

Nishanth Anand

Ph.D. Candidate in Continual Reinforcement Learning

nishanth.anand@mail.mcgill.ca | Montreal, QC, Canada

itsnva7.com | [inLinkedIn](#) | [GitHub](#) | [Google Scholar](#) | [Twitter](#)

EDUCATION

McGill University / Mila	Montreal, QC, Canada
Ph.D. in Computer Science	2019 – Present
<i>Thesis Topic: Continual Reinforcement Learning</i>	<i>CGPA: 4.0/4.0</i>
McGill University / Mila	Montreal, QC, Canada
M.Sc. in Computer Science	2017 – 2019
<i>Thesis: Temporal Credit Assignment via Traces in RL</i>	<i>CGPA: 3.92/4.0</i>
PES University	Bengaluru, India
B.E. in Telecommunications Engineering	2011 – 2015
<i>Project: Developed a pruning method for RVMs, leading to two IEEE publications</i>	<i>CGPA: 9.0/10.0</i>

PUBLICATIONS

Peer-Reviewed Conference Papers

1. Zijing Wu, et al. (incl. **Nishanth Anand**). “The basal ganglia as a distributed reinforcement learning system computing general values.” *Computational and Systems Neuroscience (COSYNE)*, 2026. (Acceptance Rate: 58%)
2. **Nishanth Anand**, Doina Precup. “Prediction and Control in Continual Reinforcement Learning.” *Conference on Neural Information Processing Systems (NeurIPS)*, 2023. (Acceptance Rate: 26.1%). [[Paper](#)]
3. **Nishanth Anand**, Doina Precup. “Preferential Temporal Difference Learning.” *International Conference on Machine Learning (ICML)*, 2021. (Acceptance Rate: 21.5%). [[Paper](#)]
4. Pierre Thodoroff*, **Nishanth Anand***, et al. “Recurrent Value Functions.” *Reinforcement Learning and Decision Making (RLDM)*, 2019. [[Paper](#)]

Workshop Papers

1. Jacob Chmura, et al. (incl. **Nishanth Anand; last author**). “AIF-GEN: Open-Source Platform and Synthetic Dataset Suite for Lifelong Reinforcement Learning on Large Language Models.” *ICML 2025 Workshop on Championing Open-source Development in ML*, 2025.

Manuscripts in Preparation & Theses

1. **Nishanth Anand**, Doina Precup. “Permanent and Transient Representations for Continual Reinforcement Learning.” *In preparation.*
2. Kanishk Jain, et al. (incl. **Nishanth Anand**) “Discovering Failure Modes in Vision Language Models (VLMs) using RL” *In preparation.*
3. Shahradsad Mohammadzadeh, et al. (incl. **Nishanth Anand; last author**). “AIF-GEN: Open-Source Platform and Synthetic Dataset Suite for Lifelong Reinforcement Learning on Large Language Models.” *In preparation.*
4. **Nishanth Anand***, Zijing Wu*, et al. “Habits for Reinforcement Learning.” *In preparation.*
5. Mohammad-sami-nur Islam, et al. (incl. **Nishanth Anand; last author**). “An Empirical Investigation into the PT Framework” *In preparation.*

6. **Nishanth Anand.** “The Permanent and Transient Framework for Continual Reinforcement Learning.” *Ph.D. Thesis, McGill University*, 2026.
7. **Nishanth Anand.** “Temporal Credit Assignment via Traces in RL.” *M.Sc. Thesis, McGill University*, 2019. [\[PDF\]](#)

RESEARCH EXPERIENCE

Ph.D. Researcher, McGill University / Mila Montreal, QC
Supervisor: Prof. Doina Precup Sep 2019 – Present

- Developed novel algorithms for continual reinforcement learning inspired by the complementary learning systems (CLS) theory from neuroscience.
- Developed a framework that decomposes value functions into permanent (foundational) and transient (fine-tuned) components to accelerate adaptation in non-stationary environments.
- Introduced Preferential Temporal Difference Learning, a method for efficient value estimation under state preferences and partial observability.
- Authored first-author papers at top-tier conferences, including **NeurIPS 2023** and **ICML 2021**.

INDUSTRY EXPERIENCE

Data Scientist, Fractal Analytics Bengaluru, India
 Jul 2015 – Jun 2017

- Developed and deployed custom AI/ML solutions to solve key business challenges for CPG clients in campaign management, inventory optimization, and predictive maintenance.
- Engineered an RL-based solution for campaign management, optimizing for long-term sales lift while reducing operational costs.
- Built a real-time anomaly detection system for wind turbine gearboxes using regression analysis, enabling proactive maintenance and preventing costly failures.

TEACHING EXPERIENCE

Co-Instructor, COMP-579: Reinforcement Learning, **McGill University**
 Winter 2026 (Upcoming)

- Co-instructing (50%) a graduate-level reinforcement learning course with Prof. Doina Precup.
- Responsibilities will include delivering lectures, holding office hours, developing course materials, and designing assignments and exams for a class of 180+ students.

Co-Instructor, INF8250AE: Introduction to Reinforcement Learning, **Polytechnique Montreal**
 Fall 2024

- Delivered one-third of the lectures for a graduate-level course, covering foundational topics like dynamic programming, Monte Carlo methods, and TD learning.
- Co-developed course assignments and exams, and held weekly office hours to support a class of 100+ students.

Instructor and Mentor, AI4Good Summer Lab Mila
 Summer 2019, 2020

- Designed and delivered the core Reinforcement Learning module for one edition of the program.

