Jonathan Yang

jonathanyang0127.github.io

Education

Stanford University Stanford, CA
Ph.D. Computer Science *August 2022 - May 2027 (Expected)*

- Advisors Chelsea Finn, Dorsa Sadigh
- Intelligence through Robotic Interaction at Scale Lab (IRIS)
- Intelligent and Interactive Autonomous Systems Group (ILIAD)

University of California, Berkeley Berkeley, CA
B.S. Electrical Engineering and Computer Science *Graduated March 2022*

- GPA: 3.93/4.0
- Relevant Coursework: Math 110: Linear Algebra, CS 170: Efficient Algorithms and Intractable Problems, CS 188: Introduction to Artificial intelligence, EECS 126: Probability and Random Processes, EECS 127: Optimization Models in Engineering, CS 189: Machine Learning, CS 285: Deep Reinforcement Learning, CS 294-190: Advanced Topics in Decision Making and Control, EE 229A: Information Theory and Coding

Research Interests

vision-based robotic control, deep reinforcement learning

Experience

Student Researcher at Google DeepMind Current

- Embodied General Reasoning Team
- Researching how to train robot policies that generalize to a broad range of environments, tasks, and embodiments.
- Studying what capabilities and forms of reasoning robots need to understand to adapt to the variability encountered in everyday life.

Publications

2D Vision Encoders Should See in 3D: Zero-shot Camera Perspective Generalization for Manipulation In Submiss
Jonathan Yang *Chelsea Finn, Dorsa Sadigh*

- Trained a robotic foundation model that can generalize to novel camera perspectives using data from from robot demonstrations and static videos.
- Showed that the robot can accomplish a task using images from a camera held by a moving human.
- https://sites.google.com/view/robot-camera-perspective

Pushing the Limits of Cross-Embodiment Learning for Manipulation and Navigation RSS 2024
Jonathan Yang*, Arjun Bhokar, Dhruv Shah, Catherine Glossop, Quan Vuong, Chelsea Finn, Dorsa Sadigh, Sergey Levine*

- Trained a robotic foundation model on manipulation, navigation, and drone datasets
- Showed that navigation data can allow for positive transfer for egocentric manipulation tasks
- https://extreme-cross-embodiment.github.io

DROID: A Large-Scale In-the-Wild Robot Manipulation Dataset RSS 2024
Droid Dataset Team

- Training large robotic foundation models across manipulation, navigation, and drone datasets
- https://droid-dataset.github.io

Open X-Embodiment: Robotic Learning Datasets and RT-X Models ICLR 2024
Open X-Embodiment Collaboration

- Large open-source effort to consolidate robotic datasets

Polybot: Training One Policy Across Robots While Embracing Variability CoRL 2023
Jonathan Yang*, Chelsea Finn, Dorsa Sadigh.*

- https://sites.google.com/view/polybot-multirobot
- Accepted to the 2023 Conference on Robot Learning

Don't Start From Scratch: Leveraging Prior Data to Automate Robotic Reinforcement Learning RSS 2022
*Homer Walke, **Jonathan Yang**, Avi Singh, Albert Yu, Aviral Kumar, Sergey Levine.*

- https://sites.google.com/view/ariel-paper
- Accepted to RSS 2022 Learning from Diverse Offline Data Workshop, Accepted to CoRL 2022 Conference

COG: Connecting New Skills to Past Experience with Offline Reinforcement Learning CoRL 2020
*Avi Singh, Albert Yu, **Jonathan Yang**, Jesse Zhang, Aviral Kumar, Sergey Levine.*

- https://sites.google.com/view/cog-rl
- Accepted to CoRL 2020, oral presentation at NeurIPS 2020 Offline RL Workshop

Research

Embodied Reasoning and Robotic Foundation Models August 2022 – Present
Stanford Artificial Intelligence Laboratory (Prof. Chelsea Finn, Dorsa Sadigh)

- Investigating the forms of reasoning and methods needed to train robotic agents to performance tasks in diverse environments and embodiments.
- Investigating what forms of structured representations robot needs in order to successfully performance generalization dexterous manipulation.
- Training large-scale autonomous driving vision-language-action models deployed on real cars.

