

Short Bio I am a recent Ph.D. graduate from UC Berkeley, where I worked with Professor Sergey Levine in the Berkeley Artificial Intelligence Research Lab. My motivation for pursuing research is to deploy machine learning models that are performant, reliable, and fair. To this end, my dissertation work focused on adapting models to be robust to distribution shifts. My early Ph.D. work spanned robotics, reinforcement learning, and deep generative modeling.

I am a lecturer for the spring 2022 semester, instructing the machine learning (CS 189) and deep learning (CS 182/282A) courses at UC Berkeley. This unique opportunity has helped solidify my foundational knowledge of machine learning as well as keep me closely familiar with the latest state of the art in deep neural networks. I am confident that both my research and teaching backgrounds will contribute greatly to my career as a research scientist.

Education **University of California, Berkeley**
Ph.D., Computer Science, August 2016 - December 2021.
Advisor: Sergey Levine.
Dissertation: Adaptation Based Approaches to Distribution Shift Problems.
B.A., Computer Science, August 2012 - May 2016.
GPA: 4.00 (Highest Distinction).

Research **University of California, Berkeley**
Graduate Student Researcher, August 2016 - December 2021.
Berkeley Artificial Intelligence Research (BAIR) Lab.

Google
Research Intern, February - May 2019.
Host: Matt Hoffman.

University of California, Berkeley
Undergraduate Researcher, January 2014 - August 2016.
UC Berkeley Robot Learning Lab (RLL).

Teaching **University of California, Berkeley**
Lecturer, Spring 2022 (present).
Introduction to Machine Learning (CS 189).
Designing, Visualizing, and Understanding Deep Neural Networks (CS 182/282A).

Course Instructor, Summer 2016.
Structure and Interpretation of Computer Programs (CS 61A).

Teaching Assistant, Spring 2016, Fall 2021 (Head Role).
Introduction to Machine Learning (CS 189/289A).

Teaching Assistant, Fall 2015.
Introduction to Artificial Intelligence (CS 188).

Teaching Assistant, Spring 2014, Fall 2014, Spring 2015.
Structure and Interpretation of Computer Programs (CS 61A).

Publications

Marvin Zhang, Sergey Levine, Chelsea Finn.

MEMO: Test Time Robustness via Adaptation and Augmentation.

Under review. arXiv 2110.09506.

Marvin Zhang*, Henrik Marklund*, Nikita Dhawan*, Abhishek Gupta, Sergey Levine, Chelsea Finn.

Adaptive Risk Minimization: Learning to Adapt to Domain Shift.

NeurIPS, 2021. arXiv 2007.02931.

Pang Wei Koh*, Shiori Sagawa*, Henrik Marklund, Sang Michael Xie, Marvin Zhang *et al.*

WILDS: A Benchmark of in-the-Wild Distribution Shifts.

ICML, 2021. arXiv 2012.07421.

Laura Smith, Nikita Dhawan, Marvin Zhang, Pieter Abbeel, Sergey Levine.

AVID: Learning Multi-Stage Tasks via Pixel-Level Translation of Human Videos.

RSS, 2020. arXiv 1912.04443.

Michael Janner, Justin Fu, Marvin Zhang, Sergey Levine.

When to Trust Your Model: Model-Based Policy Optimization.

NeurIPS, 2019. arXiv 1906.08253.

Marvin Zhang*, Sharad Vikram*, Laura Smith, Pieter Abbeel, Matthew Johnson, Sergey Levine.

SOLAR: Deep Structured Latent Representations for Model-Based Reinforcement Learning.

ICML, 2019. arXiv 1808.09105.

Yevgen Chebotar*, Karol Hausman*, Marvin Zhang*, Gaurav Sukhatme, Stefan Schaal, Sergey Levine.

Combining Model-Based and Model-Free Updates for Trajectory-Centric Reinforcement Learning.

ICML, 2017. **Best paper award**, RSS Workshop on New Frontiers for Deep Learning in Robotics, 2017. arXiv 1703.03078.

Marvin Zhang*, Xinyang Geng*, Jonathan Bruce*, Ken Caluwaerts, Massimo Vespignani,

Vytas SunSpiral, Pieter Abbeel, Sergey Levine.

Deep Reinforcement Learning for Tensegrity Robot Locomotion.

ICRA, 2017. arXiv 1609.09049.

Marvin Zhang, Zoe McCarthy, Chelsea Finn, Sergey Levine, Pieter Abbeel.

Learning Deep Neural Network Policies with Continuous Memory States.

ICRA, 2016. **Contributed talk**, NIPS Reasoning, Attention, Memory Workshop, 2015. arXiv 1507.01273.

Awards and Fellowships

NDSEG Fellowship, September 2017 - August 2021.

Neural Information Processing Systems (NeurIPS) Top Reviewer, 2019, 2021.

International Conference on Machine Learning (ICML) Top Reviewer, 2020.

NSF Graduate Research Fellowship, Declined April 2017.

UC Berkeley EECS Departmental Fellowship, August - December 2016.

UC Berkeley EECS Honors Degree Program, August 2014 - May 2016.

UC Berkeley Outstanding Graduate Student Instructor Award, May 2016.

UC Berkeley EECS Warren Dere Design Award, April 2016.

CRA Outstanding Undergraduate Researcher Honorable Mention, December 2015.

Teaching Projects

Brian Hou, **Marvin Zhang**, John DeNero.

Yelp Maps: A Restaurant Recommendation Engine.

SIGCSE Nifty Assignments, 2016.

<http://nifty.stanford.edu/2016/hou-zhang-denero-yelp-maps/>

Other Experience

University of California, Berkeley

Research Experience for Undergraduates (REU) Organizer, December 2020 - October 2021.

Berkeley Artificial Intelligence Research (BAIR) Lab.

<https://bair.berkeley.edu/reu.html>

Outreach Coordinator, September 2016 - August 2018.

UC Berkeley Robot Learning Lab (RLL).

<https://rll.berkeley.edu/outreach/>

Prism Skylabs

Engineering Intern, May - August 2013.

Mentors: Doug Johnston and Michael Fogel.