## Overview

I am a research scientist at NVIDIA AI. I obtained my PhD degree from Stanford University, co-advised by Silvio Savarese and Fei-Fei Li. My research focuses on developing systems and algorithms to enable robots to learn manipulation tasks.

## Selected Publications (see website for full list)

- NOD-TAMP: Generalizable Long-Horizon Planning with Neural Object Descriptors Shuo Cheng, Caelan Garrett\*, Ajay Mandlekar\*, Danfei Xu Preprint 2024 [paper] [website]
- RoboCasa: Large-Scale Simulation of Everyday Tasks for Generalist Robots Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, Yuke Zhu RSS 2024 [paper] [website] [code]
- IntervenGen: Interventional Data Generation for Robust and Data-Efficient Robot Imitation Learning Ryan Hoque, Ajay Mandlekar\*, Caelan Garrett\*, Ken Goldberg, Dieter Fox IROS 2024
   [paper] [website]
- Voyager: An Open-Ended Embodied Agent with Large Language Models Guanzhi Wang, Yuqi Xie, Yunfan Jiang\*, Ajay Mandlekar\*, Chaowei Xiao, Yuke Zhu, Linxi Fan†, Anima Anandkumar† TMLR 2024 [paper] [website] [code]
- MimicGen: A Data Generation System for Scalable Robot Learning using Human Demonstrations Ajay Mandlekar, Soroush Nasiriany\*, Bowen Wen\*, Iretiayo Akinola, Yashraj Narang, Linxi Fan, Yuke Zhu, Dieter Fox CoRL 2023 Oral (CoRL2023 Towards Generalist Robots Workshop) [paper] [website] [code]
- Human-in-the-Loop Task and Motion Planning for Imitation Learning Ajay Mandlekar\*, Caelan Garrett\*, Danfei Xu, Dieter Fox CoRL 2023
   Best Paper Runner-Up and Oral (Deployable@CoRL2023 Workshop) [paper] [website]
- Imitating Task and Motion Planning with Visuomotor Transformers Murtaza Dalal, Ajay Mandlekar\*, Caelan Garrett\*, Ankur Handa, Ruslan Salakhutdinov, Dieter Fox CoRL 2023
   [paper] [website] [code]

- MineDojo: Building Open-Ended Embodied Agents with Internet-Scale Knowledge Linxi Fan, Guanzhi Wang\*, Yunfan Jiang\*, Ajay Mandlekar, Yuncong Yang, Haoyi Zhu, Andrew Tang, De-An Huang, Yuke Zhu, Anima Anandkumar NeurIPS 2022 Datasets and Benchmarks Track Outstanding Paper Award [paper] [website] [code] [blog post]
- What Matters in Learning from Offline Human Demonstrations for Robot Manipulation
   Ajay Mandlekar, Danfei Xu, Josiah Wong, Soroush Nasiriany, Chen Wang, Rohun Kulkarni, Li Fei-Fei, Silvio Savarese, Yuke Zhu, Roberto Martin-Martin

   CoRL 2021
   Oral (6.5% Acceptance)
   [paper] [website] [video] [blog post] [code]
- Error-Aware Imitation Learning from Teleoperation Data for Mobile Manipulation Josiah Wong, Albert Tung, Andrey Kurenkov, Ajay Mandlekar, Li Fei-Fei, Silvio Savarese, Roberto Martin-Martin CoRL 2021 [paper] [website] [video]
- Human-in-the-Loop Imitation Learning using Remote Teleoperation Ajay Mandlekar, Danfei Xu\*, Roberto Martin-Martin\*, Yuke Zhu, Li Fei-Fei, Silvio Savarese Preprint 2021 [paper] [website] [video]
- Learning Multi-Arm Manipulation through Collaborative Teleoperation Albert Tung\*, Josiah Wong\*, Ajay Mandlekar, Roberto Martin-Martin, Yuke Zhu, Li Fei-Fei, Silvio Savarese ICRA 2021 Best Multi-Robotic Systems Paper Finalist [paper] [website] [video]
- Deep Affordance Foresight: Planning Through What Can Be Done in the Future Danfei Xu, Ajay Mandlekar, Roberto Martin-Martin, Yuke Zhu, Silvio Savarese, Li Fei-Fei ICRA 2021
   [paper] [website] [video]
- robosuite: A Modular Simulation Framework and Benchmark for Robot Learning Yuke Zhu, Josiah Wong, Ajay Mandlekar, Roberto Martin-Martin Technical Report [paper] [website] [code]
- Learning to Generalize Across Long-Horizon Tasks from Human Demonstrations Ajay Mandlekar\*, Danfei Xu\*, Roberto Martin-Martin, Silvio Savarese, Li Fei-Fei RSS 2020
   [paper] [website] [blog post] [video]
- IRIS: Implicit Reinforcement without Interaction at Scale for Learning Control from Offline Robot Manipulation Data
   Ajay Mandlekar, Fabio Ramos, Byron Boots, Silvio Savarese, Li Fei-Fei, Animesh Garg, Dieter Fox ICRA 2020
   [paper] [website] [video]