Reset-free Exploration for Lifelong Learning January 2021 – August 2022
Robotic AI and Learning Lab, Berkeley Artificial Intelligence Research (Prof. Sergey Levine)

- Investigated reset-free exploration for a new environment given a large, offline dataset of prior behavior
- Designed and scripted policies for 32 tasks consisting of picking multiple objects, placing them into containers, and manipulating drawers
- Utilizing a transformer-based task encoder, proposed an sampling algorithm that runs CEM on top of the latent space to simultaneously train and fine-tune for a new task
- Trained a WidowX200 robot to pick and place up a completely new object in a zero-shot manner, then reset this trajectory with an 80% success rate
- Showed that with only 20 demonstrated for a distinct task of picking up/placing a pot lid on a pot, multitask RL improved performance (80% success with multitask learning vs 30% without)

Transfer Learning for Dataset-Driven Reinforcement Learning March 2020 - August 2020
Robotic AI and Learning Lab, Berkeley Artificial Intelligence Research (Prof. Sergey Levine)

- Investigated transfer learning with large-scale, off-policy prior datasets in order to effectively learn long-horizon policies capable of performing multiple tasks and reasoning about initial conditions
- Developed method to chain behaviors with batch reinforcement learning outperformed oracle behavior cloning baselines by 2 times in a setting with sparse reward signal
- Demonstrated applicability to real-world robotics by training a WidowX200 robot to pick up an object inside a closed drawer given a large, diverse buffer of prior experience and target demonstrations for picking the object up from an open drawer
- Accepted to the 2020 Conference on Robotic Learning and Neurips poster sessions

Scalable and Flexible Real-World Robotic Learning September 2019 - January 2020
Robotic AI and Learning Lab, Berkeley Artificial Intelligence Research (Prof. Sergey Levine)

- Created Roboverse: a pybullet-based simulation for robotic manipulation, with a focus on realistic physics and handing a diverse set of objects without sacrificing use of contact forces.
- Developed ROS container for the Sawyer and WidowX robots featuring 3DOF/6DOF gym environments, rgb/depth observations from kinect camera, and position/velocity control with inverse kinematics
- Formulated Markovian data collection policies capable of autonomous data collection without human supervision
- Obtained an 100% success rate for end-to-end single-object robotic grasping of a toy in a 15 by 20 inch tray

Projects

Coordination and Cooperation on the Energy Grid through Hierarchical RL September - December 2021

- In collaboration with PhD candidate Lucas Spangher
- Exploited the existing hierarchical structure in the energy grid to create a single hierarchical agent to control aggregators that buy and sell energy to independent groups of houses in order to decrease overall energy demand
- Utilized a "vehicle-to-grid" mechanism, i.e. a fleet of electric vehicles that charge in one microgrid and deposit energy in another microgrid to share energy between houses under the hierarchical agent's policy
- Demonstrated using simulation that utilizing a FeUdal networks-like approach can result in lower energy prices than modeling aggregators as individual agents

RL with Low-Dimensional Supervision for Zero-Shot Domain Generalization January - March 2020

- Developed a RL algorithm building upon PPO which utilizes low-dimensional state supervision for RGB observations during training time to generalize better at test time without supervision
- Demonstrated that adding information about the fruit and wall locations for the FruitBot OpenAI Procgen environment during training can boost performance at test time

Scalable Distributed Classification for Massive Data September - December 2020

- Developed information-theoretic metric to split large datasets for distribution classification during training time and merge them during inference in order to minimally affect classification accuracy
- Demonstrated that training and deploying a distributed model on Amazon Sagemaker for the MNIST 8M and HIGGS datasets lead to a less than 0.5% and 0.02% change in error respectively for training and fine-tuning with 20 separate classifiers

Detaching Intrinsic and Extrinsic Motivation through Goal-Conditioned RL September - December 2020

- Count-based exploration bonuses in reinforcement learning can be unstable because it causes the reinforcement learning algorithm to chase a moving reward
- Posed intrinsic motivation as a goal-selection problem in order provide more directed exploration. Utilized empowerment as a metric to select new goals.
- Achieved better exploration on 2d point-mass mazes than count-based bonuses and random network distillation

Teaching

TA for CS 224R: Deep Reinforcement Learning March 2023 - June 2023

- Worked as TA for graduate deep reinforcement learning class at Stanford University
- Mentored 30 Masters and Ph.D. students on deep reinforcement learning projects

TA for CS 188: Introduction to Artificial Intelligence January 2019 - May 2022

- Worked for five semesters as course staff for AI class teaching search, reinforcement learning, and probabilistic graphical models
- Developed new project utilizing deep reinforcement learning to train a pacman agent
- Led project design and release for Fall 2021

Activities

Stanford Wind Symphony September 2022 - Present

- Clarinetist in the University Wind Symphony

Robomaster at Berkeley (ICRA Robomaster Challenge) August 2019 - March 2020

- Developed ROS gazebo simulation to model competition where robots shoot projectiles at enemies to score points
- Designed neural fictitious self-play algorithm using the PPO to improve strategy initially seeded with behavior cloning from past successful trajectories

UAVs at Berkeley December 2018 - May 2019

- Developed custom shortest path algorithm to find the optimal path of a fixed-wing drone to drop a payload while avoiding cylindrical obstacles
- Addressed the intractability of naive approach of dividing up a 3-d map into grid squares and running A* by recursively scaling grid point sizes around obstacles
- Fit a spline to the shortest path in order for the trajectory to be achievable by a fixed wing aircraft