Sriram Ganapathi Subramanian

sriramsubramanian@cunet.carleton.ca, s2ganapa@uwaterloo.ca
Research Website, LinkedIn, GitHub, Bitbucket, Google Scholar

EDUCATION

University of Waterloo, Waterloo, Ontario, Canada

- Doctor of Philosophy (PhD) in Electrical and Computer Engineering — Sep 2018 – Jul 2022
  - Research Thesis: "Multi-Agent Reinforcement Learning in Large Complex Environments"
  - **Canada AI Association's Best Doctoral Dissertation Award, 2023**
  - Supervisors: **Mark Crowley**, Department of Electrical and Computer Engineering and **Kate Larson**, Cheriton School of Computer Science
  - Focus: Multi-Agent Reinforcement Learning, Game Theory, and Reinforcement Learning
- Master of Applied Science (MASc) in Electrical and Computer Engineering — Sep 2016 – Apr 2018
  - Research Thesis: "Reinforcement Learning for Determining Spread Dynamics of Spatially Spreading Processes with Emphasis on Forest Fires"
  - Supervisor: **Mark Crowley**, Department of Electrical and Computer Engineering
  - Focus: Reinforcement Learning, Deep Learning, and Image Processing

College Of Engineering Guindy, Anna University, Chennai, Tamil Nadu, India

- Bachelor of Engineering in Geo-Informatics — Jul 2012 – Jul 2016
  - **Governor's Medal for Academic Excellence**
  - Research Thesis: "Cartographic View of Cancer"
  - Supervisor: **S.Jayalakshmi**, Department of Remote Sensing
  - Focus: Remote Sensing, GIS, and Algorithms

RESEARCH EXPERIENCE

Carleton University, Ottawa, Ontario, Canada — Jul 2025 – Present

- **Assistant Professor** at the School of Computer Science & **Canada Research Chair** (Tier II) in Artificial Intelligence.
  - I am just starting my research lab at Carleton where the focus will be on Multi-agent Systems, Reinforcement Learning, Deep Learning, AI Safety, and Generative AI.
  - I am affiliated with the **Vector Institute for Artificial Intelligence**, Toronto as a **Faculty Affiliate**. I am also affiliated with the **Schwartz Reisman Institute for Technology and Society**, Toronto as a **Faculty Affiliate**.

Vector Institute, Toronto, Ontario, Canada — Sep 2022 – Jun 2025

- Distinguished Postdoctoral Fellow
  - I conducted fundamental research on several sub-fields within reinforcement learning and deep learning. Particularly, I contributed to the fields of Inverse Constraint Reinforcement Learning (ICRL), Multi-agent Reinforcement Learning (MARL), Reinforcement Learning (RL) in non-Markovian environments, and Reinforcement Learning applications to Chemistry and Drug Discovery. My contributions included providing a confidence-aware algorithmic framework for ICRL, relaxing neighbourhood limitations in Mean Field multi-agent algorithms, establishing principled theoretical properties for RL algorithms in non-Markovian environments and environments containing humans-in-the-loop. I explored applications of RL in autonomous driving, robotics, and scientific discovery. During the last year of my postdoc, I served as an advisor/consultant to AI startups (Skyfall, ArenaX Labs) in a part-time capacity, and spent 50% of my time at the AI Engineering team at Vector on projects relevant to Vector sponsors (CIBC, RBC, Air Canada, and Loblaws). Furthermore, I served as a course instructor for one course at the University of Waterloo and conducted (as a co-instructor) an RL bootcamp for employees of Vector Sponsors (CIBC, BMO, RBC, Scotia Bank, Air Canada, and Loblaws).
  - Supervisors: **Pascal Poupart**, Faculty Member at the Vector Institute and Professor at the University of Waterloo. **Sheila McIlraith**, Faculty Member at the Vector Institute and Professor at the University of Toronto. 50% of time in the last year of my postdoc was spent in the AI Engineering team supervised by **Deval Pandya**, Vice President of AI Engineering.
  - Focus: Multi-Agent Reinforcement Learning, Reinforcement Learning, and Game Theory

Machine Learning Lab, University Of Waterloo, Waterloo, Canada — Sep 2018 – Jul 2022

