Harley Wiltzer

harley.wiltzer@mail.mcgill.ca • https://harwiltz.github.io • +1 (514) 208-6970

Education

2022-Present **Mila / McGill University** — Montreal, Canada
Ph.D in Computer Science
Supervisors: Marc G. Bellemare and David Meger
GPA: 4.0 / 4.0

2019-2021 **McGill University / Mila** — Montreal, Canada
M.Sc in Computer Science
Supervisors: David Meger and Marc G. Bellemare
Dissertation: *On the Evolution of Return Distributions in Continuous-Time Reinforcement Learning*
GPA: 4.0 / 4.0

2015-2019 **McGill University** — Montreal, Canada
B.Eng in Computer Engineering
Capstone Advisor: Shane McIntosh
GPA: 3.96 / 4.0

Publications

2025 **Convergence Theorems for Entropy-Regularized and Distributional Reinforcement Learning**
Yash Jhaveri*, Harley Wiltzer*, Patrick Shafto, Marc G. Bellemare, David Meger
Advances in Neural Information Processing Systems (NeurIPS)

2025 **Non-Adversarial Inverse Reinforcement Learning via Successor Feature Matching**
Arnav Kumar Jain, Harley Wiltzer, Jesse Farebrother, Irina Rish, Glen Berseth, Sanjiban Choudhury
International Conference on Learning Representations (ICLR)

2025 **Tractable Representations for Convergent Approximations of Distributional HJB Equations**
Julie Alhosh, Harley Wiltzer, David Meger
Reinforcement Learning and Decision Making (RLDM)

2024 **Foundations of Multivariate Distributional Reinforcement Learning**
Harley Wiltzer, Jesse Farebrother, Arthur Gretton, Mark Rowland
Advances in Neural Information Processing Systems (NeurIPS)

2024 **Action Gaps and Advantages in Continuous-Time Distributional Reinforcement Learning**
Harley Wiltzer*, Marc G. Bellemare, David Meger, Patrick Shafto, Yash Jhaveri*
Advances in Neural Information Processing Systems (NeurIPS)

2024 **Simplifying Constraint Inference with Inverse Reinforcement Learning**
Adriana Hugessen, Harley Wiltzer, Glen Berseth
Advances in Neural Information Processing Systems (NeurIPS)
Reinforcement Learning Safety Workshop—RLC

2024 **A Distributional Analogue to the Successor Representation**
Harley Wiltzer*, Jesse Farebrother*, Arthur Gretton, Yunhao Tang, André Barreto, Will Dabney, Marc G. Bellemare, Mark Rowland
International Conference on Machine Learning (ICML) *(Spotlight, top 3.5%)*
European Workshop on Reinforcement Learning (EWRL)

2023 **Policy Optimization in a Noisy Neighborhood: On Return Landscapes in Continuous Control**
Nate Rahn*, Pierluca D'Oro*, Harley Wiltzer, Pierre-Luc Bacon, Marc G. Bellemare
Advances in Neural Information Processing Systems (NeurIPS)

2022 **Distributional Hamilton-Jacobi-Bellman Equations for Continuous-Time Reinforcement Learning**
Harley Wiltzer, David Meger, Marc G. Bellemare
International Conference on Machine Learning (ICML) *(Spotlight)*

Positions Held

2023 **Microsoft Research**, *Research Intern* New York, USA
Division: *Real-World Reinforcement Learning*
Responsibilities: *Developed methods for 3D scene reconstruction from noisy monocular camera sensors using vision-based foundation models, as well as geometric latent state planning.*

2019 **Amazon Web Services**, *Software Development Engineer Intern* Vancouver, Canada
Division: *AWS Auto Scaling*
Responsibilities: *Designed a service for AWS Auto Scaling that monitors the health of the server fleets hosting thousands of AWS services (e.g., Amazon EC2, AWS DynamoDB) in each AWS region. Invented and implemented machine learning models and signal processing algorithms to detect anomalies in Auto Scaling's regional time series data. The systems I designed are still in use at AWS as of 2025, as part of their Predictive Auto Scaling service.*

2018 **Amazon**, *Software Development Engineer Intern* Vancouver, Canada
Division: *Amazon Wallet*
Responsibilities: *Designed and implemented a state of the art system for optimizing the resolution of BIN-derived payment properties without access to security-critical credit card data. Developed this system and extensive automated tests singlehandedly.*

2017 **Micro Focus**, *QA Engineer Intern* Montreal, Canada
Division: *Micro Focus Retain*
Responsibilities: *Scrum Master for a core development team working on the Retain unified archiving system. Designed and carried out tests for the Retain software suite, identified several critical security vulnerabilities that had gone unnoticed for 3 or more years.*

