

# JEFFREY WU

☎ 925-477-8372   ✉ [jgw2140@columbia.edu](mailto:jgw2140@columbia.edu)   🖥 [jeffreywu13579.github.io/](https://jeffreywu13579.github.io/)

## EDUCATION

---

**Columbia University** Sept 2024 –  
*Ph.D. in Computer Science*

**University of California, Berkeley** Aug 2023 – May 2024  
*M.S in EECS (5th Year Program)* *GPA: 4.0/4.0*

**Relevant Coursework:** Natural Language Processing, Statistical Learning Theory

**Master's Thesis:** *V-Former: Offline RL with Temporally-Extended Actions*

**University of California, Berkeley** Aug 2019 – May 2023  
*B.A in Computer Science* *GPA: 4.0/4.0*

**Relevant Coursework:** Machine Learning, Artificial Intelligence, Deep Neural Networks, Deep Reinforcement Learning, Computer Vision, Probability and Random Processes, Optimization, Digital Signal Processing, Time Series

## PUBLICATIONS

---

### **Precise and Dexterous Robotic Manipulation via Human-in-the-Loop Reinforcement Learning**

Jianlan Luo, Charles Xu, **Jeffrey Wu**, Sergey Levine

*Under Review at Science Robotics*

### **Action-Quantized Offline Reinforcement Learning for Robotic Skill Learning**

Jianlan Luo, Perry Dong, **Jeffrey Wu**, Aviral Kumar, Xinyang Geng, Sergey Levine

*Conference on Robot Learning (CoRL), 2023*

### **FMB: A Functional Manipulation Benchmark for Generalizable Robotic Learning**

Jianlan Luo\*, Charles Xu\*, Fangchen Liu, Liam Tan, Zipeng Lin, **Jeffrey Wu**, Pieter Abbeel, Sergey Levine

*International Journal of Robotics Research (IJRR), 2024*

### **Open X-Embodiment: Robotic Learning Datasets and RT-X Models**

Open X-Embodiment Collaboration

*IEEE International Conference on Robotics and Automation (ICRA), 2024*

## RESEARCH EXPERIENCE

---

**Causal Artificial Intelligence Lab** Sept 2024 –  
*Advised by Prof. Elias Bareinboim* *New York, NY*

**Robotic AI & Learning Lab (RAIL) at BAIR** Feb 2022 – Sept 2024  
*Advised by Prof. Sergey Levine* *Berkeley, CA*

### **Human-in-the-Loop Reinforcement Learning for Robotic Manipulation** *(Under Review)*

- Presented a human-in-the-loop vision-based RL system that learn near-perfect success rate policies in just 1 to 2.5 hours of training across range of tasks, including dynamic manipulation, precision assembly, and dual-arm coordination tasks
- Surpassed imitation learning baselines and prior RL approaches in success rate, training time, and cycle time

### **Offline RL with Temporally-Extended Actions** *(Master's Thesis)*

- Proposed using a transformer-based policy trained with an advantage-weighted sequence modeling objective to predict action chunks to more effectively model multimodal and non-Markovian policies
- Improved performance on Robomimic tasks and outperform prior work on datasets with multiple action frequencies

### **Action-Quantized Offline Reinforcement Learning** *(CoRL 2023)*

- Proposed learning state-conditioned action discretization with a VQ-VAE in conjunction with discrete formulations of offline RL methods (BRAC, CQL, IQL) to reduce approximation errors in offline RL constraints and regularizers
- Improved performance up to 2-3x compared to the continuous offline RL counterpart methods in D4RL and Robomimic

### **FMB: Functional Manipulation Benchmark** *(IJRR 2024)*

- Provided a benchmark consisting of easily reproducible 3D-printed objects and boards that focus on functional manipulation skills, such as grasping, reorientation, and assembly, with a provided imitation learning framework
- Enabled evaluation of generalization across objects and positions, and offered a testbed for methods for solving complex multi-stage assembly tasks

## INDUSTRY EXPERIENCE

---

### **Amazon Web Services (AWS)**

May 2021 – Aug 2021

*Software Development Engineer Intern*

*East Palo Alto, CA*

- Migrated AWS S3 accesses in AWS Redshift from the original in-house library to the official AWS C++ SDK while maintaining existing functionality
- Carried out performance benchmarking and performance analysis

## TEACHING

---

### **EE122: Introduction to Communication Networks (UC Berkeley)**

*Teaching Assistant*

*Spring 2023*

- Hosted weekly office hours, led discussion sections, answered student questions on Ed discussion forum, created class materials, graded homeworks and exams

### **EECS127: Optimization Models in Engineering (UC Berkeley)**

*Reader/Tutor*

*Spring 2022, Fall 2022*