- PhD Student, Electrical and Computer Engineering Department
  - I considered two fundamental problems in multi-agent reinforcement learning (MARL) that prevent its wide application in complex real-world problems. The first is the issue of sample efficiency, and the second is the issue of scaling MARL to large environments. My contributions included relaxing a number of limiting assumptions in prior work in this area and providing new algorithms that are more applicable to real-world environments. Further, I designed and open-sourced a Gym-style environment for training autonomous agents to run chemistry laboratories.
  - Supervisors: **Mark Crowley**, Associate Professor, University of Waterloo, and **Kate Larson**, Professor, University of Waterloo
  - Additional Collaborators: **Pascal Poupart**, Professor, University of Waterloo, **Matthew E. Taylor**, Associate Professor, University of Alberta, **Isaac Tamblyn**, Associate Professor, University of Ottawa, and **Colin Bellinger**, Research Officer, National Research Council (NRC) of Canada

- Focus: Multi-Agent Reinforcement Learning, Reinforcement Learning, and Game Theory

Borealis AI (Royal Bank of Canada), Waterloo, Canada — Sep 2018 – May 2019

- Machine Learning Researcher (Part-time)
  - Project: Mean field Multi-Agent Reinforcement Learning Applied to Dynamic Financial Environments.
  - This project involved developing highly scalable reinforcement learning algorithms for financial applications. My role included designing the learning algorithms, implementing and testing them in open-source and proprietary test beds, and deriving necessary theoretical guarantees to prove that the algorithms are principled.
  - Supervisors: **Pascal Poupart**, Research Lab Director, Borealis AI, Waterloo and **Matthew E. Taylor**, Research Lab Director, Borealis AI, Edmonton
  - Focus: Multi-Agent Reinforcement Learning, Deep Learning, and Fintech

Borealis AI (Royal Bank of Canada), Edmonton, Canada — May 2018 – Aug 2018

- Research Intern - Machine Learning
  - Project: Stock Market Trading using Reinforcement Learning.
  - This project involved exploring the applicability of reinforcement learning (RL) for stock trading. I was tasked with modeling the stock trading scenario as a Markov decision process, implementing and testing a number of state-of-the-art RL algorithms, and designing a novel algorithm for this context.
  - Applied Research Team
  - Supervisor: **Nidhi Hegde**, Applied Research Team Lead, Borealis AI, Edmonton
  - Focus: Reinforcement Learning and Fintech

Machine Learning Lab, University Of Waterloo, Waterloo, Canada — Sep 2016 – Apr 2018

- Master's Student, Electrical and Computer Engineering Department
  - In this research, I bench-marked existing reinforcement learning algorithms on a novel domain pertaining to wildfire prediction using real-world datasets. Additionally, I designed a new reinforcement learning algorithm that combines the advantages of existing algorithms and outperforms the state-of-the-art physics-based simulators in the wildfire domain.
  - Supervisor: **Mark Crowley**, Associate Professor, University of Waterloo
  - Focus: Reinforcement Learning, Deep Learning, and Image Processing

Denso International America, Detroit, United States and **University of Waterloo**, Waterloo, Canada

- Researcher (at University of Waterloo) — Apr 2017 – Jan 2018
  - Project: Predictive Collaborative Automated Drive for Crash Avoidance.
  - This project involves enabling an ensemble of self-driving cars to learn predictive collaborative driving alongside human-driven cars for navigation and crash avoidance. My role in the project involved using RL algorithms like Deep Q-Learning to make the learning modules for the vehicles. The project was showcased as a demonstrative display at the Consumer Electronics Show (CES)-2018 in Las Vegas (from 8th January to 12th January).
  - Supervisors: **Sebastian Fischmeister**, Associate Professor, University of Waterloo
    William Melek, Professor, University of Waterloo
    Mark Crowley, Assistant Professor, University of Waterloo
  - Focus: Reinforcement Learning and Autonomous Driving

Wilfrid Laurier University, Waterloo, Canada — May 2017 – Aug 2017

- Research Assistant
  - Project: Spatial Pattern Comparison Statistics.
  - This project involved the research and development of various spatial pattern comparison statistics as a package in R. My role consisted of implementing the package and reproducing results from the corresponding research publications.
  - Supervisor: **Colin Robertson**, Associate Professor, Wilfrid Laurier University
  - Focus: Spatial Statistics, Pattern Comparison, and Geo-computation

