

JACKY LIANG

jackyliang42@gmail.com ♦ www.jacky.io

EDUCATION

Carnegie Mellon University

August 2018 - December 2022

PhD

Robotics Institute

University of California Berkeley

August 2014 - December 2017

Bachelor of Science

GPA: 3.93

Electrical Engineering and Computer Science, High Honors

RESEARCH

Research Scientist, Google DeepMind

June 2023 - Present

Developing and applying foundation models for robotics.

PhD Student, Carnegie Mellon University, IAM Lab

August 2018 - December 2022

Co-advised by Professor Oliver Kroemer and Professor Maxim Likhachev.

Researched enabling robust and generalizable robot manipulation by incorporating structured priors into learning-based methods like reinforcement learning and self-supervised learning [5-6, 10-16].

My thesis “Learning with Structured Priors for Robust Robot Manipulation” can be viewed at www.jacky.io/links/thesis. A recording of the defense can be found at www.jacky.io/links/defense

Research Intern, Google Brain

May 2022 - September 2022

Worked on using Large Language Models (LLMs) to reason and plan robot tasks, applying action-oriented multi-modal foundation models to robotics. In [17] we used language-based reasoning with grounded visual feedback to enable more robust high-level task planning. In [18] we used code-generation as the medium for LLMs to express robot policies, enabling zero-shot completion of a wide-variety of language-specified tasks. Hosted by Andy Zeng and Pete Florence.

Research Intern, Nvidia

May 2019 - August 2019

Worked on using GPU-accelerated physics simulation and BioTac contact feedback for vision-free, in-hand object pose tracking during dexterous manipulation [9]. Hosted by Ankur Handa.

Research Intern, Nvidia

January 2018 - June 2018

Worked on large-scale distributed Deep Reinforcement Learning with GPU-accelerated robotics simulation [4]. Hosted by Viktor Makoviychuk.

Undergraduate Research Assistant, UC Berkeley, Autolab

October 2015 - December 2017

Advised by Professor Ken Goldberg.

Dexterity Network (Dex-Net) - We used analytical grasp metrics to supervise convolutional neural networks to rate grasp quality from depth images [2]. This is one of the earliest works applying large-scale deep learning to plan robust robot grasps.

Deep Imitation Learning - Deep learning of vision-feedback policies for bilateral manipulation tasks from human demonstrations collected via a custom real-time teleoperation system [3].

- [23] **Jacky Liang**, Fei Xia, Wenhao Yu, Andy Zeng, Montserrat Gonzalez Arenas, Maria Attarian, Maria Bauza, Matthew Bennice, Alex Bewley, Adil Dostmohamed, Chuyuan Kelly Fu, Nimrod Gileadi, Marissa Giustina, Keerthana Gopalakrishnan, Leonard Hasenclever, Jan Humplik, Jasmine Hsu, Nikhil Joshi, Ben Jyenis, Chase Kew, Sean Kirmani, Tsang-Wei Edward Lee, Kuang-Huei Lee, Assaf Hurwitz Michaely, Joss Moore, Ken Oslund, Dushyant Rao, Allen Ren, Baruch Tabanpour, Quan Vuong, Ayzaan Wahid, Ted Xiao, Ying Xu, Vincent Zhuang, Peng Xu, Erik Frey, Ken Caluwaerts, Tingnan Zhang, Brian Ichter, Jonathan Tompson, Leila Takayama, Vincent Vanhoucke, Izhak Shafran, Maja Mataric, Dorsa Sadigh, Nicolas Heess, Kanishka Rao, Nik Stewart, Jie Tan, Carolina Parada *Learning to Learn Faster from Human Feedback with Language Model Predictive Control*. February 2024. <https://arxiv.org/abs/2402.11450>
- [22] Soroush Nasiriany, Fei Xia, Wenhao Yu, Ted Xiao, **Jacky Liang**, Ishita Dasgupta, Annie Xie, Danny Driess, Ayzaan Wahid, Zhuo Xu, Quan Vuong, Tingnan Zhang, Tsang-Wei Edward Lee, Kuang-Huei Lee, Peng Xu, Sean Kirmani, Yuke Zhu, Andy Zeng, Karol Hausman, Nicolas Heess, Chelsea Finn, Sergey Levine, Brian Ichter. *PIVOT: Iterative Visual Prompting Elicits Actionable Knowledge for VLMs*. February 2024. <https://arxiv.org/abs/2402.07872>
- [21] Chengshu Li, **Jacky Liang**, Andy Zeng, Xinyun Chen, Karol Hausman, Dorsa Sadigh, Sergey Levine, Li Fei-Fei, Fei Xia, Brian Ichter. *Chain of code: Reasoning with a language model-augmented code emulator*. December 2023. <https://arxiv.org/abs/2312.04474>

