

Education

University of California, Berkeley

M.S. Electrical Engineering and Computer Sciences

B.A. Computer Science | Cumulative GPA: 3.85

Berkeley, CA

May 2026

May 2025

- **Undergraduate Coursework:** Foundations of LLMs, Intro to Machine Learning, Microfabrication, Computer Graphics, Convex Optimization, Computer Architecture, Robotic Manipulation and Interaction, Nanorobotics, Data Structures and Algorithms, Advanced Algorithms, Discrete Math and Probability Theory, Designing Information Devices, Foundations of Data Science
- **Graduate Coursework:** Deep Reinforcement Learning (Top 11 out of 220 students), Computer Vision

Experience

Nvidia

Software Engineer Intern

Santa Clara, CA

May 2025 – August 2025

- Working on mapping and vision transformers with autonomous vehicles

Berkeley Artificial Intelligence Research (BAIR) Lab

Undergraduate Researcher

Berkeley, CA

January 2022 – Present

- Conducting research in robotic grasping and manipulation and computer vision techniques in 3D reconstruction and tracking advised by Professor [Ken Goldberg](#); published papers in robotics conferences including CoRL, ICRA, IROS and CASE

Head Teaching Assistant for UC Berkeley Graduate Robotics Courses

January 2023 – Present

- Running discussion sections and hosting office hours for the introductory upper-division/graduate robotics class of **300 students** and the advanced graduate robotics course of **100 students**
- Developed an [entirely new robotics lab](#) for students that takes them through the entire process of building a simple robotic system using object segmentation, PID control, and trajectory planning with bezier curves

Research & Publications

R2R2R: Scaling Robotic Manipulation Data Without Robot Hardware - CoRL 2025 (Oral)

April 2025

- Real2Render2Real (R2R2R) is a scalable pipeline for generating data to train generalist manipulation policies without dynamics simulation or teleoperation
- R2R2R can generate thousands of robot trajectories in IsaacLab with just a smartphone scan of an object and a recording of a person manipulating the object
- Modern RGB + Proprioception based imitation learning frameworks and VLA models (Vanilla Diffusion Policy, π_0 -FAST) can then be trained on R2R2R data without requiring **any teleoperation data!**

Creating Digital Twins via Bimanual Robot Handoff and Gaussian Splat Merging - IROS 2025

March 2025

- A bimanual robot picks up and scans objects from multiple viewpoints, then it hands off the object to the other arm getting views of previously occluded points. We then merge the scans into a single, visually-accurate 3DGS model

Gaussian Splatting for Tracking Human and Robot Manipulation of Objects - ICRA 2025

September 2024

- Built a state-of-the-art computer-vision algorithm and robotic system for densely mapping and tracking objects in the real world for robots to work with humans on long-horizon tasks
- Implemented on a UR5 robot with a wrist-mounted depth camera and tested on sequential pick and place with **12** consecutive successes and 6DOF tool-in-gripper visual-servoing

Language Embedded Gaussian Splatting for Mobile Robotics - IROS 2024 (Oral)

June 2024

- Heavily modified power and network system of mobile robot with custom 305 degree FOV camera setup to autonomously traverse large-scale environments (workspaces, kitchens, cafeterias) and build a 3D map of the environment with language-embeddings in real-time
- We use the language embeddings in the map to query the map for the presence of a given item and if it is present the robot will traverse to the object and retrieve it

Storytelling With Action Guidance - EMNLP Findings 2024

June 2024

- Reframing the task of storytelling with large language models (LLMs) to a search problem through a two-model feedback loop: one LLM generates story content, and another, auxiliary LLM, is used to choose the next best “action” to steer the story’s future direction
- This approach produces stories that are **85% more preferred** than previous works and allows small open-source models, such as Llama-2-7B, to outperform GPT-3.5-Turbo

Onboard Tracking of 160+ MPH Autonomous Racecars Using Multi-Input EKF

May 2023

- Developed a computer vision pipeline in C++ with ROS2 to preprocess incoming LiDAR and RADAR readings on fully autonomous race cars traveling 160+ mph at rates of 20Hz and 5.5Hz respectively
- Fused these sensor measurements with an extended Kalman filter that produces readings on a car’s position, orientation and velocity at 200Hz with an average positional error of 17cm

Skills

- **Programming Languages:** Java, C, C++, Python, SQL, HTML, CSS, JavaScript, RISC-V
- **Tools & Libraries:** PyTorch, ROS, ROS2, OpenCV, Docker, Git, React, Tensorflow, IsaacLab