking for industrial processes

Google Ads May 2019 - Aug 2019
Software Engineering Intern *Mountain View, CA*

Luminar Technologies May 2018 - Aug 2018
Software Engineering Intern *Palo Alto, CA*

PUBLICATIONS

[C7] Anthony Liang*, Yigit Korkmaz*, Jiahui Zhang, Minyoung Hwang, Abrar Anwar, Sid Kaushik, Aditya Shah, Alex Huang, Luke Zettlemoyer, Dieter Fox, Yu Xiang, Anqi Li, Andreea Bobu, Abhishek Gupta, Stephen Tu†, Erdem Bıyık†, Jesse Zhang†. "Robometer: Scaling General-Purpose Robotic Reward Models via Trajectory Comparisons". *Robotics: Science and Systems (RSS) 2026*

[C6] Anthony Liang, Jonathan Berant, Adam Fisch, Abhimanyu Goyal, Kalpesh Krishna†, Jacob Eisenstein†. "Plantain: Plan-Answer Interleaved Reasoning". *International Conference on Machine Learning (ICML) 2026 (Spotlight)*

[C5] Matthew Hong*, Anthony Liang*, Kevin Kim, Harshitha Rajaprakash, Jesse Thomason†, Erdem Bıyık†, Jesse Zhang†. "HAND Me the Data: Fast Robot Adaptation via Hand Path Retrieval" *International Conference on Robot Automation (ICRA) 2026*

[C4] Anthony Liang*, Pavel Czempin*, Matthew Hong, Yutai Zhou, Erdem Bıyık, Stephen Tu. "CLAM: Continuous Latent Action Models for Robot Learning from Unlabeled Demonstrations" *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2026*

[C3] Anthony Liang, Chih-wei Hsu, Yinlam Chow, Guy Tennenholtz, Erdem Bıyık, Craig Boutilier. "DynaMITE-RL: A Dynamics Model for Improved Temporal Meta Reinforcement Learning", *International Conference on Machine Learning (ICML) AutoRL Workshop 2024, Conference on Neural Information Processing Systems (NeurIPS) 2024*

[C2] Anthony Liang, Jesse Thomason, Erdem Bıyık. "ViSaRL: Visual Reinforcement Learning Guided By Human Saliency", *Spotlight talk at ICRA Pretraining for Robotics (PT4R) Workshop, IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) 2024*

[C1] Wilka Carvalho, Anthony Liang, Kimin Lee, Sungryull Sohn, Honglak Lee, Richard L. Lewis, Satinder Singh. "Reinforcement Learning for Sparse-Reward Object-Interaction Tasks in a First-person Simulated 3D Environment", *International Joint Conferences on Artificial Intelligence (IJCAI) 2021*

[T5] Anthony Liang, Yigit Korkmaz, Jiahui Zhang, Jesse Zhang, Abrar Anwar, Sidhant Kaushik, Yufei Wang, Yu Xiang, David Held, Dieter Fox, Abhishek Gupta, Stephen Tu, Erdem Biyik. "SPUR: Scaling Reward Learning from Human Demonstrations". *CoRL Eval&Deploy Workshop 2025*

[T4] Anthony Liang, Pavel Czempin, Yutai Zhou, Stephen Tu, Erdem Bıyık. "In-Context Generalization to New Tasks From Unlabeled Observation Data", *ICML In-Context Learning Workshop 2024*

[T3] Ishika Singh, Anthony Liang, Mohit Shridhar, Jesse Thomason. "Self-Supervised 3D Representation Learning for Robotics", *ICRA Pretraining for Robotics (PT4R) 2023*

[T2] Anthony Liang, Ishika Singh, Karl Pertsch, Jesse Thomason. "Transformer Adapters for Robot Learning", *CoRL Workshop on Pretraining for Robot Learning 2022*

[T1] Wilka Carvalho, Anthony Liang, Kimin Lee, Sungryull Sohn, Richard L. Lewis, Satinder Singh, Honglak Lee. "ROMA: A Relational, Object-Model Learning Agent for Sample-Efficient Reinforcement Learning", *ICML Workshop on Object-Oriented Learning 2020*

TEACHING

University of Southern California

CSCI 445: Foundation of Robotics Fall 2026
CSCI 699: Robot Learning Fall 2024
CSCI 499: Natural Language for Interactive AI Fall 2022

University of Michigan, Ann Arbor

EECS 442: Computer Vision	Winter 2021
EECS 498: Algorithmic Robotics	Fall 2020
EECS 504: Graduate Computer Vision	Winter 2020
EECS 280: Introduction to Programming and Data Structures	Fall 2018 - Fall 2019

Summer STEM Institute Research Mentor — Summer 2021

SERVICES

Reviewer:

- ICML, ICLR, ICRA, NeurIPS, HRI, RA-L, AAAI, RO-MAN, CVPR

Workshop Organizer:

- Structured Physical Intelligence (SPIN) at CoRL 2026
- Human-in-the-Loop Robot Learning: Teaching, Correcting, and Adapting at RSS 2025

Mentorship and Outreach:

- USC UGrad Mentoring Program
- CURVE: Mentored USC undergrads through fellowship program.

STUDENT MENTORSHIP

Toyota Research Institute

- Zheyuan Hu — Long-context AR VLAs
- Nico Espinosa Dice — Pretraining for post-training

USC Masters Students

- Matthew Hong — HAND
- Dhanush Kumar Penmetsa — Gaze for robot teleoperation
- Ziyi Liu — Preference-based RL

USC Undergraduate Students

- Yixi Quan (Undergrad, USC) — Pretrained Video-LLMs for Robot Learning
- Jaiv Doshi (Undergrad, USC) — Human-intervention reinforcement learning
- Junu Song (USC CURVE Fellowship) — Real-world robot navigation

Fellowship Programs and External Collaborators

- Mihir Rao (Undergrad, Berkeley) — Trace VLAs
- Sankalp (Sunny) Agrawal (Undergrad, USC SURE Program) — Meta-RL with task descriptors
- Shreya Ramanujam (Undergrad, IIT) — Gaze for robot teleoperation

INVITED TALKS

"Scaling Reward Models via Preference Comparisons" — 2026

- GRAIL Lab, NYU — (Hosts: Irmak Güzey, Lerrel Pinto)
- Apple Robotics Reading Group — (Host: Soheil Habibian)

HONORS AND AWARDS

- NSF Graduate Research Fellowship Honorable Mention — 2020