Wilfrid Laurier University, Waterloo, Canada — Apr 2015 – Aug 2015

- Research Intern (Mitacs Globalink)
  - Project: Geo+Social Analytics for Healthy Urban Environments
  - This project involved finding important influences of the spatial environment a person interacts with, on his/her mental health. My role comprised implementing the computational tools and statistical functions needed in the project as an API. C++ was used as the language of development. Additionally, I also worked on developing visualization modules to represent the derived results.
  - Supervisor: **Colin Robertson**, Associate Professor, Wilfrid Laurier University
  - Focus: Cartography, Geographic Modelling, and Geo-computation

GRANTED AFFILIATIONS

- Faculty Affiliation at the Vector Institute for AI — Jul 2025
- Faculty Affiliation at the Schwartz Reisman Institute for Technology and Society — Jul 2025

COMPETITIVE GRANTS AWARDED

- NSERC Discovery Grant (192,500 CAD).
  Role: Sriram Ganapathi Subramanian was the sole PI for this grant. Jul 2025
- Canada Research Chair (Tier II) in Artificial Intelligence (600,000 CAD).
  Role: Sriram Ganapathi Subramanian was the sole PI for this grant. Jul 2025
- Coefficient Giving - Technical AI Safety Grant (412,500 USD).
  Role: Sriram Ganapathi Subramanian was the co-PI for this grant (Pascal Poupart from the University of Waterloo served as the PI). Dec 2025

AWARDS & SCHOLARSHIPS

- Canadian AI Association (CAIAC) Best Doctoral Dissertation Award (2000 CAD) May 2023
- Waterloo AI Scholarship (5000 CAD) May 2022
- University of Waterloo Research Excellence Award (2000 CAD) Feb 2022
- Vector Institute Postgraduate Affiliation (12000 CAD) Apr 2021
- S.P. Pasupalak Scholarship in Robotics and Artificial Intelligence (1500 CAD) Sep 2020
- AI for Earth Microsoft Grant (15000 USD) May 2018
- Faculty of Engineering Award, University of Waterloo. (1500 CAD, multiple times) May 2017
- Hackathon - EngHacks (2000 CAD) Mar 2017
- Hackathon - HackforHealth (2000 CAD) Jan 2017
- Mitacs Graduate Fellowship (15000 CAD) May 2016
- Mapp your way, Environmental Systems Research Institute (ESRI) (75000 INR) Jan 2016
- Governor's Medal, Anna University Aug 2015
- Mitacs Global Link Research Award (6000 CAD) May 2015

JOURNAL PAPERS UNDER SUBMISSION

[1] Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson, Mark Crowley, and Pascal Poupart "Revisiting Neighbourhoods in Mean Field Reinforcement Learning", *Under Submission to TMLR*, Mar 2026.

[2] Sriram Ganapathi Subramanian, Hao Wang, Yue Yu, Marc St. Aubin, Sharon O'Sullivan, Lawrence Wan, Luis Ricardez-Sandoval, and Pascal Poupart, "Distributional Reinforcement Learning using Feedback from External Knowledge Sources", *Under Submision to Artificial Intelligence Journal (AIJ)*, Mar 2026.

JOURNAL PUBLICATIONS

[1] Hao Wang, Sriram Ganapathi Subramanian, Gustavo Sutter, Pascal Poupart, Luis Ricardez-Sandoval, "A density functional theory-guided discovery of bimetallic transition-metal-doped cerium oxide catalysts for the reverse water gas shift reaction", *Applied Surface Science*, Apr 2026.

[2] Guiliang Liu, Sheng Xu, Shicheng Liu, Ashish Gaurav, Sriram Ganapathi Subramanian, and Pascal Poupart, "A Comprehensive Survey on Inverse Constrained Reinforcement Learning: Definitions, Progress and Challenges", *Transactions on Machine Learning Research (TMLR)*, May 2025.

[3] Chris Beeler, Sriram Ganapathi Subramanian, Kyle Sprague, Mark Baula, Nouha Chatti, Amanuel Dawit, Xinkai Li, Nicholas Paquin, Mitchell Shahen, Zihan Yang, Colin Bellinger, Mark Crowley, and Isaac Tamblyn, "ChemGymRL: An Interactive Framework for Reinforcement Learning for Digital Chemistry", *Digital Discovery*, Feb 2024.

[4] Su Zhang, Srijita Das, Sriram Ganapathi Subramanian, and Matthew E. Taylor, "Two-Level Actor-Critic Using Multiple Teachers", *Transactions on Machine Learning Research (TMLR)*, Oct 2023.

