

Evgenii Nikishin

Montreal, QC, Canada
nikishin.evg@gmail.com
evgenii-nikishin.github.io

A final-year PhD student studying agents and reasoning. I currently work on improving the capability of language models to self-correct generated code based on execution feedback.

Education

- 2021–Now Ph.D. student in Computer Science (GPA: 4.3/4.3).
Mila, University of Montreal
Advisors: [Pierre-Luc Bacon](#), [Aaron Courville](#)
- 2019–2020 Ph.D. student in Operations Research.
Cornell University, transferred to Mila
- 2017–2019 M.Sc. in Computer Science, with honors (GPA: 8.9/10.0).
Higher School of Economics (HSE) and
Skolkovo Institute of Science and Technology
Advisor: [Dmitry Vetrov](#)
- 2013–2017 B.Sc. in Computer Science, with honors (GPA: 4.9/5.0).
Lomonosov Moscow State University
Advisor: [Alexander D'yakonov](#)

Experience

- Jul 2022– Research Scientist Intern, **DeepMind**.
- Dec 2022 Understanding loss of plasticity in deep RL. Hosts: [Junhyuk Oh](#), [André Barreto](#)
Resulted in one oral paper at ICML 2023 and one spotlight paper at NeurIPS 2023.
- May 2020– Research Intern (remote), **Mila, University of Montreal**.
- Dec 2020 Addressing the objective misalignment in model-based RL. Hosts: [Pierre-Luc Bacon](#), [Yoshua Bengio](#)
Resulted in a publication at AAAI 2022.
- Sep 2017– Research Assistant, **Bayesian Methods Research Group** and **Samsung AI Lab**.
- Aug 2019 Improving stability and transfer in deep RL. Supervisor: [Dmitry Vetrov](#)
Resulted in two workshop papers at UAI 2018 and NeurIPS 2019.
- Jul 2018– Research Fellow, **ETH Zürich**.
- Aug 2018 Applying offline RL for personalized medical treatment. Host: [Gunnar Rätsch](#)
Selected as one of 16 participants of the [Summer Research Fellowship](#) program from among 1,000+ applicants.
- Jul 2016– Research & Development Intern, **Kaspersky Lab**.
- Dec 2016 Implemented an autoencoder-based feature extractor for textual representations of executable files.
The features were later used in production for detecting malicious software.

Selected Awards

- 2023 University of Montreal End of PhD scholarship.
- 2023 Co-author and co-recipient of Google Institute Research Program Grant.
- 2023 University of Montreal AI scholarship.
- 2022 ICML outstanding reviewer.
- 2021, 2022 DIRO excellence scholarship.
- 2021 ICLR outstanding reviewer.
- 2020 Mitacs-Mila short-term research grant.
- 2019 Cornell graduate fellowship.
- 2019 NeurIPS 2019 travel award.
- 2019 MLSS 2019 full travel grant.
- 2016 State academic excellence scholarship.
- 2016 Second place in Yandex Chatbot Hackathon.
- 2015 Bronze medal in [Driver Telematics Analysis](#) competition on Kaggle.

Media Coverage

- 2024 [How Selective Forgetting Can Help AI Learn Better](#), Quanta Magazine.

Invited Talks

- 2023 The Promise of Foundation Models for Reinforcement Learning.
University of Alberta
- 2023 Deep Reinforcement Learning with Plasticity Injection.
ML Collective
Higher School of Economics
Bellairs Research Institute
- 2023 A Summary of Perspectives on Plasticity Loss.
University of Alberta
- 2022 The Primacy Bias in Deep Reinforcement Learning.
Higher School of Economics
DeepMind
Vector Institute, University of Toronto
University of Oxford and Waymo
Microsoft Research
ML Collective
- 2021 Control-Oriented Model-Based Reinforcement Learning with Implicit Differentiation.
Contextual Optimization Workshop, University of Montreal
Cornell Tech
AI Research Institute Moscow
Amsterdam ML Lab
Stanford University
ML Collective
Google Brain
University of Bristol

