Montreal, Canada
nishanth.anand@mail.mcgill.ca
itsnva7.com
nishanth-anand-vemgal-185a5757
NishanthVAnand
itsNVA7

Nishanth Anand

Education

2019-Present **Ph.D. in Computer Science**, *McGill University, Montreal, Canada*, CGPA: **4.0/4.0**.

2017-2019 **M.Sc. in Computer Science**, *McGill University, Montreal, Canada*, CGPA: **3.92/4.0**.

2011-2015 **B.E. in Telecommunications Engineering**, *PESIT (now PES University), Bengaluru, India*, CGPA: **9.0/10.0**.

Publications

NeurIPS 2023 **Nishanth Anand**, Doina Precup, *Prediction and Control in Continual Reinforcement Learning.* (Acceptance rate: 26.1%). [PDF]

ICML 2021 **Nishanth Anand**, Doina Precup, *Preferential Temporal Difference Learning.* (Acceptance rate: 21.5%). [PDF]

2019 **Nishanth Anand**, *Temporal Credit Assignment via Traces in RL, M.Sc. thesis.* [PDF]

RLDM 2019 Pierre Thodoroff*, **Nishanth Anand***, Lucas-Page Caccia, Doina Precup, Joelle Pineau, *Recurrent Value Functions.* [PDF] (Also accepted at SPiRL workshop, ICLR 2019).

ADCOM 2016 HS Karthik, **Nishanth Anand**, J. Manikandan, *Stock Market Prediction using Optimum Threshold based Relevance Vector Machines.* (Acceptance rate: <40%).

INDICON 2015 **Nishanth Anand**, J. Manikandan, *Sparse representation using optimum threshold based relevance vector machine.* (Acceptance rate: 35.6%). [PDF]

INDICON 2015 **Nishanth Anand**, J. Manikandan, *SAR image compression using Relevance Vector Machines.* (Acceptance rate: 35.6%). [PDF]

Academic Research Experience

Sept 2019 -Present **Ph.D. Thesis**, *McGill University and Mila*, Montreal, Canada.

- Researching and developing algorithms for continual reinforcement learning. Supervisor: Prof. Doina Precup.
- I introduced Preferential Temporal Difference Learning, a method to estimate value function efficiently in the presence of state preferences and partial observability.
- Inspired by complementary learning systems (CLS) theory, I introduced an algorithm to decompose the value function into two components and learning them at different timescales for continual reinforcement learning.

Jan 2018 - Aug 2019 **Master's Thesis**, *McGill University and Mila*, Montreal, Canada.

- I developed Recurrent Learning for Reinforcement Learning — a new deep reinforcement learning method to compute value function efficiently under partial observability.
- I unified various eligibility traces in reinforcement learning using Infinite Impulse Response (IIR) filters and demonstrated its advantages empirically.

Sept 2013 - May 2015 **Bachelor's degree project**, *CORI lab and PES University*, Bengaluru, India.

- I introduced an optimum threshold-based pruning for relevance vector machines and applied it to compress images.

Industry Experience

July 2015 - June 2017 **Data Scientist**, *Fractal Analytics*, Bengaluru, India.

- My responsibilities included: solving clients' business problems using AI-ML-based approaches, developing proof of concepts and scaling promising solutions, and analyzing and drawing insights from real-world data.

Selected Projects

Feb 2020 - April 2020 **Matrix Splitting Methods for Prediction in RL**, *Course project*, Excursions in Reinforcement Learning.

I developed sample-based learning algorithms of Gauss-Seidel, Successive Over Relaxation, and Richardson's matrix splitting methods for the prediction problem in reinforcement learning. These methods performed better than the vanilla temporal difference learning algorithm on small-scale problems.

Feb 2019 - April 2020 **IIR Filters as a General Framework for Eligibility Traces**, *Course project*, Deep Learning Theory.

I generalized several versions of eligibility traces — a basic mechanism for temporal credit assignment in reinforcement learning — using IIR filters (a particular type of filters commonly used in communication networks). The general framework performed better on basic reinforcement learning problems. This work began as a course project and it eventually became a chapter in my Master's thesis upon further research and development.

Dec 2015 - April 2017 **Optimizing Campaign Management**, *Fractal Analytics*, Bengaluru, India.

We developed a reinforcement learning-based solution to optimize the campaign management of a business firm. Our solution reduced campaign management costs while increasing sales by optimizing for a longer horizon when compared with their existing solution.

Sept 2016 - Jan 2017 **Anomaly detection**, *Fractal Analytics*, Bengaluru, India.

We implemented an end-to-end, real-time anomaly detection model based on regression analysis for a windmill gearbox. Our solution detected anomalies effectively in the gearbox in advance, which provided sufficient time for the operators to intervene.

Academic/Volunteering Activities

May 2024 **Lead organizer of the Mila RL Workshop**.

- I co-led all the efforts in organizing the one-day RL workshop at Mila, whose goal was to review the RL contributions of the Mila community in the last year. The professors also shared their research agenda for the upcoming year.
- I was responsible for coordinating with the professors, securing funds, drafting the program, making announcements at Mila, coordinating with the event team members and vendors to arrange refreshments, leading a team of volunteers for various tasks on the event day, and hosting the speakers.

June-Dec 2023 **Co-organizer of the New in ML workshop, NeurIPS 2023**.

