

ALEXANDER LI

Homepage: alexanderli.com.

Email: alexanderli@cmu.edu

EDUCATION

Carnegie Mellon University

Ph.D. in Machine Learning

Advisor: Deepak Pathak

GPA: 4.00 /4.00

August 2020 - present

University of California, Berkeley

B.S., M.S. in Electrical Engineering and Computer Science

GPA: 4.00/4.00, *Graduated with Highest Honors*

August 2016 - May 2020

HONORS/AWARDS

National Science Foundation Graduate Research Fellowship	2020
1st place, Citadel San Francisco Invitational Data Open	2019
Mark D. Weiser Excellence in Computing Scholarship at UC Berkeley	2019
Accel Scholar	2018
UC Berkeley Alumni Association Leadership Award	2017
UC Berkeley Regents' and Chancellor's Scholar	2016

EXPERIENCE

Carnegie Mellon Machine Learning Department

Ph.D. Student

August 2020 - present

Pittsburgh, PA

- Advised by Prof. Deepak Pathak
- Collaborators: Alexei Efros
- Research focus: understanding and changing inductive biases in deep learning optimization.

Berkeley Artificial Intelligence Research Lab

Undergraduate Researcher

January 2018 - May 2020

Berkeley, CA

- Advised by Prof. Pieter Abbeel
- Collaborators: Lerrel Pinto, Carlos Florensa, Ignasi Clavera
- Research focus: deep reinforcement learning, particularly hierarchical and multi-task algorithms.

Machine Learning at Berkeley

Project Manager, AutoQuad

Aug 2017 - May 2018

Berkeley, CA

- Led team of six members to improve long-term route planning for autonomous quadcopters.
- Designed computer vision and reinforcement learning algorithms to train drones to reach goals.

Machine Learning Engineer, Wish Inc.

Berkeley, CA

- Developed computer vision algorithms to score user-uploaded images at production-level timescales.

PUBLICATIONS AND PREPRINTS

Understanding Collapse in Non-Contrastive Siamese Representation Learning

Alexander Li, Alexei A. Efros, Deepak Pathak

Coming soon to arxiv

Functional Regularization for Reinforcement Learning via Learned Fourier Features

Alexander Li, Deepak Pathak

Neural Information Processing Systems (NeurIPS) 2021

Generalized Hindsight for Reinforcement Learning

Alexander Li, Lerrel Pinto, Pieter Abbeel

Neural Information Processing Systems (NeurIPS) 2020

Sub-policy Adaptation for Hierarchical Reinforcement Learning

Alexander Li*, Carlos Florensa*, Ignasi Clavera, Pieter Abbeel

International Conference on Learning Representations (ICLR) 2020

Autoregressive Models: What Are They Good For?

Murtaza Dalal*, **Alexander Li***, Rohan Taori*

NeurIPS 2019 Workshop on Information Theory and Machine Learning

Sunspot Rotation and the M-Class Flare in Solar Active Region NOAA 11158

Alexander Li, Yang Liu

Solar Physics 2015

TEACHING

Teaching Assistant

Carnegie Mellon 10-708: Deep Reinforcement Learning

Fall 2021

Berkeley CS 294-158: Deep Unsupervised Learning

Spring 2020

Berkeley EECS 126: Probability and Random Processes

Fall 2019

Berkeley CS 188: Artificial Intelligence

Fall 2018, Spring 2019

Academic Intern

Berkeley CS 189: Machine Learning

Spring 2018

Reader

Berkeley CS 70: Discrete Mathematics and Probability Theory

Fall 2017

PROFESSIONAL SERVICE

Reviewer:

International Conference on Machine Learning (ICML)

2021, 2022

TECHNICAL STRENGTHS

PyTorch, JAX, NumPy

COURSEWORK

Carnegie Mellon University

Selected graduate coursework: Computer Vision (A+), Convex Optimization (A+), Probabilistic Graphical Models (A+), Computational Game Solving (A)

University of California, Berkeley

Selected graduate coursework: Computer Vision (A+), Deep Reinforcement Learning (A), Deep Unsupervised Learning (A+), Theoretical Statistics (A), NLP (S), ML for Biology (S)

Selected undergraduate coursework: Machine Learning (A+), Artificial Intelligence (A+), Algorithms (A+), Probability/Stochastic Processes (A+), Convex Optimization (A+), Operating Systems (A), Internet Architecture (A+)