Talks

2025 **The RL Sofa** (Slides) Mila—Quebec AI Institute

2025 **Shafto Lab** (Slides) Rutgers University

2024 **Mobile Robotics Mega Meeting** (Slides) McGill University

2023 **Risk-Aware RL Workshop** (Slides) Université de Montréal

2023 **Reinforcement Learning Discussion Group** (Slides) Microsoft Research, NYC

2023 **Reinforcement Learning Theory Workshop** (Recording, Slides) University of Alberta

2022 **The RL Sofa** (Slides) Mila—Quebec AI Institute

Honors and Scholarships

2025 **TMLR Expert Reviewer**, *Transactions on Machine Learning Research*

2023 **Doctoral Research Scholarship**, *Fonds de recherche du Québec (FRQNT)*

2023 **Postgraduate Scholarhip—Doctoral**, *NSERC*

2020 **Alexander Graham Bell Scholarship**, *NSERC*

2019 **British Association Medal**, *McGill University*

2018 **Peter P. Sebestyen Award**, *McGill University*

2017 **W.G. McBride Scholarship**, *McGill University*

2016 **Brian Cullen Award**, *McGill University*

2016 **Douglas Macaulay Scholarship**, *McGill University*

2015 **J.W. McConnell Scholarship**, *McGill University*

Service

I served as a reviewer or referee for the following conference and journals:

2025 Annals of Statistics, AISTATS, ICML, Neural Computation, NeurIPS, RLC (Senior Reviewer, Technical Reviewer), TMLR (Expert Reviewer)

2024 ICLR, ICML (Best Reviewer Award), ICRA, JMLR, NeurIPS (Reviewer Award), RLC, TMLR

2023 AISTATS, Artificial Intelligence Journal, CoLLAs, ICLR, ICML, JMLR, NeurIPS (Reviewer Award)

2022 ICML, NeurIPS (Reviewer Award)

Teaching Experience

2022 **Teaching Assistant**, *McGill University, School of Computer Science*
Course: `COMP579`, *Reinforcement Learning*
Responsibilities: *Graded exams and assignments, held office hours, and led tutorial sessions.*

2022 **Teaching Assistant**, *Polytechnique Montréal*
Course: `INF8250e`, *Reinforcement Learning*
Responsibilities: *Designed the assignment on Deep RL for continuous control with PyTorch, graded exams and assignments, held office hours, and let tutorial sessions.*

2022 **Teaching Assistant**, *McGill University, School of Computer Science*
Course: `COMP551`, *Applied Machine Learning*
Responsibilities: *Graded exams and assignments, held office hours, and led tutorial sessions.*

2021 **Teaching Assistant**, *McGill University, School of Computer Science*
Course: `COMP551`, *Applied Machine Learning*
Responsibilities: *Graded exams and assignments, held office hours, led tutorial sessions, and helped design the final course project.*

Contributions to Open Source

2021-2022 **JuliaReinforcementLearning** https://juliareinforcementlearning.org
Contributions: *Implementation of reinforcement learning algorithms and library methods, bug fixes, and documentation.*

2020-2021 **Jax** https://github.com/google/jax
Contributions: *Identified a number of bugs that cause memory leaks, JIT cache misses, and inconsistent behavior in JIT-compiled code.*

2021 **Nixpkgs** https://github.com/NixOS/nixpkgs
Contributions: *Implemented package derivations for the Nix package manager to provide the* `dm_haiku` *Jax neural network Python package, as well as related packages in the Jax ecosystem, to Nix users for truly reproducible software builds.*

2020 **K-9 Mail** https://github.com/DestructiveReasoning/k-9
Contributions: *Implemented support for XOAUTH2 authentication for Office365 email accounts in the popular free / open source K-9 Mail Android email client.*

2018-2020 **Gentoo Linux** https://github.com/harwiltz/gentoo
Contributions: *Maintenance of a collection of software packages for the Gentoo Linux distribution, including packages that provide support for AMD ROCm drivers for relatively old GPUs, as well as machine learning libraries built with support for these drivers.*

2019 **SageRank** https://github.com/harwiltz/SageRank
Contributions: *Designed and implemented a recommendation engine for research papers, based on the library of papers accumulated by a user.*

Technical Skills

Programming Languages: *Python, Julia, Scala, Java, Haskell, Ruby, GNU Octave / Matlab, C, C++, Rust, Javascript, Bash*

ML Libraries: *Jax, PyTorch, Flux.jl, ReinforcementLearning.jl*

Misc. Software: *git, Slurm, Nix, Docker, GNU/Linux*

Portfolio: https://github.com/harwiltz, https://github.com/DestructiveReasoning

Languages

English: *Fluent*

French: *Competent*