

Sean Kirmani

sean@kirmani.ai • kirmani.ai • San Francisco, CA

EXPERIENCE

Google Brain — Senior Research Engineer

Apr 2023 - Present

Machine learning research in vision, language, and robotics. Mountain View, California.

Everyday Robots — Technical Lead, Semantic Perception

Nov 2021 - Apr 2023

Everyday Robots spun out of X in November 2021. Mountain View, California.

- Vision-language model lead at Everyday Robots. Introduced and deployed first vision and language model (CLIP) in production for robot visual question answering. Scaled diffusion models to create synthetic data for CLIP. Landed on-robot open-vocabulary object detector to detect novel objects.
- Designed and built a multi-sensor (camera and lidar), open-vocabulary panoptic segmentation model.
- Full-stack ML engineer: end-to-end ownership of entire ML flywheel from data collection to inference. Built model automation pipeline for data collection, training, evaluation, and on-robot deployment for all perception models.

X (formerly Google[x]) — Senior Software Engineer, Everyday Robots

July 2018 - Nov 2021

Working on perception for robots. Mountain View, California.

- Early engineer on the perception team. Expert in bringing research to production in real world systems.
- Created the lidar panoptic segmentation model and RGB-D camera panoptic segmentation model (with associated automation flywheel) and deployed to robot fleet.
- Trained multimodal vision and action models, resulting in publication at ICRA.
- Filed 5 patents and published 2 papers.
- Built the first 3D object tracker.

X (formerly Google[x]) — Software Engineering Intern, Robotics

May 2017 - Aug 2017

Worked on perception for human-robot interaction. Part of the first X intern class. Converted to full-time. Mountain View, California.

Google VR/AR — Software Engineering Intern, Project Tango

May 2016 - Aug 2016

Worked on experimental augmented reality. Created environmental lighting system allowing more photorealistic lighting and reflections in augmented reality for Tango SDK. Google Developer Blog post available at kirmani.ai/go/tango. Experience in computer graphics, computational photography, and computer vision. Mountain View, California.

Google — Software Engineering Intern, Chrome for Android

May 2015 - Aug 2015

Intern on Chrome for Android. Wrote test infrastructure for sign-in authentication tests. Created parametrizable testing framework. All code is open source as part of Chromium. Viewable at kirmani.ai/go/chromium. Mountain View, California.

Accordion Health — Software Engineer

Aug 2014 - Jan 2015

Acquired by Evolent Health, Inc. in June 2017. Early stage startup to use machine learning for health care data analytics. Clustered co-morbidity for several sets of patients.

SKILLS & INTERESTS

Interests Computer Vision, Natural Language Processing, Robot Learning, Deep Reinforcement Learning
Familiar Technologies Tensorflow, JAX, Robot Operating System, OpenGL, Vim, Git, Android

EDUCATION

The University of Texas at Austin

Bachelor of Science, Computer Science

Bachelor of Science, Electrical and Computer Engineering (Robotics focus)

Honors Thesis: Deep Reinforcement Learning for Aerial Obstacle Avoidance Using Monocular RGB Images

PUBLICATIONS

Practical Imitation Learning in the Real World via Task Consistency Loss

Mohi Khansari, Daniel Ho, Yuqing Du, Armando Fuentes, Matthew Bennice, Nicolas Sievers, [Sean Kirmani](#), Yunfei Bai, Eric Jang

In the proceedings of the *International Conference on Robotics and Automation (ICRA)*, 2023.

PRISM: Pose Registration for Integrated Semantic Mapping

Justin Hart, Rishi Shah, [Sean Kirmani](#), Nick Walker, Kathryn Baldauf, Nathan John, Peter Stone

In the proceedings of the *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, 2018.

Passive Demonstrations of Light-Based Robot Signals for Improved Human Interpretability

Rolando Fernandez, Nathan John, [Sean Kirmani](#), Justin Hart, Jivko Sinapov, Peter Stone

In the proceedings of the *IEEE International Conference on Robot and Human Interactive Communication (RO-MAN)*, 2018.

PATENTS

Object Association Using Machine Learning Models

Patent Number: US20220388175A1

[Sean Kirmani](#), Guy Satat, Michael Quinlan

Fusing Multiple Depth Sensing Modalities

Patent Number: US20220366590A1

Guy Satat, Michael Quinlan, [Sean Kirmani](#), Anelia Angelova, Ariel Gordon

Label transfer between data from multiple sensors

Patent Number: US20220268939A1

Sarah Najmark, [Sean Kirmani](#)

Shared Dense Network with Robot Task-Specific Heads

Patent Number: US20210181716A1

Michael Quinlan, [Sean Kirmani](#)

Engagement detection and attention estimation for human-robot interaction

Patent Number: US11436869B1

[Sean Kirmani](#), Michael Quinlan, Sarah Coe

RESEARCH

Robotics Lab at UT Austin, Dr. Peter Stone and Dr. Justin Hart

Dec 2017 - Apr 2018

Research on semantic mapping and social navigation with the Building Wide Intelligence (BWI) lab. Austin, Texas.

Robotics Lab at UT Austin, Dr. Andrea Thomaz and Dr. Scott Niekum

Jan 2016 - May 2017

Research in human robot interaction in the Personal Autonomous Robotics Lab (PeARL) and Socially Intelligent Machines (SiM) Lab. Experience in behavior architectures, perception, planning, and machine learning. Austin, Texas.

Wireless Networking & Communication Group, Dr. Joydeep Ghosh

Aug 2014 - Jan 2016

Selected by Professor Joydeep Ghosh in the University of Texas Electrical and Computer Engineering department in the Intelligent Data Exploration and Analysis Laboratory (IDEAL). Lab focuses on machine learning and data mining. Research on making self-driving cars a safe reality using distributed machine learning through wireless communication. Austin, Texas.

EXTENDED RESUME AT KIRMANI.AI/RESUME
