

Romina Abachi

romina@alumni.utoronto.ca

[Github](#)

[Scholar](#)

[LinkedIn](#)

Graduating MSc. with 3+ years of experience in Reinforcement Learning and Machine Learning in research and production settings, and in multi-disciplinary teams.

Skills ML RL Deep learning frameworks (PyTorch, Jax) Python C++ AI in Video Games Cloud Computing (Google Cloud Platform) Computer Vision Optimization Git Docker Natural Language Processing Information Retrieval RAG LangChain HuggingFace

Publications

-
- **Abachi, Romina**, Mohammad Ghavamzadeh, and Amir-massoud Farahmand. "Policy-Aware Model Learning for Policy Gradient Methods." arXiv preprint arXiv:2003.00030 (2020)
 - Nikishin, Evgenii, **Romina Abachi**, Rishabh Agarwal, and Pierre-Luc Bacon. "Control-Oriented Model-Based Reinforcement Learning with Implicit Differentiation." arXiv preprint arXiv:2106.03273 (2021).
 - **Abachi, Romina***, Claas Voelcker*, Animesh Garg, and Amir-massoud Farahmand. "VIPer: Iterative Value-Aware Model Learning on the Value Improvement Path." Decision Awareness in Reinforcement Learning Workshop at the International Conference on Machine Learning (ICML) 2022. (* equal contribution)
 - **Abachi, Romina**, and Amir-massoud Farahmand. "Optimistic Risk-Aware Model-based Reinforcement Learning." The 15th European Workshop on Reinforcement Learning (EWRL 2022).
 - Voelcker, Claas A., Arash Ahmadian, **Romina Abachi**, Igor Gilitschenski, and Amir-massoud Farahmand. " λ -AC: Learning latent decision-aware models for reinforcement learning in continuous state-spaces." arXiv preprint arXiv:2306.17366 (2023).

Work Experience

Research Engineer

Montreal QC

Ubisoft -- Rainbow Six Siege

Dec 2024 -- Present

- Leading the research and development of novel methods using large models for cheat detection given player behaviour data and tracking data.
- Coordinating with multiple teams for data collection and preprocessing.

R&D Programmer Intern

Toronto ON

Ubisoft La Forge

Mar -- Nov 2024

- Participated in the development and implementation of a project to benchmark Imitation Learning and Reinforcement Learning methods for navigation in video games without a NavMesh.
- Designed maps in Godot for use in training and collecting offline data.
- Worked on the integration of RL algorithms for navigation in Rainbow 6 Siege and an unannounced game.
- Carried out discussions with production groups to design a benchmark according to their requirements.

Data Scientist

Toronto ON

Quickplay

Jan 2023 -- Feb 2024

- Developed and deployed an NLP conversational assistant based on Retrieval Augmented Generation for a client's internal content catalog. Finetuned Large Language Models (OpenAI, HuggingFace, open-source ones), semantic search, and ANN alg with reusable embeddings.
- Developed PoC for time series prediction and anomaly detection of Quality of Experience attributes.
- Experimented with Scene Detection algorithms for movies and TV series, improving the existing approach from 14% test accuracy to 50% on the MovieNet benchmark.
- Determined requirements for A/B testing on the front page layout of the app for a client.

Applied Research Scientist for Industrial Applications of RL

Toronto ON

Vector Institute (jointly with Telus Comm., Linamar Corp., Loblaws Ltd.)

May 2019 -- Jan 2020

- Assisted in teaching a class of 6-10 industry participants the foundations of reinforcement learning, created assignments and conducted reading group sessions
- Working individually with each company's group to guide the design of a simulation for their project and implement it in Python, and to come up with possible solutions
- Projects were: allocation of trucks to gates at a distribution centre, datacenter cooling, and short-distance fleet management. Some results presented [here](#).

Research Scientist Intern

Toronto ON

Borealis AI - Fundamental Team

May -- Oct 2018

- Studied literature on verifiable robustness bounds to adversarial attacks, improving on the state-of-the-art robustness bounds based on convex relaxations by ~8% by replacing the activation function in the dual network by a dual version of leaky ReLU.

Education

M.Sc. in Computer Science

Toronto ON

University of Toronto, Artificial Intelligence

Feb 2020 -- Nov 2024

- Supervisors: Amir-masssoud Farahmand and Sheila McIlraith
- Relevant courses: Minimizing Expectations, Knowledge Representation and Reasoning, Algorithms in Private Data Analysis, Algorithms in Collective Decision Making
- Teaching Assistant for Introduction to RL and Introduction to ML

M.A.Sc. in Electrical and Computer Engineering

Toronto ON

University of Toronto, Probabilistic and Statistical Systems Lab

Sept 2017 -- Jan 2020

- Supervisors: Amir-masssoud Farahmand and Brendan Frey
- Thesis: "Policy-Aware Model-based Reinforcement Learning"
- Relevant courses: Statistical Learning Theory, Convex Optimization, Probabilistic Learning and Reasoning

B.A.Sc. in Electrical Engineering, High Honours, with PEY

Toronto ON

University of Toronto

Sept 2012 -- Jun 2017

- Capstone project: "Quadcopter-based Solar Panel Cleaning", supervised by: Professor Olivier Trescases
- Relevant courses: Machine Learning, Digital Communications, Algorithms & Data Structures, Statistics
- Cumulative GPA: 3.92/4.00

Highlighted Research Projects

Stochastic Dueling Contextual Bandits for Discovery of Collective Art Preferences

Toronto ON

University of Toronto

Jan 2022 -- March 2022

- On a quest to find the "Most preferred painting" of a group of people, we created a tool using a StyleGAN model and stochastic dueling contextual bandits to gradually find the collective preferences of different individuals by moving in the latent space of a generative model (inspired by [this](#) project).

Robust Planning for Model-based Reinforcement Learning

Toronto ON

University of Toronto

Jan 2021 -- April 2021

- Studied robust planning problem given a probabilistic constraint on the dynamics model
- Using the modified differential method of multipliers, formulated the constrained optimization in terms of gradient descent updates. Results showed improvement over using the MLE for planning.

Learning Planning Models from Partially-Observed Data

Toronto ON

University of Toronto

Sept 2020 -- Jan 2021

- Studied the performance of symbolic planning with partially-observed image data. Combined the Latplan framework for learning symbolic models for planning directly from image inputs (Asai and Fukunga, 2018) and the GP-VAE (Fortuin et al, 2020) for missing-input imputation. Thus we leveraged developments in the probabilistic imputation of multi-variate time series to improve planning performance with missing data.