- The New in ML workshop is a part of NeurIPS affinity workshops focussed on promoting the representation of marginalized communities in machine learning. In particular, the New in ML workshop provides a forum for early research students to develop the skills required to thrive in research through a series of talks, a panel discussion, a mentorship session, and a call for papers.
- My primary responsibilities were around the paper review process: drafting instructions for the call for papers and reviewers, and hosting and managing the paper review system on OpenReview. **I moderated the panel discussion on slow science during the workshop.** I also assisted other organizing members in the planning and logistics of the workshop, such as finding reviewers, mentors and speakers.

Nov 2022 - Aug 2023 **EDI Commissioner at McGill CS Graduate Society**, *McGill*, Montreal.

- The purpose of the CSGS is to serve the needs of graduate students registered in the School of Computer Science at McGill University through social events and talks. The role of the EDI commissioner, an elected position in the CSGS, is to promote and organize EDI-related activities in the department.
- During my tenure, I worked on improving the representation of women graduate students by organizing "Women in CS" talks and socials. Women in CS talks, open to both undergraduate and graduate students, featured a reputed speaker from industry or academia who shared their research work along with their experience as a woman in tech. Socials provided an informal and intimate space for women students to make new connections.

June-July 2019/2020 **Teaching Assistant and Mentor at AI4Good Summer Lab**, *Mila*, Montreal.

- AI4Good Summer Lab is a yearly summer school whose goal is to promote the representation of women in machine learning. Each year around 30 participants are selected to undergo training in machine learning through courses and projects.
- I was a teaching assistant and a mentor for two consecutive editions, where I held office hours daily to clarify the teaching material and answered students' questions. I also guided a group of students on their project.

May 2020 - present **Co-organizer of the RL sofa meetings**, *Mila*, Montreal.

- "The RL Sofa" is a one-hour weekly reinforcement learning meeting between Mila and Amii's RL communities to share and discuss their work. These sessions feature presentations by Mila students, and occasionally by renowned experts like Doina Precup and Yoshua Bengio. The primary focus is on showcasing completed projects and recent RL research, with an interactive format that encourages questions and constructive feedback from the audience.
- I have co-organized over 100 talks so far, and my responsibilities include: finding speakers, booking a venue, making announcements, moderating the meeting and Q&A session, and looking after other logistics.

2017-2022 **Teaching assistant at McGill University**.

- Introduction to Reinforcement Learning (Winter 2022, Winter 2020), Applied Machine Learning (Fall 2021), Introduction to Machine Learning (Fall 2018), Computers and Society (Winter 2021), Algorithms and Data Structures (Winter 2018), and Introduction to Operating Systems (Fall 2017)

2022-2024 **Reviewing Experience**.

- ACM Computing Surveys, 2024,
- Journal Proceedings of the Royal Society A, 2024,
- Reinforcement Learning Conference (RLC) 2024 (**Senior Reviewer**),
- International Conference on Machine Learning (ICML) 2024,
- New in ML workshop, NeurIPS 2023 (**Senior Reviewer**),
- Conference on Lifelong Learning Agents (CoLLAs) 2023,
- Conference on Lifelong Learning Agents (CoLLAs) 2022,
- All things Attention workshop, NeurIPS 2022.

Talks

Jan 2024 *Decomposing Value Function along Histories*, Weekly group meetings, Mila.

Dec 2023 *Prediction and Control in Continual RL*, NeurIPS 2023, New Orleans.

Dec 2023 *Prediction and Control in Continual RL*, Rich Sutton's group meeting, Amii.

Nov 2023 *Prediction and Control in Continual RL*, RL sofa meeting, Mila.

Oct 2023 *On Proto Value Functions*, Weekly group meetings, Mila.

Sept 2023 *Discussions on Continual Reinforcement Learning*, Weekly group meetings, Mila.

March 2022 *TD(λ), eligibility traces and DRL*, Deep RL reading group, Mila.

July 2021 *Preferential Temporal Difference Learning*, ICML 2021, Remote.

July 2020 *TD(λ) Convergence Proof with Function Approximation*, RL theory reading group, Mila.

July 2020 *Q-learning and Deep Q-learning*, Lecture, AI for good summer school.

June 2020 *Preferential Temporal Difference Learning*, RL sofa meeting, Mila.

April 2020 *Introduction to Deep Reinforcement Learning*, Guest Lecture, PES University.

Jan 2020 *Dynamic Programming in Reinforcement Learning*, Lecture, COMP-767, McGill University.

Nov 2019 *Additive Approach to Lambda Returns*, RL sofa meeting, Mila.

Oct 2019 *Meta Gradient Reinforcement Learning*, RL theory reading group, Mila.

May 2019 *Advances in Deep Reinforcement Learning*, Lecture, MMA, McGill University.

March 2019 *Recurrent Value Functions*, RL sofa meeting, Mila. (Also presented at the DLRL Summer School 2019, Edmonton; RLDM Conference 2019, Montreal; SpiRL Workshop, ICLR 2019, New Orleans.)

Skills

Python, R, C, Pytorch, Numpy, Scikit-Learn, Pandas, Comet ML, LaTex, Slurm, Linux, Git.

Relevant Coursework

Graduate	Reinforcement Learning, Excursions in Reinforcement Learning, Machine Learning, Probabilistic Graphical Models, Mathematical Foundations for Machine Learning, Theoretical Principles of Deep Learning, Probabilistic Analysis of Algorithms, Matrix Computation.
Undergrad (selected)	Data Structures, Information theory, Signals and Systems, Micro Controllers, Control Systems, Linear Algebra, Calculus, Probability.