[5] Ken Ming Lee, Sriram Ganapathi Subramanian, and Mark Crowley, "Investigation of Independent Reinforcement Learning Algorithms in Multi-Agent Environments", *Frontiers in Artificial Intelligence*, Sep 2022.

[6] Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson, and Mark Crowley, "Multi-Agent Advisor Q-Learning", *Journal of Artificial Intelligence Research (JAIR)*, 74 (2022), 1-74, May 2022.

[7] Piyush Jain, Sean Coogan, Sriram Ganapathi Subramanian, Mark Crowley, Stephen Taylor, and Mike Flannigan, "A review of machine learning applications in wildfire science and management", *Environmental Reviews*, Canadian Science Publishing, Volume 28, Number 4, Dec 2020.

[8] Sriram Ganapathi Subramanian and Mark Crowley, "Using Spatial Reinforcement Learning to Build Forest Wildfire Dynamics Models from Satellite Images" *Frontiers in ICT*, Apr 2018.

HIGHLY REFEREED CONFERENCE PUBLICATIONS

[1] Sriram Ganapathi Subramanian, Toryn Klassen, and Sheila McIlraith, "Multi-Agent Reinforcement Learning with Reward Machines for Mixed Cooperative-Competitive Environments", *To Appear in Reinforcement Learning Conference (RLC 2026)*, Montreal, Canada (Acceptance Rate: 30%), May 2026.

[2] Chandler Smith et al. "Evaluating Generalization Capabilities of LLM-Based Agents in Mixed-Motive Scenarios Using Concordia", *Neural Information Processing Systems (NeurIPS) - Datasets and Benchmarks Track 2025*, San Diego, USA (Acceptance Rate: 25%), Dec 2025. Note: This paper has 78 authors including Sriram Ganapathi Subramanian.

[3] Shuhui Zhu, Baoxiang Wang, Sriram Ganapathi Subramanian, and Pascal Poupart, "Learning to Negotiate via Voluntary Commitment", *Artificial Intelligence and Statistics Conference (AISTATS 2025)*, Mai Khao, Thailand (Acceptance Rate: 30%), May 2025.

[4] Sriram Ganapathi Subramanian, Galen Liu, Mohammed Elmahgiubi, Kasra Rezaee, and Pascal Poupart, "Confidence Aware Inverse Constrained Reinforcement Learning", *International Conference on Machine Learning (ICML 2024)*, Vienna, Austria (Acceptance Rate: 27%), Jul 2024.

[5] Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson, and Mark Crowley, "Learning from Multiple Independent Advisors in Multi-agent Reinforcement Learning", *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2023)*, London, UK (Acceptance Rate: 23%), May 2023.

[6] Sriram Ganapathi Subramanian, Matthew E. Taylor, Mark Crowley, and Pascal Poupart, "Decentralized Mean Field Games", *AAAI Conference on Artificial Intelligence (AAAI 2022))*, Vancouver, BC, Canada (Acceptance Rate: 15%), Feb 2022.

[7] Sriram Ganapathi Subramanian, Matthew E. Taylor, Mark Crowley, and Pascal Poupart, "Partially Observable Mean Field Reinforcement Learning", *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2021)*, London, UK (Acceptance Rate: 24%), May 2021.

[8] Sriram Ganapathi Subramanian, Pascal Poupart, Matthew E. Taylor, and Nidhi Hegde, "Multi Type Mean Field Reinforcement Learning", *International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2020)*, Aukland, New Zealand, (Acceptance Rate: 23%), May 2020.

OTHER CONFERENCE PUBLICATIONS

[1] Agustinus Kristiadi, Felix Strieth-Kalthoff, Sriram Ganapathi Subramanian, Vincent Fortuin, Pascal Poupart, and Geoff Pleiss, "How Useful is Intermittent, Asynchronous Expert Feedback for Bayesian Optimization?", *Sixth Symposium on Advances in Approximate Bayesian Inference*, Vienna, Austria, Jul 2024.

[2] Sriram Ganapathi Subramanian, Matthew E. Taylor, Kate Larson, and Mark Crowley, "Multi-Agent Advisor Q-Learning (Extended Abstract)", *International Joint Conference on Artificial Intelligence (IJCAI 2023), Journal Track*, Macao, SAR China, Aug 2023.

[3] Su Zhang, Srijita Das, Sriram Ganapathi Subramanian, and Matthew E. Taylor, "Two-Level Actor-Critic Using Multiple Teachers", *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2023)*, London, UK *(Extended Abstract)*, Aug 2023.

