

Kenny Young

kjyoung@ualberta.ca | kenjyoung.github.io | [Github](#) | [Google Scholar](#)

Education

PhD candidate with Richard Sutton in Computing Science, University of Alberta

Expected graduation date: April 2024

Obtained MSc Computing Science, University of Alberta, 2019 (4.0 GPA)

Obtained BSc Honors Computing Science, University of Alberta, 2016 (4.0 GPA)

Obtained BSc Honors Mathematical Physics, University of Alberta, 2014 (3.9 GPA)

Relevant Work Experience

Machine Learning Intern at Borealis AI Vancouver, May-September 2020

Supervisors: Mohammad Ahmed and Andreas Lehrmann

Explored and developed methods for obtaining calibrated uncertainty estimates in deep neural networks.

Subject Matter Expert at Coursera, January-August 2019

Supervisors: Martha White and Adam White

Prepared scripts for the University of Alberta reinforcement learning specialization (<https://www.coursera.org/specializations/reinforcement-learning>).

Machine Learning Researcher at Borealis AI Edmonton, January-April 2018

Supervisor: Matt Taylor

Primarily worked on meta-learning based hyperparameter tuning for reinforcement learning control. I presented a paper on this work at IJCAI 2019.

Research Assistant at University of Alberta, Summer 2015

Supervisor: Ryan Hayward

Primarily developed a C++ program to efficiently solve positions in a variant of the board game Hex. I presented a paper on this work at the International Conference on Computers and Games 2016.

Research Assistant at University of Alberta, Summer 2013

Supervisors: Carsten Krauss and Kalpana Singh

Programmed camera control, computer vision and database interface software for the Sudbury Neutrino Observatory calibration camera system.

Research Philosophy

I'm broadly interested in uncovering the general principles behind intelligent behaviour and figuring out how to implement them algorithmically. To this end, I am currently particularly interested in exploring the potential of model-based reinforcement learning and planning to allow agents to synthesize their experiences into novel predictions in order to efficiently learn how to act in the world. I am fascinated by recent advances in large language models but believe reinforcement learning and planning remain the most promising general techniques to enable agents to autonomously self-improve beyond what is possible by imitation of human-generated data and direct human feedback.

Relevant Software Projects

Github page: <https://github.com/kenjyoung>

MinAtar: Miniaturized ATARI Games for Reinforcement Learning Experiments

(<https://github.com/kenjyoung/MinAtar>)

AlphaZero style learning on procedurally generated mazes using the MCTX framework

(https://github.com/kenjyoung/mctx_learning_demo)

DreamerV2 implemented in JAX with support for multiple random seeds on one GPU

(https://github.com/kenjyoung/dreamerv2_JAX)

Learning to play Hex with deep Q-learning

(<https://github.com/kenjyoung/Neurohex>)

Programming and Software

Python, JAX, PyTorch, TensorFlow, theano, C++, git, LaTeX

Publications

Google Scholar Page: <https://scholar.google.com/citations?user=zI2uHi8AAAAJ>

Kenny Young, Aditya Ramesh, Louis Kirsch, Jürgen Schmidhuber, *The Benefits of Model-Based Generalization in Reinforcement Learning*. International Conference on Machine Learning (2023). <https://arxiv.org/abs/2211.02222>

Tian Tian, **Kenny Young**, Richard S. Sutton, *Doubly-Asynchronous Value Iteration: Making Value Iteration Asynchronous in Actions*. Conference on Neural Information Processing Systems (2022). <https://arxiv.org/abs/2207.01613>

Kenny Young, *Hindsight Network Credit Assignment: Efficient Credit Assignment in Networks of Discrete Stochastic Units*. Association for the Advancement of Artificial Intelligence Conference (2022). <https://ojs.aaai.org/index.php/AAAI/article/view/20874>

Kenny Young, Baoxiang Wang, Matthew E. Taylor, *Metatrace Actor-Critic: Online Step-Size Tuning by Meta-gradient Descent for Reinforcement Learning Control*. International Joint Conferences on Artificial Intelligence (2019). <https://arxiv.org/abs/1805.04514>

