Remote, Canada

Toronto, Canada

(intermittent)

Sep. 2020-Apr. 2025

Sep. 2023-Feb. 2024

# Bryan Chan

 $\boxtimes$ bryan.chan@ualberta.ca<br/>  ${\ensuremath{\mathscr{O}}}$  chanb.github.io  $% \ensuremath{\mathsf{in}}$  chanb<br/>py  $\ensuremath{\mathbb{Q}}$  chanb

# Education

University of Alberta PhD in Computing Science, Statistical Machine Learning	2022 - Present
$\circ$ Supervisory Committee: Dale Schuurmans (supervisor), Csaba Szepesvári, and András György	
$\circ$ Research Interests: In-context learning and reinforcement learning (RL)	
University of Toronto MSc in Applied Computing, Computer Science	2018 – 2019
• Supervisors: Florian Shkurti, Animesh Garg, and James Bergstra	
$\circ$ Research: Average Reward Reinforcement Learning for System Optimization in Robotics Application $\swarrow$	
University of Toronto HBSc in Computer Science	2013 – 2018
$\circ$ Graduated with high distinction, specializing in software engineering (Co-op)	

# Experience

Machine Learning Researcher/Research Intern	Toronto, Canada
Kindred AI (Acquired by Ocado Technology)	May 2019–Oct. 2024
	(intermittent)

- $\circ\,$  Developed a novel average-reward RL algorithm to optimize a robotic bin-picking system, improving target KPI by 30 units per hour on held-out data
- $\circ~$  Designed and deployed RL pipelines to train  ${\sim}30$  production robots in parallel, with safeguards implemented via bandits to prevent performance degradation
- Built teleoperation systems for data collection and human-in-the-loop control to recover from failures; currently used in production deployment
- $\circ\,$  Built simulation frameworks to test manipulation strategies and hardware designs, enabling rapid experimentation by cross-functional teams
- $\circ\,$  Full-time employment from April 2020–June 2022

## Student Researcher

 $Google \ DeepMind$ 

- Studied sample efficiency of RT-X variants via CLIP-based demonstration diversity metrics
- Scaled training transformer models on TPU clusters using Jax and FSDP for distributed inference optimization

# Sessional Instructor

University of Toronto

 Designed and taught ML and AI courses (120+ students/term), covering classical ML, deep learning, and decision-making under uncertainty

# Selected Publications/Projects

Please see Google Scholar 🗹 for the complete list of publications and GitHub 🗹 for more open-source projects

## Toward Understanding In-context vs. In-weight Learning

International Conference on Learning Representations, 2025

Theoretical and empirical analysis of when transformers and LLMs exhibit emergent in-context learning

## JaxL 🗹

A modular framework in Jax that unifies training ML/RL models; supporting scalable, composable pipelines for academic and applied research.

# RL Sandbox 🗹

PyTorch-based toolkit for fast development and ablation of RL/IL algorithms in discrete/continuous environments; used in several publications.

# Technical Skills

Languages/Frameworks/Infra.: Python, C, Go, Jax, PyTorch, TensorFlow, FSDP, Docker, GCP ML Techniques: Classical ML, online/offline RL, imitation learning, in-context learning, ML theory