

Sriram Ganapathi Subramanian

sriram.subramanian@vectorinstitute.ai, 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

Vector Institute, Toronto, Ontario, Canada

Sep 2022 – Present

- Postdoctoral Fellow
 - I am conducting fundamental research on several sub-fields within reinforcement learning and deep learning. Particularly, I am working on Inverse Constraint Reinforcement Learning, Multi-agent Reinforcement Learning, Reinforcement Learning in non-Markovian environments, and Reinforcement Learning applications to Chemistry and Drug Discovery.
 - 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.
 - 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

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 Azure 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 PUBLICATIONS

- [1] 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* (To Appear), Feb 2024.
- [2] 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.
- [3] Ken Ming Lee, Sriram Ganapathi Subramanian, and Mark Crowley, "Investigation of Independent Reinforcement Learning Algorithms in Multi-Agent Environments", *Journal of Frontiers in Artificial Intelligence*, Sep 2022.
- [4] 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.
- [5] 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.
- [6] Sriram Ganapathi Subramanian and Mark Crowley, "Using Spatial Reinforcement Learning to Build Forest Wildfire Dynamics Models from Satellite Images" *Journal of Frontiers in ICT - Environmental Informatics*, Apr 2018.

HIGHLY REFEREED CONFERENCE PUBLICATIONS

- [1] 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.
- [2] 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.
- [3] 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.
- [4] 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)*, Auckland, New Zealand, (**Acceptance Rate: 23%**), May 2020.

OTHER CONFERENCE PUBLICATIONS

- [1] 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.
- [2] 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.
- [3] 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**
- [4] 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.

- [5] Sriram Ganapathi Subramanian, Jaspreet Singh Sambee, Benyamin Ghogh, and Mark Crowley, “Decision Assist for Self-driving Cars”, *Canadian Conference on Artificial Intelligence*, Toronto, ON, Canada, May 2018.
- [6] 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.
- [7] 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, Colin Bellinger, Mark Crowley, 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, 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, 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.

SUBMITTED WORKS

- [1] Sriram Ganapathi Subramanian, Galen Liu, Mohammed Elmahgiubi, Kasra Rezaee, and Pascal Poupart, “Confidence Aware Inverse Constrained Reinforcement Learning”, *Under submission to ICML*, Feb 2024.
- [2] Sriram Ganapathi Subramanian, Toryn Klassen, and Sheila McIlraith, “Multi-Agent Reinforcement Learning with Reward Machines for Mixed Cooperative-Competitive Environments”, *Under submission to ICML*, Feb 2024.

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.S, Parthipan Muthu; Patent Number: 11581/2015-CO/L, “Mine closure using remote sensing aid”, Indian Software Copyright, Filing Date: Feb 2016.

TEACHING EXPERIENCE

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
- International Conference on Machine Learning (ICML) - PC Member - 2 times
- Conference on Neural Information Processing Systems (NeurIPS) - PC Member - 3 times
- Conference on Autonomous Agents and Multi-Agent Systems (AAMAS) - PC Member - 2 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 - 5 times
- Various workshops in ICML, NeurIPS, ICLR, AAMAS, AAAI and IJCAI - PC Member - Multiple times

**STUDENTS
MENTORED**

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

Kateryna Nekhomiazh, Master’s student at the University of Toronto.
 ▪ Research project: Reward Machines Sep 2023 – Present

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

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

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

Nicholas Paquin, Undergraduate student at the University of Waterloo.
 ▪ Project: Software development of Chemgym RL May 2019 – Apr 2020

- INVITED TALKS (SELECTED)**
- 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)
 - 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)
- 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
- SERVICE**
- Expectations Workshop, University of Waterloo**
- Student Mentor Apr 2018
- Responsibilities included:
- Instructing future teaching assistants on the rules and responsibilities of the same at the University of Waterloo.
 - Organizing workshop activities for training students to take up teaching assistant roles at the University of Waterloo.
- Tech Forum, College of Engineering Guindy, Anna University**
- Student Director of Projects Aug 2016 – Aug 2017
- Responsibilities included:
- Guiding student projects on campus.
 - Run the UNESCO patronaged technical event of College of Engineering Guindy, called Kurukshetra.
- LANGUAGES**
- Tamil: Native language
 - English: Fluent (speaking, reading, writing)
 - Hindi: Intermediate (reading); basic (speaking, writing)
 - Sanskrit: Basic (speaking, reading, writing)