[4] Sushrut Bhalla, Sriram Ganapathi Subramanian, and Mark Crowley, "Deep Multi Agent Reinforcement Learning for Autonomous Driving", *Canadian Conference on Artificial Intelligence*, Ottawa, Canada, May 2020. Best paper award nominee

[5] Sushrut Bhalla, Sriram Ganapathi Subramanian, and Mark Crowley, "Training Cooperative Agents for Multi-Agent Reinforcement Learning", *International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2019)*, Montreal, Canada *(Extended Abstract)*, May 2019.

[6] Sriram Ganapathi Subramanian, Jaspreet Singh Sambee, Benyamin Ghojogh, and Mark Crowley, "Decision Assist for Self-driving Cars", *Canadian Conference on Artificial Intelligence*, Toronto, ON, Canada, May 2018.

[7] Sriram Ganapathi Subramanian and Mark Crowley, "Combining MCTS and A3C for Prediction of Spatially Spreading Processes in Forest Wildfire Setting", *Canadian Conference on Artificial Intelligence*, Toronto, ON, Canada, May 2018.

[8] Sriram Ganapathi Subramanian and Mark Crowley, "Learning Forest Wildfire Dynamics from Satellite Images using Reinforcement Learning", *Conference on Reinforcement Learning and Decision Making (RLDM 2018)*, Ann Arbor, MI, USA, May 2018.

WORKSHOP PAPERS

[1] Chris Beeler, Sriram Ganapathi Subramanian, Kyle Sprague, Mark Crowley, Colin Bellinger, and Isaac Tamblyn, "Demonstrating ChemGymRL: An Interactive Framework for Reinforcement Learning for Digital Chemistry", *AI for Accelerated Materials Design - Neural Information Processing Systems (NeurIPS 2023)*, New Orleans, USA, Dec 2023.

[2] Chris Beeler, Sriram Ganapathi Subramanian, Kyle Sprague, Colin Bellinger, Mark Crowley, and Isaac Tamblyn, "ChemGymRL: An Interactive Framework for Reinforcement Learning for Digital Chemistry", *AI for Science - Neural Information Processing Systems (NeurIPS 2023)*, New Orleans, USA, Dec 2023.

[3] Volodymyr Tkachuk, Sriram Ganapathi Subramanian, and Matthew E. Taylor, "The Effect of Q-function Reuse on the Total Regret of Tabular, Model-Free, Reinforcement Learning", *Adaptive Learning Agents Workshop - International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS 2021) (long talk)*, London, UK, May 2021.

[4] Sai Krishna Gottipati, Yashaswi Pathak, Rohan Nuttall, Sahir, Raviteja Chanduru, Ahmed Touati, Sriram Ganapathi Subramanian, Matthew E. Taylor, and Sarath Chandar, "Maximum Reward Formulation in Reinforcement Learning", *Neural Information Processing Systems (NeurIPS)* workshop on Deep RL (Oral), Vancouver, Canada, Nov 2020.

[5] Sriram Ganapathi Subramanian and Mark Crowley, "A Complementary Approach to Improve Wild Fire Prediction Systems" *Neural Information Processing Systems (NeurIPS)* workshop on AI for Social Good, Montreal, Canada, Dec 2018.

PATENTS

[1] Sriram Ganapathi Subramanian, Pascal Poupart, Matthew E. Taylor and Nidhi Hegde; Patent Number: US20200279136A1, "System and method for multi-type mean field reinforcement machine learning", US Patent, Owner: Royal Bank of Canada, Filing Date: Mar 2019.

[2] Zhiyuan Du, Joseph Lull, Rajesh Malhan, Sriram Ganapathi Subramanian, Sushrut Bhalla, Jaspreet Sambee and Mark Crowley; Patent Number: US11131992B2, "Multi-Level Collaborative Control System With Dual Neural Network Planning For Autonomous Vehicle Control In A Noisy Environment", US Patent, Owner: DENSO International America, Filing Date: Nov 2018.

[3] Sriram Ganapathi Subramanian, Keshav Pameshwaran, Sivakumar Senthilkumar, Parthipan Muthu; Patent Number: 11581/2015-CO/L, "Mine closure using remote sensing aid", Indian Software Copyright, Filing Date: Feb 2016.