- Mentored teams of undergraduate students on capstone projects applying ML to social good problems, guiding them from ideation to final presentation.

Teaching Assistant, McGill University

2017 – 2022

- Courses: Introduction to RL, Applied ML, Introduction to ML, Algorithms & Data Structures.

RESEARCH MENTORSHIP & SUPERVISION

Lead Supervisor, AIF-GEN Project

2024 – Present

- Supervised a student team in the creation of **AIF-GEN**, a novel platform for generating synthetic preference data to study the continual alignment of LLMs via reinforcement learning.
- This work culminated in a publication at the **ICML 2025 CodeML workshop** (with my mentee as first author and myself as senior author), and a full journal article is in preparation for **JMLR**.

Co-supervisor (with Prof. Doina Precup), MSc Thesis

2024 – Present

- Co-supervising an MSc student’s thesis investigating architectural choices (e.g., Mixture of Experts, Attention) and neural plasticity mechanisms for my continual RL framework.

Research Mentor, Neuro-Inspired RL Projects

2024 – Present

- Providing mentorship and research direction to two junior PhD students on projects exploring brain-inspired RL.
- This collaboration resulted in two co-authored manuscripts on modeling the striatum’s functions and learning variable timescales for adaptation.

Reinforcement Learning Advisor, Inter-Lab Collaboration

2024 – Present

- Serving as the primary RL advisor for a PhD student in another lab on a project to discover novel failure modes in Vision-Language Models (VLMs) through adversarial question generation.

SELECTED TECHNICAL PROJECTS

LLM-based State Features for Reinforcement Learning

(Unpublished Ph.D. Project)

- Investigated the use of frozen, pre-trained LLMs as general-purpose state feature extractors for RL.
- Built a pipeline to feed symbolic observations (as text) into an LLM and trained a linear Q-function on its intermediate layer representations.
- Achieved strong performance on small-scale MiniGrid tasks, identifying key scaling challenges for more complex environments like Craftax.

Permanent & Transient Successor Features (SFs)

(Unpublished Ph.D. Project)

- Explored a novel extension of my core continual RL framework to Successor Features (SFs).
- Designed a method to decompose SFs into permanent and transient components to handle non-stationary rewards and transition dynamics simultaneously.
- Preliminary results demonstrated faster adaptation compared to baseline SF models in changing environments.

IIR Filters as a General Framework for Eligibility Traces

(M.Sc. Project → Thesis Chapter)

- Unified various eligibility trace mechanisms in RL under the mathematical framework of Infinite Impulse Response (IIR) filters.
- This work began as a course project and was developed further to become a core chapter of my M.Sc. thesis.

LEADERSHIP & PROFESSIONAL SERVICE

- **Lead Organizer**, Mila RL Workshop May 2024
- Co-led the organization of a one-day workshop showcasing RL research from the Mila community, managing logistics, funding, and a team of volunteers.
- **Co-organizer**, New in ML Workshop @ NeurIPS 2023 Dec 2023
- Managed the end-to-end paper review process on OpenReview for an affinity workshop promoting marginalized communities in ML. Moderated a panel on “slow science.”
- **EDI Commissioner**, McGill Computer Science Graduate Society 2022 – 2023
- Organized a “Women in CS” speaker series and networking socials to improve the representation and community for women graduate students.
- **Co-organizer**, The RL Sofa Meetings (Mila & Amii) May 2020 – Present
- Co-organized over 150+ weekly talks for the joint Mila/Amii RL communities, featuring students and renowned experts.
- **Peer Review:**
- **Senior Reviewer:** RLC (2024, 2025, 2026), New in ML @ NeurIPS (2023)
- **Reviewer:** NeurIPS (2025), ICML (2024, 2025, 2026), ICLR (2025, 2026), AAAI (2025), CoLLAs (2022, 2023), ACM Computing Surveys (2024), Proceedings of the Royal Society A (2024)

INVITED TALKS & PRESENTATIONS

- **Poster Presentation**, “Prediction and Control in Continual RL,” *NeurIPS 2023*, New Orleans, LA. Dec 2023
- **Invited Talk**, “Prediction and Control in Continual RL,” *Rich Sutton’s Group Meeting*, Amii, Remote. Dec 2023
- **Poster Presentation**, “Preferential Temporal Difference Learning,” *ICML 2021*, Remote. Jul 2021
- **Guest Lecture**, “Introduction to Deep Reinforcement Learning,” *PES University*, Remote. Apr 2020
- **Guest Lecture**, “Dynamic Programming in Reinforcement Learning,” *COMP-767, McGill University*, Montreal, QC. Jan 2020
- **Poster Presentation**, “Recurrent Value Functions,” *RLDM Conference 2019*, Montreal, QC. Mar 2019

AWARDS & HONORS

- PaliGemma Academic Program GCP Credit Award (\$5,000) 2024
- McGill University Graduate Excellence Award (\$25,000 / year) 2019 – Present
- MHRD Scholarship for Excellent Academic Performance 2012 – 2015

TECHNICAL SKILLS

- **Programming:** Python (Expert), Bash
- **AI/ML Frameworks:** PyTorch, JAX/Flax, Hugging Face (Transformers, Datasets), Gymnasium (OpenAI Gym), NumPy, Pandas, Scikit-learn
- **RL Environments:** Atari, Craftax, MiniGrid, Mujoco
- **Developer Tools:** Git, Slurm, Weights & Biases, Docker, LaTeX