

AYUSH JAIN

Montreal, Canada

+1 438 871 2395 \diamond ayush.jain@mail.mcgill.ca \diamond jainayus@mila.quebec \diamond Website

RESEARCH INTERESTS

- Reinforcement Learning (RL)
- Temporal Abstraction and Hierarchical Reinforcement Learning
- Continual Learning
- Skill-learning
- Safety in AI
- Off-Policy Learning
- Eligibility Traces

EDUCATION

McGill University, Montreal

Ph.D., Computer Science

Sept 2020 - *

Master of Science, Computer Science

GPA: 3.80/4.00

Sept 2015 - Dec 2017

Mila

Supervisor: *Doina Precup*

NSIT, University of Delhi, India

August 2008 - June 2012

Bachelor Of Engineering, Computer Engineering

Overall Percentage: 78.88/100

Ranked among **top 5** in department of 120+ students

Approx. GPA: 3.82/4.00

Supervisor: Satish Chand

PUBLICATIONS

1. **Ayush Jain** and Doina Precup. 2018. **Eligibility Traces for Options**. In Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (**AAMAS 2018**), Stockholm, Sweden, July 10-15, 2018, IFAAMAS, 9 pages.
2. Ayush Jain. 2018. **Eligibility Traces for Options**. Master's thesis, McGill University, 2018.
3. Arushi Jain, Gandharv Patil, **Ayush Jain**, Khimya Khetarpal and Doina Precup. 2021. **Variance Penalized On-Policy and Off-Policy Actor-Critic**. In Proceedings of the 35th **AAAI** Conference on Artificial Intelligence, held virtually, Feb 2-9, 2021, pp. 7899-7907.
4. Arushi Jain, **Ayush Jain** and Doina Precup. 2018. **Safe Actor-Critic**[Paper][Poster][Slide] Accepted in Safety, Risk and Uncertainty in RL, Uncertainty in AI (**UAI 2018**) Workshop. Accepted in Women in ML (WiML), (**NeurIPS 2018**) Workshop.

RESEARCH AND WORK EXPERIENCES

Mila

Sept 2020 - *

Graduate Student

- Working towards developing a hierarchical skill-based solution that evolves continually and transfers knowledge from one task to another, from one or more domains.

SPORTLOGiQ

August 2017 - August 2020

Machine Learning Researcher

- Developed a model checking tool and an explainable MDP for Hockey with factored representation. Incorporating minimum sufficient explanations into model predictions.
- Developed a model to accurately predict events in Soccer from partially observed broadcast feed of a game.
- Modelled hockey games as MDP to devise sound metrics for the game. Learned policies- how teams play in different scenarios in NHL through imitation learning. Learned unique value functions by rewarding defensive/offensive play, saving goal, scoring etc. Goal is to learn intrinsic reward function of each team.

- Developed model to predict odds of a team winning in a game based on historical game reports and expected team performances.

Reasoning and Learning Lab, McGill University
Research Assistant advised by Prof. Doina Precup

February 2016 - August 2017

- Worked on temporal abstraction and hierarchical reinforcement learning.
- Thesis: tested the utility of eligibility traces with options and searched for good ways of doing multi-step intra-option updates. Proposed three algorithms, based on off-policy methods - importance sampling, tree backup and $\text{retrace}(\lambda)$, for using eligibility traces with options.

Adobe Systems
Computer Scientist

July 2012 - August 2015

- Released major version of Adobe Drive, a cloud based digital-asset synchronization service with version control. I was responsible for functionality allowing virtual collections of cloud assets; major improvements to Drive's kernel extension; and platform support for Windows 8, OSX Yosemite and El Capitan releases.
- Worked on a page description languages translation engine - PDFtoPS. I worked on different font technologies, vector imaging and color management and shipped major features in font handling, imaging, and reliability improvements.

National Informatics Centre, Govt. of India
Research Intern advised by Shashi Kant Sharma

May 2011 - July 2011

- Worked on fingerprint identification methods. This resulted in a library for web-based fingerprint enrollment, feature extraction, and matching algorithm which ensures interoperability of devices and long term storage of data with technology independence.

NSIT, University of Delhi
Summer Research Intern advised by Dr. Anand Gupta

June 2010 - August 2010

- Worked on detection of pedestrians and moving objects. I investigated performance of a model which combined body part detectors (for heads and legs) to detect total number of pedestrians in a video feed.

TEACHING EXPERIENCES

Teaching Assistant, McGill University

Courses: COMP 551 Applied Machine Learning; COMP 202 Fundamentals of Programming; COMP 208 Computers in Engineering

RESEARCH PROJECTS

- **Safe Actor Critic (2018):** Designed a novel, generic **on-policy and off-policy safe actor-critic** framework to learn safe policies that computes a direct estimate of the variance in the return which meets certain safety requirements apart from the performance demands.
- **Eligibility Trace based Methods to Supplement Backpropagation Through Time in RNNs (2016):** Experimented on a recurrent neural network which used BPTT on a temporal record over latent variables. Our experiment confirmed theoretical understanding - traces help distribute credit to a greater depth leading to faster convergence.
- **Learning Better Word Embeddings with Morphological Knowledge (2015):** Experimented on a neural network which simultaneously learns word and morpheme embeddings - capturing the explicit relationship among morphemes. I was able to derive better representation for rare morphological derivatives of words and word embeddings learnt performed better on analogical task.

- **A Generalized Architecture for EEG Data Analysis (2015):** Deploying machine learning techniques on EEG data requires a highly tuned architecture specific to task at hand. There is no common architecture to handle all EEG data irrespective of the task. We experimented on designing a generalized, task-insensitive architecture for EEG data. Used deep learning techniques to propose generic frameworks for statistical EEG data set, which don't require any hand crafted nor domain specific features for EEG classification.
- **Self Learning in Games using TD(λ) (2012):** Developed a self-learning AI Othello agent using NN to approximate large state-action pairs with TD(λ) updates. Studied the effectiveness of merging TD(λ) and Co-Evolutionary Learning.
- **Automated Car Number Plate Recognition System (2011):** Developed an automated license plate recognition system in Indian scenario. I trained a pattern recognition system as plate and non-plate regions using SVMs.

TECHNICAL SKILLS

Programming Languages
Tool & Frameworks

Python, C, Java
Pytorch, L^AT_EX