CONFERENCE PUBLICATIONS

- [20] Montserrat Gonzalez Arenas, Ted Xiao, Sumeet Singh, Vidhi Jain, Allen Z Ren, Quan Vuong, Jake Varley, Alexander Herzog, Isabel Leal, Sean Kirmani, Dorsa Sadigh, Vikas Sindhwani, Kanishka Rao, **Jacky Liang**, Andy Zeng. *How to prompt your robot: A promptbook for manipulation skills with code as policies*. International Conference on Robotics and Automation (ICRA). May 2024.
- [19] Open X-Embodiment Collaboration *Open X-embodiment: Robotic Learning Datasets and RT-X Models*. International Conference on Robotics and Automation (ICRA). May 2024. <https://arxiv.org/abs/2310.08864v4>
- [18] **Jacky Liang**, Wenlong Huang, Fei Xia, Peng Xu, Karol Hausman, Brian Ichter, Pete Florence, Andy Zeng. *Code as Policies: Language Model Programs for Embodied Control*. International Conference on Robotics and Automation (ICRA). May 2023. **Outstanding Paper in Robot Learning** <https://arxiv.org/abs/2209.07753>
- [17] Wenlong Huang*, Fei Xia*, Ted Xiao*, Harris Chan, **Jacky Liang**, Pete Florence, Andy Zeng, Jonathan Tompson, Igor Mordatch, Yevgen Chebotar, Pierre Sermanet, Noah Brown, Tomas Jackson, Linda Luu, Sergey Levine, Karol Hausman, Brian Ichter. *Inner Monologue: Embodied Reasoning through Planning with Language Models*. Conference on Robot Learning (CoRL). December 2022. <https://arxiv.org/abs/2207.05608>
- [16] **Jacky Liang**, Xianyi Cheng, Oliver Kroemer. *Learning Preconditions of Hybrid Force-Velocity Controllers for Contact-Rich Manipulation*. Conference on Robot Learning (CoRL). December 2022. <https://arxiv.org/abs/2206.12728>
- [15] **Jacky Liang***, Mohit Sharma*, Alex LaGrassa, Shivam Vats, Saumya Saxena, Oliver Kroemer. *Search-Based Task Planning with Learned Skill Effect Models for Lifelong Robotic Manipulation*. International Conference on Robotics and Automation (ICRA). May 2022. <https://arxiv.org/abs/2109.08771>
- [14] Vicky Zeng, Timothy E. Lee*, **Jacky Liang***, Oliver Kroemer. *Visual Identification of Articulated Object Parts*. International Conference on Intelligent Robots and Systems (IROS). September 2021. <https://arxiv.org/abs/2012.00284>

- [13] **Jacky Liang**, Oliver Kroemer. *Contact Localization for Robot Arms in Motion without Torque Sensing*. International Conference on Robotics and Automation (ICRA). May 2021. <https://arxiv.org/abs/2011.03142>
- [11] Mohit Sharma*, **Jacky Liang***, Jialiang Zhao, Alex LaGrassa, Oliver Kroemer. *Learning to Compose Hierarchical Object-Centric Controllers for Robotic Manipulation*. Conference on Robot Learning (CoRL). **Oral Presentation** November 2020. <https://arxiv.org/abs/2011.04627>
- [10] **Jacky Liang**, Saumya Saxena, Oliver Kroemer. *Learning Active Task-Oriented Exploration Policies for Bridging the Sim-to-Real Gap*. Robotics: Science and Systems (RSS). July 2020. <https://arxiv.org/abs/2006.01952>
- [9] **Jacky Liang**, Ankur Handa, Karl Van Wyk, Viktor Makoviychuk, Oliver Kroemer, Dieter Fox. *In-Hand Object Pose Tracking via Contact Feedback and GPU-Accelerated Robotic Simulation*. International Conference on Robotics and Automation (ICRA). May 2020. <https://arxiv.org/abs/2002.12160>
- [8] Ankur Handa, Karl Van Wyk, Wei Yang, **Jacky Liang**, Yu-Wei Chao, Qian Wan, Stan Birchfield, Nathan Ratliff, Dieter Fox. *DexPilot: Vision Based Teleoperation of Dexterous Robotic Hand-Arm System*. International Conference on Robotics and Automation (ICRA). May 2020. <https://arxiv.org/abs/1910.03135>
- [6] Austin S. Wang, Wuming Zhang, Daniel Troniak, **Jacky Liang**, Oliver Kroemer. *Homography-Based Deep Visual Servoing Methods for Planar Grasps*. International Conference on Intelligent Robots and Systems (IROS). November 2019.
- [5] Jialiang Zhao, **Jacky Liang**, Oliver Kroemer. *Towards Precise Robotic Grasping by Probabilistic Post-grasp Displacement Estimation*. Field and Service Robotics (FSR). August 2019. <https://arxiv.org/abs/1909.02129>
- [4] **Jacky Liang***, Viktor Makoviychuk*, Ankur Handa*, Nuttapong Chentanez, Miles Macklin, Dieter Fox. *GPU-Accelerated Robotic Simulation for Distributed Reinforcement Learning*. Conference on Robot Learning (CoRL). October 2018. <https://arxiv.org/abs/1810.05762>
- [3] **Jacky Liang**, Jeffrey Mahler, Michael Laskey, Pusong Li, Ken Goldberg. *Using dVRK Teleoperation to Facilitate Deep Learning of Automation Tasks for an Industrial Robot*. Conference on Automation Science and Engineering (CASE). Xian, China. August 2017. **Finalist, Best Student Paper Award**
- [2] Jeffrey Mahler, **Jacky Liang**, Sherdil Niyaz, Michael Laskey, Richard Doan, Xinyu Liu, Juan Aparicio Ojea, Ken Goldberg. *Dex-Net 2.0: Deep Learning to Plan Robust Grasps with Synthetic Point Clouds and Analytic Grasp Metrics*. Robotics: Science and Systems (RSS). MIT Cambridge, MA. July 2017.
- [1] Menglong Guo, David V. Gealy, **Jacky Liang**, Jeffrey Mahler, Aimee Goncalves, Stephen McKinley, Ken Goldberg. *Design of Parallel-Jaw Gripper Tip Surfaces for Robust Grasping*. International Conference on Robotics and Automation (ICRA). Singapore. May 2017.