TEACHING EXPERIENCE

Generative AI and Large Language Models, Instructor, Carleton University	Jan 2026 – Apr 2026
Reinforcement Learning Bootcamp, Instructor, Vector Institute	Mar 2024 – Jun 2024
Introduction to Artificial Intelligence, Instructor, University of Waterloo	May 2023 – Aug 2023
Reinforcement Learning, Teaching Assistant, University of Waterloo	Aug 2021 – Dec 2021
Reinforcement Learning, Teaching Assistant, Vector Institute	Mar 2021 – Jun 2021
Reinforcement Learning, Teaching Assistant, University Of Waterloo	May 2020 – Aug 2020
Fundamentals of Programming, Teaching Assistant, University Of Waterloo	Sep 2019 – Dec 2019
Data and Knowledge Modelling, Teaching Assistant, University Of Waterloo	Jan 2018 – Apr 2018
Algorithms and Data Structures, Teaching Assistant, University Of Waterloo	Jan 2017 – Apr 2017

PROGRAM COMMITTEE MEMBER/ REVIEWER

- AAAI conference on Artificial Intelligence (AAAI) - PC Member - 3 times, Area Chair - 1 time
- International Conference on Machine Learning (ICML) - PC Member - 3 times
- Conference on Neural Information Processing Systems (NeurIPS) - PC Member - 4 times
- Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) - PC Member - 4 times
- International Joint Conference on Artificial Intelligence (IJCAI) - PC Member - 3 times
- Journal of Artificial Intelligence Research (JAIR) - Reviewer - 3 times
- Journal of Autonomous Agents and Multi-agent Systems (JAAMAS) - Reviewer - 3 times
- International Journal of Wildland Fire - Reviewer - 3 times
- Neural Networks - Reviewer - 2 times
- Transaction on Machine Learning Research (TMLR) - Reviewer - 8 times
- Various workshops in ICML, NeurIPS, ICLR, AAMAS, AAAI and IJCAI - PC Member - Multiple times

STUDENTS MENTORED

Kunwar Nir, I-Cureus Undergraduate Research Assistant at Carleton University.
Research project: Inverse Constraint Reinforcement Learning — May 2026 – Aug 2026

Kamal Yassin, Honours Project Student at Carleton University.
Research project: Machine Learning-Based Interpretation of Cardiotocography — May 2026 – Aug 2026

Ryan Chung Kam Chung, Honours Project Student at Carleton University.
Research project: Adaptive Communication Budgeting for Multi-Agent LLM Reasoning — May 2026 – Aug 2026

Julie Wechsler, Honours Thesis Student at Carleton University.
Research project: Optimizing Moment Retrieval for Educational Videos — Sep 2025 – Apr 2026

Ahmed Al-Obaidi, Honours Project Student at Carleton University.
Research project: A Training-Time Safety Dynamics for Constrained RL — Jan 2026 – Apr 2026

Adnan El Assadi, Honours Project Student at Carleton University.
Research project: Adaptive Reward Shaping for Text-Based Adventure Games — Sep 2025 – Dec 2025

Shuhui Zhu, Doctoral student at the University of Waterloo.
Research project: Communication in Multi-agent Reinforcement Learning — Sep 2023 – Dec 2024

Kateryna Nekhomiazh, Master's student at the University of Toronto.
Research project: Reward Machines — Sep 2023 – Sep 2024

Cynthia Huang, Doctoral student at the University of Waterloo.
Research project: Multi-agent Reinforcement Learning in Fintech — Nov 2022 – Sep 2023

Yanting Miao, Doctoral student at the University of Waterloo.
Research project: Offline Reinforcement Learning — Nov 2022 – Sep 2023

Kyle Sprague, Undergraduate student at the University of Ottawa.
Research project: Algorithmic development in ChemgymRL — Sep 2022 – Aug 2023

Su Zhang, Doctoral student at Washington State University.
Research project: Two-level Actor-Critic Using Multiple Teachers — Feb 2022 – Jul 2023

Nouha Chatti, Master's student at the University of Waterloo.
Research project: Curriculum learning in Chemgym RL — Oct 2021 – Apr 2022

Vlad Tkachuk, Undergraduate student at the University of Waterloo.
Research project: Effect of Q-function Reuse on Tabular, Model-Free Reinforcement Learning — Oct 2021 – Apr 2022

Ken Ming Lee, Undergraduate student at the University of Waterloo.
Research project: Investigating independent reinforcement learning algorithms in multi-agent environments — Aug 2020 – Aug 2021

