

Michael Zhang

michael@cs.toronto.edu ◇ <https://www.cs.toronto.edu/~michael/>

EDUCATION

University of Toronto Doctor of Philosophy, Computer Science Advised by Professor Jimmy Ba	<i>2018 - 2025</i>
University of California, Berkeley Master of Science, Electrical Engineering and Computer Science Advised by Professor Pieter Abbeel	<i>2017 - 2018</i> GPA: 4.000
University of California, Berkeley Bachelor of Arts, Computer Science Graduated with Highest Honors (top 3% of class)	<i>2014 - 2017</i> GPA: 4.000

RESEARCH AND INDUSTRY EXPERIENCE

Vector Institute <i>Student Researcher</i>	2018 - Present <i>Toronto, Ontario</i>
· Research on deep learning and reinforcement learning.	
Tesla <i>Autopilot Intern</i>	June 2022 - August 2022 <i>Palo Alto, California</i>
· Research and development on algorithms for control	
Google <i>Research Intern</i>	Feb 2020 - June 2020 <i>Toronto, Ontario</i>
· Research on offline reinforcement learning advised by Mohammad Norouzi.	
NextAI <i>Scientist in Residence</i>	2019 <i>Toronto, Ontario</i>
· Provide scientific and technical support to start-ups selected in NextAI accelerator.	
LinkedIn <i>Software Engineer Intern</i>	Summer 2018 <i>Dublin, Ireland</i>
· Built natural language models to understand what members write about on LinkedIn.	
Berkeley Artificial Intelligence Research <i>Student Researcher</i>	2016 - 2018 <i>Berkeley, California</i>
· Develop reinforcement learning algorithms for training real-world robots, object tracking, and solving sparse reward tasks.	
Workday <i>Data Science Intern</i>	Summer 2016 <i>Pleasanton, California</i>
· Model and predict problems that may occur for companies going live on the platform.	

PUBLICATIONS

Decomposed Prompting to Answer Questions on a Course Discussion Board
Brandon Jaipersaud, Paul Zhang, Jimmy Ba, Andrew Petersen, Lisa Zhang, Michael R. Zhang
International Conference on Artificial Intelligence in Education 2023

Multi-Rate VAE: Train Once, Get the Full Rate-Distortion Curve

Juhan Bae, Michael R. Zhang, Michael Ruan, Eric Wang, So Hasegawa, Jimmy Ba, Roger Grosse
International Conference on Learning Representations (ICLR) 2023 - top 5% of accepted papers

Learning Domain Invariant Representations in Goal-conditioned Block MDPs

Beining Han, Chongyi Zheng, Harris Chang, Keiran Paster, Michael R. Zhang, Jimmy Ba
Neural Information Processing Systems (NeurIPS) 2021

Analyzing Monotonic Linear Interpolation in Neural Network Loss Landscapes

James Lucas, Juhan Bae, Michael R. Zhang, Stanislav Fort, Richard Zemel, Roger Grosse
International Conference on Machine Learning (ICML) 2021

Autoregressive Models for Offline Reinforcement Learning

Michael R. Zhang, Tom Le Paine, Ofir Nachum, Cosmin Paduraru, George Tucker, Ziyu Wang, Mohammad Norouzi
International Conference on Learning Representations (ICLR) 2021

Benchmarks for Deep Off-Policy Evaluation

Justin Fu, Mohammad Norouzi, Ofir Nachum, George Tucker, Ziyu Wang, Alexander Novikov, Mengjiao Yang, Michael R. Zhang, Yutian Chen, Aviral Kumar, Cosmin Paduraru, Sergey Levine, Tom Le Paine
International Conference on Learning Representations (ICLR) 2021

Objective Social Choice: Using Auxiliary Information to Improve Voting Outcomes

Silviu Pitis, Michael R. Zhang
International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) 2020

Lookahead Optimizer: k steps forward, 1 step back

Michael R. Zhang, James Lucas, Geoffrey Hinton, Jimmy Ba
Neural Information Processing Systems (NeurIPS) 2019

Reverse Curriculum Generation for Reinforcement Learning

Carlos Florensa, David Held, Markus Wulfmeier, Michael Zhang, Pieter Abbeel
Conference on Robot Learning (CoRL) 2017

