ALEXANDER LI

Homepage: alexanderli.com.
Email: alexanderli@cmu.edu

EDUCATION

Carnegie Mellon University *August 2020 - May 2025*
Ph.D. in Machine Learning

- Research focus: generalization, generative models, learning from Internet-scale data.
- Advised by Prof. Deepak Pathak
- Close collaborators: Prof. Alexei Efros

University of California, Berkeley *August 2016 - May 2020*
B.S., M.S. in Electrical Engineering and Computer Science
GPA: 4.00/4.00, Graduated with Highest Honors

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

WORK EXPERIENCE

Anthropic October 2024 - present
Member of Technical Staff *San Francisco, CA*

- Pretraining team

Meta AI Research June 2023 - Jan 2024
Research Scientist Intern *San Francisco, CA*

- Worked with Xinlei Chen
- Research focus: pretraining, finetuning, Transformers

Stanford Artificial Intelligence Lab June 2022 - Sep 2022
Visiting Researcher *Stanford, CA*

- Worked with Ananya Kumar and Prof. Tengyu Ma
- Research focus: improving self-supervised representation learning for OOD generalization

HONORS/AWARDS

Two Sigma PhD Fellowship Runner-Up (2nd place out of 160 nominees)	2023
Stability AI Compute Grant	2023
National Science Foundation Graduate Research Fellowship	2020
1st place, Citadel San Francisco Invitational Data Open	2019
Mark D. Weiser Excellence in Computing Scholarship	2019
Quantedge Award for Academic Excellence	2019
Accel Scholar	2018
UC Berkeley Alumni Association Leadership Award	2017
Edward Kraft Award for Freshmen	2016
UC Berkeley Regents' and Chancellor's Scholar (top $< 2\%$ of incoming students)	2016

PUBLICATIONS AND PREPRINTS

Prism: Discrete Diffusion Is Generalized Autoregression
Alexander Li, Deepak Pathak
To appear

Generative Classifiers Avoid Shortcut Solutions
Alexander Li, Ananya Kumar, Deepak Pathak
International Conference on Learning Representations (ICLR) 2020
Oral at ICML 2024 Workshop on Structured Probabilistic Inference and Generative Modeling

Who Needs Features? On the Surprising Effectiveness of Attention Transfer for Vision Transformers
Alexander Li, Yuandong Tian, Beidi Chen, Deepak Pathak, Xinlei Chen
Neural Information Processing Systems (NeurIPS) 2024

An Introduction to Vision-Language Modeling
Florian Bordes, Richard Yuanzhe Pang, Anurag Ajay, **Alexander Li**, et al.

Diffusion-TTA: Test-time Adaptation of Discriminative Models via Generative Feedback
Mihir Prabhudesai*, Tsung-Wei Ke*, **Alexander Li**, Deepak Pathak, Katerina Fragkiadaki
Neural Information Processing Systems (NeurIPS) 2023

Your Diffusion Model is Secretly a Zero-Shot Classifier
Alexander Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, Deepak Pathak
International Conference on Computer Vision (ICCV) 2023

Internet Explorer: Targeted Representation Learning on the Open Web
Alexander Li*, Ellis Brown*, Alexei A. Efros, Deepak Pathak
International Conference on Machine Learning (ICML) 2023

Understanding Collapse in Non-Contrastive Siamese Representation Learning
Alexander Li, Alexei A. Efros, Deepak Pathak
European Conference on Computer Vision (ECCV) 2022

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

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 16-884: Deep Learning for Robotics	Fall 2022
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 2021, 2022, 2023, 2024, 2025
Conference on Neural Information Processing Systems (top reviewer in 2023) 2022, 2023, 2024
International Conference on Learning Representations 2025

TECHNICAL STRENGTHS

PyTorch, JAX, NumPy