

# Bryan Chan

✉ bryan.chan@ualberta.ca    🌐 chanb.github.io    in chanbpy    📄 chanb

## Education

---

**University of Alberta** *PhD in Computing Science, Statistical Machine Learning* 2022–Present

- Supervisory Committee: Dale Schuurmans (supervisor), Csaba Szepesvári, and András György
- 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
- Research: [Average Reward Reinforcement Learning for System Optimization in Robotics Application](#) [🔗](#)

**University of Toronto** *HBSc in Computer Science* 2013–2018

- 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)*

- 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
- Designed and deployed RL pipelines to train ~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 to develop imitation-learning (IL) models for production deployment
- Built simulation frameworks to test manipulation strategies and hardware designs, enabling rapid experimentation by cross-functional teams
- Full-time employment from April 2020–June 2022

**Student Researcher** *Remote, Canada*  
*Google DeepMind* *Sep. 2023–Feb. 2024*

- 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** *Toronto, Canada*  
*University of Toronto* *Sep. 2020–Apr. 2025*  
*(intermittent)*

- Designed and taught ML and AI courses (120+ students/term), covering classical ML, deep learning, and decision-making under uncertainty
- Managed multiple teaching assistants to effectively create new assignments, tests, and exams; and delivering hands-on tutorials on AI applications including transformers for LLMs and diffusion for image generation

**Programmer Analyst Intern** *Toronto, Canada*  
*Ontario Teachers' Pension Plan* *Sep. 2015–Aug. 2017*  
*(intermittent)*

- Standardized automatic web service deployment pipelines across the IT department, and modularizing and constructing reusable templates for users, thereby reducing the deployment time from few hours to minutes
- Developed RESTful and SOAP-based web services using .NET to retrieve complex business data for internal clients
- Built a web application using .NET and Angular for managers to monitor real-time status of services in different environments

## Publications/Projects

---

Please see [Google Scholar](#) and [GitHub](#) for more projects

### [Toward Understanding In-context vs. In-weight Learning](#)

**B. Chan**<sup>\*</sup>, X. Chen<sup>\*</sup>, A. György, D. Schuurmans

*International Conference on Learning Representations (ICLR)*, 2025 (\* equal contribution)

### [AVG-DICE: Stationary Distribution Correction by Regression](#)

F. Che, **B. Chan**, C. Ma, A. R. Mahmood

*Reinforcement Learning Conference (RLC)*, 2025

### [Efficient Imitation Without Demonstrations via Value-Penalized Auxiliary Control from Examples](#)

T. Ablett, **B.Chan**, H. Wang, and J. Kelly

*IEEE International Conference on Robotics and Automation*, 2025

### [Offline-to-online Reinforcement Learning for Image-based Grasping with Scarce Demonstrations](#)

**B. Chan**, A. Leung, J. Bergstra

*CoRL Workshop on Mastering Robot Manipulation in a World of Abundant Data*, 2024

### [Mitigating the Curse of Horizon in Monte-Carlo Returns](#)

A. Ayoub<sup>\*</sup>, D. Szepesvari<sup>\*</sup>, F. Zanini<sup>\*</sup>, **B. Chan**<sup>\*</sup>, D. Gupta, D. Schuurmans, B. C. da Silva

*Reinforcement Learning Conference*, 2024 (\* equal contribution)

### [A Statistical Guarantee for Representation Transfer in Multitask Imitation Learning](#)

**B. Chan**, K. Pereida<sup>\*</sup>, J. Bergstra<sup>\*</sup>

*NeurIPS Workshop on Robot Learning*, 2023 (\* equal advising)

### [Learning from Guided Play: Improving Exploration in Adversarial Imitation Learning with Simple Auxiliary Tasks](#)

T. Ablett, **B.Chan**, and J. Kelly

*IEEE Robotics and Automation Letters* (with IROS), 2023

### [Heteroscedastic Uncertainty for Robust Generative Latent Dynamics](#)

O. Limoyo, **B. Chan**, F. Marić, B. Wagstaff, A. R. Mahmood, and J. Kelly

*IEEE Robotics and Automation Letters* (with IROS), 2020

### [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, AWS

**ML Techniques:** Classical ML, online/offline RL, IL, in-context learning, ML theory