**Equal Contribution*

OTHER PUBLICATIONS

- [12] Kevin Zhang*, Mohit Sharma*, **Jacky Liang***, Oliver Kroemer. *A Modular Robotic Arm Control Stack for Research: Franka-Interface and FrankaPy*. November 2020. <https://arxiv.org/abs/2011.02398>
- [7] Nathan Zhang*, **Jacky Liang***, Amanda Tomlinson*, Frank Boensch, Anant Sahai. *Undergraduate-Led Survey Class to Improve CS Education for New Students*. SIGCSE '20: Proceedings of the 51st ACM Technical Symposium on Computer Science Education. February 2020.

**Equal Contribution*

OPEN-SOURCE SOFTWARE

Language Model Programs, Interactive colab notebook that implements language model programs for a simulated table-top manipulation domain.

<https://colab.research.google.com/drive/124TE4TsGYrvduzeDclufyvwwc2qbbrE>

FrankaPy, Python interface for the Franka Emika Panda robot arm.

<https://github.com/iamlab-cmu/frankapy>

IsaacGym-Utils, Python library that provides simpler APIs for the Nvidia Isaac Gym robot simulator.

<https://github.com/iamlab-cmu/isaacgym-utils>

Simple-ZMQ, Python library that uses zmq to send arbitrary objects over a network.

https://github.com/jacky-liang/simple_zmq

Async-Savers, Python library for asynchronously saving data in shards.

https://github.com/jacky-liang/async_savers

Data-Learning-Boilerplate, Boilerplate Python code for collecting data and training neural networks using Weights and Biases and PyTorch Lightning.

<https://github.com/jacky-liang/data-learning-boilerplate>

YuMiPy, Python interface for the ABB YuMi robot arm.

<https://github.com/BerkeleyAutomation/yumipy>

OUTREACH AND SERVICE

Area Chair, Conference on Robot Learning

April 2024 - Present

Editor, Last Week in AI (<https://lastweekin.ai/>)

November 2018 - Present

Mentor, CMU Project Ignite. Led local high school students in a project that uses derivative-free optimization to improve cookie recipes.

Spring 2021

Mentor, CMU AI Mentoring Program

Fall 2019, Fall 2020, Fall 2021

Editor, Berkeley AI Research Blog (bair.berkeley.edu/blog)

August 2017 - December 2017

Tutoring and Education Officer, UC Berkeley Eta Kappa Nu

January 2016 - May 2017

TEACHING

CMU

January 2021 - May 2021

Teaching Assistant

- 16-350 Planning Techniques for Robotics. Topics include graph search, sampling-based planning, symbolic planning.

CMU

January 2020 - May 2020

Teaching Assistant

- 16-662 Robot Autonomy. Topics include task-space control, motion planning, task planning, grasping.

UC Berkeley

August 2016 - December 2017

Undergraduate Student Instructor

- CS188 Introduction to Artificial Intelligence. Topics include search, games, probabilistic graphical models, reinforcement learning, and deep learning.
- EE16A Design of Information Devices and Systems. Topics include linear algebra, signals and systems, and circuit design.

AWARDS AND HONORS

Outstanding Paper in Robot Learning, International Conference on Robotics and Automation
May 2023

Graduate Research Fellow, National Science Foundation *August 2018*

Finalist, Best Student Paper, Conference on Automation Science and Engineering *August 2017*

Member, UC Berkeley Tau Beta Pi, Engineering Honors Society *December 2016*

Member, UC Berkeley Eta Kappa Nu, EECS Honors Society *May 2016*

Regent's and Chancellor's Scholar, UC Berkeley *August 2014*

STUDENTS MENTORED

Jeff Tan, Learning robust 3D representations for manipulation. *February 2021 - May 2022*

Vicky Zeng, First-author on *Visual Identification of Articulated Object Parts*. Currently PhD Student at Johns Hopkins University *March 2020 - May 2021*