Zihan Yang, Undergraduate student at the University of Waterloo.
Project: Software development of Chemgym RL — Sep 2019 – Apr 2020

Mark Baula, Undergraduate student at the University of Waterloo.
Project: Software development of Chemgym RL — May 2019 – Apr 2020

Amanuel Dawit, Undergraduate student at the University of Waterloo.
Project: Software development of Chemgym RL — May 2019 – Apr 2020

Xinkai Li, Master's student at the University of Waterloo.
Project: Software development of Chemgym RL — May 2019 – Apr 2020

INVITED TALKS (SELECTED)

- From Autonomy to Accountability: Engineering Constrained Cooperation in Multi-Agent Systems - Data Day, Carleton University - (March 2026)

- Artificial Intelligence in the Physical World - Science Cafe, Ottawa Public Library - (March 2026)
- Real-world multi-agent reinforcement learning - ServiceNow, Montreal (Oct 2025)
- The AI Policy Lab: Simulating Our Way to Smarter Government - Innovation Morning, Carleton University Kanata - (Sept 2025)
- Accelerating training in multi-agent reinforcement learning through action advising - U.C. Berkeley: Multi-Agent Learning Seminar (Aug 2023)
- Multi-agent Reinforcement Learning in Large Complex Environments - Canadian AI, Montreal, 2023 (Doctoral Dissertation Award talk - June 2023)
- Decentralized Mean Field Games – University of Alberta – Alberta Machine Intelligence Institute (Amii) AI seminar (Aug 2022)
- Multi-agent Reinforcement Learning in Large, Complex Environments – Vector Institute (Aug 2022)
- Multi-agent Advisor Q-learning – Vector Institute (Feb 2022)
- Recent Advances in Mean-Field RL Methods – University of Alberta – Alberta Machine Intelligence Institute (Amii) AI seminar (May 2021).
- Multi Type Mean Field Reinforcement Learning – A.I. Socratic Circles (AISC), Toronto (Sep 2020)
- Reinforcement Learning in SSP Domains – Borealis AI Toronto (Apr 2018)
- Algorithmic Analysis and Improvements in Multi-Agent Reinforcement Learning for Partially Observable Setting – University of Waterloo AI Seminar (Mar 2018)

MEDIA

- Waterloo News: Transforming CO2 into valuable products (Dec 2023)
- Waterloo News: Recent electrical and computer engineering grad, Sriram Ganapathi Subramanian, wins Canadian Artificial Intelligence Association's Best Doctoral Dissertation Award. (Jun 2023)
- Alberta News: Scaling Up AI for Success (Feb 2022)
- The Record: Waterloo professor says AI is a useful tool to help fight wildfires. (Dec 2019)
- The Scientist: Artificial Intelligence Tackles a World of Images. (May 2019)
- Waterloo News: Fighting wildfires with artificial intelligence. (Nov 2018)
- Wired: How Supercomputers Can Help Fix Our Wildfire Problem. (Nov 2018)
- Digital: AI for Earth: Microsoft's Most Compelling Environmental Projects. (Sep 2018)
- Waterloo News: Waterloo Innovation Summit addresses climate change at Vancouver event. (Nov 2018)

SERVICE

- Co-op report marker (Fall 2025 & Winter 2026)
- Reviewer for the MindBridge Challenge at Carleton University (Summer 2025 & Winter 2026)
- Reviewer for the Vector Scholarship in AI at the Vector Institute for the 2026 cohort (Winter 2026).
- Chair for the PhD Proposal Defence of Marzi Heidari (Supervisor: Yuhong Guo)
- Internal/External Examiner for the Master's defense of Bahareh Abolhasanzadeh (Supervisor: Masoud Barati)
- Member of Promotion and Tenure Committee for the School of Computer Science in Fall 2025 (Chair: Pat Morin)
- Organizer - Reinforcement Learning Reading Group, Vector Institute, Winter & Spring 2026
- Part of the awards committee for DSAAI students from August 2025 – April 2026. Duties involved ranking award competition applications (towards NSERC and externally funded grant awards) from DSAAI grad students.

EDI INITIATIVES

Mentor - Indigenous and Black Engineering and Technology PhD Project Canada (IBET PhD Project),

- Mentoring PhD students from underrepresented communities in STEM. Oct 2022 – Present

Instructor - Each One Teach One (EOTO) Campaign - India,

- Teaching and creating opportunities for underprivileged children in India. Sep 2012 – Feb 2016