Publications

- 2024 [Maxwell's Demon at Work: Efficient Pruning by Leveraging Saturation of Neurons.](#)
Simon Dufort-Labbé, Pierluca D'Oro, **Evgenii Nikishin**, Razvan Pascanu, Pierre-Luc Bacon, Aristide Baratin
We show that neuron saturation, which traditionally was viewed as undesirable, could instead lead to sparse yet accurate networks.
ArXiv.
- 2024 [The Curse of Diversity in Ensemble-Based Exploration.](#)
Zhixuan Lin, Pierluca D'Oro, **Evgenii Nikishin**, Aaron Courville
We demonstrate that individual ensemble members in RL exhibit surprisingly low performance and propose a mechanism to address the limitation.
ICLR 2024.
- 2023 [Deep Reinforcement Learning with Plasticity Injection.](#)
Evgenii Nikishin, Junhyuk Oh, Georg Ostrovski, Clare Lyle, Razvan Pascanu, Will Dabney, André Barreto
We propose an intervention for diagnosing the loss of plasticity phenomenon in RL and dynamically growing neural networks in RL for increasing computational efficiency.
NeurIPS 2023 (**Spotlight**); also Reincarnating RL workshop, ICLR 2023 (**Spotlight**).
- 2023 [Understanding Plasticity in Neural Networks.](#)
Clare Lyle, Zeyu Zheng, **Evgenii Nikishin**, Bernardo Avila Pires, Razvan Pascanu, Will Dabney
An analysis of plasticity loss showing its relation to pathologies in loss landscapes and demonstrating efficiency of layer normalization for mitigating it.
ICML 2023 (**Oral**).
- 2023 [Sample-Efficient Reinforcement Learning by Breaking the Replay Ratio Barrier.](#)
Pierluca D'Oro*, Max Schwarzer*, **Evgenii Nikishin**, Pierre-Luc Bacon, Marc Bellemare, Aaron Courville
We achieve model-free state-of-the-art performance by applying resets for scaling of the number of updates per step
ICLR 2023 (**Oral**); also Deep RL workshop, NeurIPS 2022. [Code](#)
- 2022 [The Primacy Bias in Deep Reinforcement Learning.](#)
Evgenii Nikishin*, Max Schwarzer*, Pierluca D'Oro*, Pierre-Luc Bacon, Aaron Courville
We identify a damaging tendency of deep RL agents to overfit early experiences and propose a simple yet powerful remedy based on periodic resetting of a part of the agent
ICML 2022; also RLDM 2022. [Code](#)
- 2021 [Control-Oriented Model-Based Reinforcement Learning with Implicit Differentiation.](#)
Evgenii Nikishin, Romina Abachi, Rishabh Agarwal, Pierre-Luc Bacon
A model learning method for RL that directly optimizes the sum of rewards instead of likelihood, a proxy to the agent's objective
AAAI 2022; also Unsupervised RL workshop, ICML 2021. [Code](#)

- 2020 [Quantifying and Understanding Adversarial Examples in Discrete Input Spaces](#).
Volodymyr Kuleshov, **Evgenii Nikishin**, Shantanu Thakoor, Tingfung Lau, Stefano Ermon
A framework for attacks based on synonyms and theoretical analysis of adversarial examples in genomics, language, and graphs domains
ArXiv.
- 2019 [Unsupervised Domain Adaptation with Shared Latent Dynamics for Reinforcement Learning](#).
Evgenii Nikishin, Arsenii Ashukha, Dmitry Vetrov
Domain adaptation via learning shared dynamics in a latent space and adversarial matching of the latent representations of states
Bayesian Deep Learning workshop, NeurIPS 2019. [Code](#)
- 2018 [Improving Stability in Deep Reinforcement Learning with Weight Averaging](#).
Evgenii Nikishin, Pavel Izmailov, Ben Athiwaratkun, Dmitrii Podoprikin, Timur Garipov, Pavel Shvechikov, Dmitry Vetrov, Andrew Gordon Wilson
Averaging weights during training to reduce the variance of policy gradient estimates
Uncertainty in Deep Learning workshop, UAI 2018

Service

Organizing: [Decision Awareness in RL](#) workshop at ICML 2022 (lead organizer), [DeepBayes](#) summer school 2018, 2019.

Reviewer: NeurIPS 2023, ICML 2023, ICLR 2023 Reincarnating RL workshop, ICML 2022, ICLR 2022, NeurIPS 2021 workshop proposals, NeurIPS 2021, ICML 2021, ICLR 2021, AAAI 2021, ICML 2020 BIG workshop, IndabaX-AI4D innovation grants 2020, ICLR 2020.

Volunteer: Mila WiML Mentor 2023, ICLR 2020, NeurIPS 2019 WiML workshop.

Admissions: Cornell CS PhD applicants starting in Fall 2020.

Technical Knowledge

Languages	Python, C/C++, SQL
ML Libraries	PyTorch, JAX, vLLM, HuggingFace, SciPy stack, scikit-learn
Technologies	Git, Bash, Docker, \LaTeX

Teaching

- Sep 2020– Dec 2020 Teaching Assistant, **Cornell University**.
ORIE 4350, Introduction to Game Theory
- Sep 2017– Dec 2017 Practice Lecturer, **Higher School of Economics**.
Machine Learning classes for 3rd year CS students
- Jul 2017 Teaching Assistant, **Coursera**.
Quizzes and assignments for the Bayesian Methods for Machine Learning course.
The specialization of the course received Coursera [Outstanding Educator Award](#)

Additional Education

- Aug 2020 (Virtual) [DL and RL Summer School](#), Mila, Canada.
- Jul 2019 [Machine Learning Summer School](#), Gatsby Unit, UK.
- Jul 2019 [Reinforcement Learning Summer School](#), INRIA, France.
- Mar 2017 [Spring school "Structural Inference"](#), Humboldt University, Germany.
- Feb 2017 [Winter school on Reinforcement Learning](#), MIPT, Russia.