

Viet Nguyen

viet@cs.toronto.edu | cs.toronto.edu/~viet/

EDUCATION

Vector Institute, University of Toronto

2023 – 2027

Doctor of Philosophy, Computer Science

Toronto, ON

- **Supervisor:** Prof. Rahul Krishnan
- **Interests:** compositional generalization, AI for healthcare, multimodal LLMs, cognition, ultimate frisbee.
- **Expected completion date:** May 2027

Montreal Institute of Learning Algorithms (Mila), McGill University

2021 – 2023

Master of Science, Computer Science

Montreal, QC

- **Supervisor:** Prof. Doina Precup
- **Thesis:** “Noisy Hierarchical Reinforcement Learning”
- **Interests:** reinforcement learning, learning theory, continual learning, consciousness, intelligence.
- **Completion date:** August 2023

McGill University

2018 – 2021

Bachelor of Science, Honors Probability and Statistics, Minors Computer Science

Montreal, QC

RESEARCH

Machine Learning, Theory and Applications - prof. Rahul Krishnan

Fall 2023 – Now

Vector, University of Toronto - Research Assistant, PhD Student

Toronto, ON

- **Compositional generalization, theory and applications in transformer-based models:** Generalization guarantees for compositional machine learning tasks via compound space analysis. Applications in targeted data augmentation for robust predictions. Applications in finetuning for LLMs and foundation models to specific compositional downstream tasks for exponential increases in efficiency.
- **Statistical guarantees for Post-Deployment Deterioration Monitoring (PDDM):** Sample complexity analysis for disagreement-based post-deployment deterioration monitoring and detection due to deployment distribution change. Designed and executed large-scale experiments on synthetic data to estimate the fidelity of the theoretical results with empirical observations for performance degradation.
- **Sequential decision making with expert demonstrations using maximum entropy:** Extensive empirical analysis on the performance of the proposed algorithm in challenging environments such as Deep Sea and Frozen Lake. Designed and executed experiments on large-scale distributed computing nodes using Jax.

Reinforcement Learning Theory - prof. Doina Precup

Spring 2020 – Summer 2023

Mila, McGill University - Research Assistant, Master's Student

Montreal, QC

- **Hierarchical Reinforcement Learning:** Efficiency and sample complexity of learning MDPs when there are repeated structures. Constructed guarantees on learning efficiency based on similarities among substructures, as well as explored novel metrics in quantifying substructure similarity. (2022-2023)
- **Efficient representations in reinforcement learning:** model-based sampling of representations in the continuous state space setting for block-MDPs and low-rank MDPs for faster and more efficient learning. (2021-2022)
- **Deep exploration in reinforcement learning:** high-probability regret bounds for novel posterior sampling-based algorithms, conclusive experimental results (2020-2021)
- Developed and maintained a deep reinforcement learning api written in pytorch for quick and easy algorithm prototyping, testing, and interpretation.

Learning Theory, Neural Networks, applied PDEs - prof. Abbas Khalili

Winter 2020

McGill University, Research Student

Montreal, QC

- Researched the landscape geometry of the neural network optimization objective from the perspective of gradient flow in the space of probability measures endowed with the Kantorovich metric.
- Employed various concentration inequalities, functional analysis, and propagation of chaos type arguments to prove convergence of stochastic gradient descent to a limiting PDE.

PUBLICATIONS

- **Nguyen, V.**, Shui, C., Giri, V., Arya, S., Verma, A., Razak, F., Krishnan, R. *Provable Post-Deployment Deterioration Monitoring*. Under review, ICLR 2025.
- Balazadeh, V., Chidambaram, K., **Nguyen, V.**, Krishnan, R., Syrgkanis, V. *Sequential Decision Making with Expert Demonstrations under Unobserved Heterogeneity*. NeurIPS 2024.
- **Nguyen, V.** *Noisy Hierarchical Reinforcement Learning*. McGill University, School of Computer Science, Thesis, 2023.
- Ishfaq, H., Cui, Q., **Nguyen, V.**, Ayoub, A., Yang, Z., Wang, Z., R. Precup, D., Yang, L.F. *Randomized Exploration for Reinforcement Learning with General Value Function Approximation*. ICML 2021.
- Ishfaq, H., Yang, Z., Lupu, A., **Nguyen, V.**, Liu, M., Islam, R. Precup, D., Wang, Z. *Provably Efficient Policy Optimization via Thompson Sampling*. Deep Reinforcement Learning Workshop at NeurIPS 2020, BayLearn 2020 Workshop.

SELECTED EXPERIENCES

Research and Development Engineer Intern

Summer 2020

FPT Canada Inc.

Montreal, QC

- Researched reinforcement learning-based solutions and optimizations to the Vehicle Routing Problem (VRP), a generalization of the Traveling Salesman Problem (TSP).
- Created sampling environment for randomized VRPs, designed the reinforcement learning optimization objectives.
- Tested common neural network-based Q-learning algorithms as well as policy gradient-based methods, assessed and evaluated their performances against each other.
- Reproduced linear programming-based solvers from the literature and achieved similar results.

Research Project - prof. Prakash Panangaden

Fall 2019

McGill University

Montreal, QC

- Researched continuous neural networks: formulation and equivalence to regular neural networks.
- Studied convergence to Gaussian processes in the infinite width limit, gradient flow in function space w.r.t. the Neural Tangent Kernel (NTK), and universal approximation of operators and functionals by function machines.
- Proved universal approximation of functional continuous neural networks.
- **Delivered presentation at the Seminary on Undergraduate Mathematics in Montreal (SUMM).**

McGill NeurotechX 2019, ML team

Winter 2019

Self-Driving Wheelchair

Montreal, QC

- Researched state-of-the-art models and feature construction pipelines in classification of motorimagery signals.
- Used numpy, scipy, matplotlib, seaborn for the analysis and interpretation of EEG time series signals and spectrograms of frequency distributions over time.
- Re-implemented milestone works from the literature in PyTorch, ran benchmarks on different EEG datasets, both simulated and real.
- **Placed First in the NeurotechX Competition, 2019.**

OTHER WORKS AND PRESENTATIONS

- **Nguyen, V.**, Hu, E. *Value Iteration-based Provably Efficient Exploration*. January, 2021
- Hu, E., **Nguyen, V.** *Provable Efficiency: Finding Regret Bounds in Reinforcement Learning*. December, 2020
- **Nguyen, V.** *On the Analysis of Stochastic Gradient Descent in Neural Networks via Gradient Flows*. May, 2020
- **Nguyen, V.**, *On the Concentration of Measure in Orlicz spaces of exponential type*. April, 2020
- Hu, E., Huang, L., **Nguyen, V.** *Neural Networks: A Continuum of Potential*. December, 2019
- Zelaya, M.C., Leech, M., **Nguyen, V.** *Fader Networks: A Heuristic Approach*. April, 2019

TEACHING

Basic Reinforcement Learning, SUMS, McGill University <i>Instructor</i>	Winter 21
Foundations of Machine Learning, University of Montreal <i>Teaching Assistant</i>	Fall 20, Winter 21, Fall 21
Math. Foundations for Machine Learning <i>Teaching Assistant</i>	Fall 21
Theory of Machine Learning <i>Teaching Assistant</i>	Winter 22
Theory of Machine Learning <i>Teaching Assistant</i>	Winter 23
AI4Good <i>Teaching Assistant</i>	Summer 23
Introduction to Machine Learning, University of Toronto <i>Teaching Assistant</i>	Fall 23, Winter 24

AWARDS

DeepMind Graduate Award	2023
School Of Computer Science Teaching Assistant Award	Winter 22