Kenny Young, Tian Tian, *MinAtar: An Atari-Inspired Testbed for Thorough and Reproducible Reinforcement Learning Experiments*. The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (2019). <https://arxiv.org/abs/1903.03176>

Kenny Young, Richard S. Sutton, *Learning What to Remember with Online Policy Gradient Over a Reservoir*. NeurIPS Workshop on Reinforcement Learning under Partial Observability (2018). https://www.ias.informatik.tu-darmstadt.de/uploads/Team/JoniPajarinen/RLPO2018_paper_7.pdf

Craig Sherstan, Brendan Bennett, Dylan R. Ashley, **Kenny Young**, Adam White, Martha White, Richard S. Sutton, *Comparing Direct and Indirect Temporal-Difference Methods for Estimating the Variance of the Return*. Conference on Uncertainty in Artificial Intelligence (2018). <http://auai.org/uai2018/proceedings/papers/35.pdf>

Craig Sherstan, **Kenny Young**, Dylan Ashley, Brendan Bennett, Adam White, Martha White, Rich Sutton, *Direct Estimation of the Variance of the λ -Return with Temporal-Difference Methods*. The Multi-disciplinary Conference on Reinforcement Learning and Decision Making (2017). <https://arxiv.org/abs/1801.08287>

Kenny Young, Ryan Hayward, Gautham Vasan, *Neurohex: A Deep Q-learning Hex Agent*. IJCAI Computer Games Workshop (2016). <https://arxiv.org/abs/1604.07097>

Kenny Young, Ryan Hayward, *A Reverse Hex Solver*. International Computer Games Association conference (2016). <https://arxiv.org/abs/1707.00627>

Unpublished Work

Kenny Young, Richard S. Sutton, *Iterative Option Discovery for Planning, by Planning*. arXiv preprint arXiv:2310.01569 (2023). <https://arxiv.org/abs/2310.01569>

Kenny Young, Richard S. Sutton, *Understanding the Pathologies of Approximate Policy Evaluation when Combined with Greedification in Reinforcement Learning*. arXiv preprint arXiv:2010.15268 (2020). <https://arxiv.org/abs/2010.15268>

Kenny Young, *Variance Reduced Advantage Estimation with δ Hindsight Credit Assignment*. arXiv preprint arXiv:1911.08362 (2019). <https://arxiv.org/abs/1911.08362>

Kenny Young, *Learning What to Remember: Strategies for Selective External Memory in Online Reinforcement Learning Agents*. Master's Thesis (2019). <https://era.library.ualberta.ca/items/7b9ed64d-9fc6-4cce-9ed7-d4bd8f91892d>

**Select
Courses
Completed**

Reinforcement Learning Theory:

Instructor: Csaba Szepesvari

Mathematical analysis of the reinforcement learning problem and algorithms.

Stochastic Approximation:

Instructors: Ajin Joseph and Martha White

Mathematical foundations for analysis of many machine learning algorithms.

Deep Learning:

Instructor: Dale Schuurmans

Developments in deep learning and insights from the instructor's own research.

Bandit Algorithms:

Instructor: Csaba Szepesvari

Theoretical analysis and application of algorithms for bandit problems.

**Academic
Honors**

Andrew Stewart Memorial Graduate Prize, 2023

Alberta Graduate Excellence Scholarship, 2022

Alberta Innovates Graduate Student Scholarship, 2019

NSERC-Canada Graduate Scholarship-Doctoral, 2019

NSERC-Canada Graduate Scholarship-Master's, 2017

Queen Elizabeth II Graduate Scholarship, 2017

NSERC-Undergraduate Student Research Award, 2016

NSERC-Undergraduate Student Research Award, 2015

NSERC-Undergraduate Student Research Award, 2013

**Reviewing
Work**

NeurIPS (2019, 2020, 2021)

ICML (2020, 2021)

ICLR (2020, 2023)

AAAI (2020, 2021)

Hobbies

Rock Climbing

Skiing

Hiking