• Scaling Robot Supervision to Hundreds of Hours with RoboTurk: Robotic Manipulation Dataset through Human Reasoning and Dexterity Ajay Mandlekar, Jonathan Booher, Max Spero, Albert Tung, Anchit Gupta, Yuke Zhu, Animesh Garg, Silvio Savarese. Li Fei-Fei **IROS 2019 Best Cognitive Robotics Paper Finalist** [paper] [website] [blog post]

• RoboTurk: A Crowdsourcing Platform for Robotic Skill Learning through Imitation Ajay Mandlekar, Yuke Zhu, Animesh Garg, Jonathan Booher, Max Spero, Albert Tung, Julian Gao, John Emmons, Anchit Gupta, Emre Orbay, Silvio Savarese, Li Fei-Fei **CoRL 2018** [paper] [website] [video]

# **Community Service**

### **Workshop Organizer**

- Co-organized Overlooked Aspects of Imitation Learning: Systems, Data, Tasks, and Beyond workshop at RSS 2022.
- Co-organized Advances and Challenges in Imitation Learning for Robotics workshop at RSS 2020.

#### **Peer Review**

• Reviewer for NeurIPS, ICML, RSS, CoRL, IROS, CVPR, IEEE T-RO.

## Education

### **Stanford University**

NDSEG Fellow PhD student, Electrical Engineering M.S. Computer Science, Class of 2018  $GPA \cdot 41$ 

#### **California Institute of Technology**

B.S. Electrical Engineering, B.S. Computer Science Class of 2016 Officer of Tau Beta Pi Engineering Honor Society GPA: 4.1 Programming Languages: Python, C, Matlab, R, OCaml, Haskell, Swift, IA32, x86, VHDL

# Work Experience

## **NVIDIA Seattle Robotics Lab**

Research Intern

• Developed IRIS algorithm for learning from offline crowdsourced robot manipulation data. Supervised by Dieter Fox.

### **Stanford Vision and Learning Lab**

Stanford EE PhD

• Advised by Silvio Savarese and Fei-Fei Li

#### Apple

Summer 2015

Summer 2019

Winter 2017 - Present

#### Software Engineering Intern on iOS Location and Motion

• Implemented some deep learning methods and presented to Craig Federighi, the SVP of Software.

### **Caltech High-Speed Integrated Circuits**

Research Intern, Robert J. McEliece and David Rutledge SURF Fellow

- Spent a summer in Professor Ali Hajimiri's lab. Implemented a wireless feedback mechanism using a Raspberry Pi, Android tablet, Arduino Nano.
- Designed both hardware (PCB Design, digital and analog circuit development) and software (control, sensing, and optimization) interfaces and implemented wireless and serial communication protocols.

#### **SLAC National Accelerator Laboratory**

Research Intern

- Studied accelerator physics relating to the operation of the Free Electron Laser and the production of extremely high-power x-rays used by biologists.
- Investigated the theory of harmonic lasing in order to try to make the shift towards shorter wavelength, higher energy radiation.

### Teaching

- CS 332 Advanced Survey of Reinforcement Learning (Teaching Assistant Stanford)
- CS 231N Convolutional Neural Networks for Visual Recognition (Teaching Assistant Stanford)
- ACM 95ab Complex Analysis, Differential Equations (Teaching Assistant Caltech)
- CS 24 Introduction to Computing Systems (Teaching Assistant Caltech)

## **Selected Coursework**

- CS 124 Operating Systems
  Built a command shell, bootloader, priority scheduler, MLFQ scheduler on top of Pintos.
- CS 155/156 Machine Learning and Data Mining
  - Netflix Challenge placed 3rd in class using a blend of Time-SVD++, RBM, and kNN models.
- EE 52/90/91 Embedded Systems Hardware, Analog Electronics Projects
  PCB Design, built an MP3 player, a "smart" dog bowl and a function generator.
- EE 111/112 Digital Signal Processing
  DSP Systems, Difference Equations, Z-Transform, Multirate Systems, Filter Banks, Filter Design.
- EE 113/114 Feedback and Control Circuits, Analog Circuit Design
  Basic feedback circuits, compensation, PID, digital control, amplifier design, SPICE simulation.
- EE 189 Design and Construction of Biodevices
  Built a pulse oximeter and a PCR machine.
- EE 364ab Convex Optimization

Summer 2014

Summer 2013