Probabilistically Safe Policy Transfer

David Held, Zoe McCarthy, Michael Zhang, Fred Shentu, Pieter Abbeel.
IEEE International Conference on Robotics and Automation (ICRA) 2017

PREPRINTS

Report Cards: Qualitative Evaluation of Language Models Using Natural Language Summaries

Blair Yang, Fuyang Cui, Keiran Paster, Jimmy Ba, Pashootan Vaezipoor, Silviu Pitis, Michael R. Zhang
Neurips 2024 Socially Responsible Language Modelling Research Workshop (spotlight)

Using Large Language Models for Hyperparameter Optimization

Michael R. Zhang, Nishkrit Desai, Juhan Bae, Jonathan Lorraine, Jimmy Ba
Neurips 2023 Foundation Models for Decision Making Workshop

Unlearnable Algorithms for In-context Learning

Andrei Muresanu, Anvith Thudi, Michael R. Zhang, Nicolas Papernot

Robustness to Adversarial Gradients: A Glimpse Into the Loss Landscape of Contrastive Pre-training

Phil Fradkin, Lazar Atanackovic, Michael R. Zhang
ICML 2022 Pre-training Workshop

TEACHING

Lecturing

University of Toronto CSC311: Introduction to Machine Learning: Winter 2023

Teaching Assistant

University of Toronto CSC311: Introduction to Machine Learning: Fall 2021 (head TA)

University of Toronto CSC2547: Introduction to Reinforcement Learning: Winter 2020

Vector Institute: Deep Learning Applications: Fall 2020

University of Toronto CSC421/2516: Neural Networks and Deep Learning: Winter 2019

University of Toronto CSC2541: Deep Reinforcement Learning: Fall 2018

UC Berkeley CS170: Introduction to Algorithms: Spring 2018

UC Berkeley CS189: Introduction to Machine Learning: Spring 2017, Fall 2017

UC Berkeley CS70: Discrete Mathematics and Probability: Spring 2016, Fall 2016

Mentoring:

Blair Yang (Undergraduate → Research Intern with Roger Grosse)

Fuyang Cui (Undergraduate → Research Intern with Sheila McIlraith)

Nishkrit Desai (Undergraduate → Research Intern with Chris Maddison)

Michael Yuan (Undergraduate → Clearpath)

Eric Wang (Undergraduate → AMD)

Michael Lizzi (Undergraduate → Ambi Robotics)

Zhiyu Liang (Undergraduate → MS at Yale)

SELECTED HONORS

Schwartz Reisman Fellow, 2023-2024

NSERC Canada Graduate Scholarship - Doctoral, 2021-2024

Ontario Graduate Scholarship, 2021 (respectfully declined)

National Science Foundation Graduate Research Fellowship, 2018 (respectfully declined)

Siebel Scholar, 2018

National award to 90 graduate students for outstanding academic performance and leadership

UC Berkeley Outstanding Graduate Student Instructor Award, 2017

Award given to the top 10% of teaching assistants based on faculty nominations and reviews

UC Berkeley Leadership Award

Merit-based scholarship for leadership impacting academic, work, and community

LEADERSHIP AND SERVICE

Paper reviewing ICLR (2020-23), ICML {2020,2022-23}, Neurips {2021-23}, TMLR

Toronto CS Graduate Application Assistance Program 2021 - Present

Established a student-led mentorship initiative which supports prospective graduate applicants through feedback and guidance, improving accessibility of the graduate admissions process.

AI Safety Reading Group Jan 2023 - Present

Organize reading group/seminar on topics in responsible and safe AI

Vector Institute Foundation Models for Science Oct 2024

Led breakout session

Self-Organizing Conference on Machine Learning Nov 2019

Led breakout session on deep learning and social chair

Her Code Camp Volunteer	Sep 2019
Robot Learning Lab Outreach	2017 - 2018
Coordinate lab outreach program to provide tours and robot demo to visitors of various ages	
Berkeley AI & AI4All Camp Volunteer	June 2017
Machine Learning at Berkeley	2017
Led technical consulting project to delight images with Unity Technologies	
Computer Science Mentors	2016
Develop materials and led small group tutoring sessions	