Ivan Anokhin

Montreal
✉ i.anokhin.mm@gmail.com
avecplezir.github.io/ivananokhin

I am PhD student currently interested in scalable reinforcement learning systems; asynchronous neural network computation; RL for LLM agents; mid- and long-term memory mechanisms.

Education

08.2023 – 09.2027 PhD in Computer Science, Mila, University of Montreal, Advisor: Irina Rish, Topic: Scalable Reinforcement Learning Systems.

2017 – 2019 Master in Data Science, Skolkovo Institute of Science and Technology, Advisor: Dmitry Yarotsky, Thesis: Loss Surface of a Deep Neural Network.

2012 – 2016 BSc in Applied Mathematics and Computer Science, St. Petersburg State University.

Selected Publications

ICLR 2026 **Learning From the Past with Cascading Eligibility Traces**.
Tokiniaina Raharison Ralambomihanta*, **Ivan Anokhin***, Roman Pogodin*, Samira Ebrahimi Kahou, Jonathan Cornford, Blake Aaron Richards
Paper, Code

ICML2025 workshop **AIF-GEN: Open-Source Platform and Synthetic Dataset Suite for Lifelong Reinforcement Learning on Large Language Models**.
Jacob Chmura*, Shahrad Mohammadzadeh*, **Ivan Anokhin**, Jacob-Junqi Tian, Mandana Samiei, Taz Scott-Talib, Irina Rish, Doina Precup, Reihaneh Rabbany, Nishanth Anand
Paper, Code

ICLR 2025 **Handling Delay in Real-Time Reinforcement Learning**.
Ivan Anokhin, Rishav Rishav, Matthew Riemer, Stephen Chung, Irina Rish, Samira Ebrahimi Kahou
Paper, Code, Blog post

NeurIPS 2023 **Thinker: Learning to Plan and Act**.
Stephen Chung, **Ivan Anokhin**, David Krueger.
Project Page

Experience

08.2023–now PhD, **Mila**.
Worked on asynchronous computation for neural networks to improve speed/latency. Developed parallel layer-wise inference for real-time/on-device RL; accepted to ICLR 2025. Developed a parallel layer-wise backward pass, which resulted in a bio-inspired credit-assignment method; accepted to ICLR 2026. Currently exploring asynchronous reasoning with recurrent Mixture of Experts.
Contributed to continual LLM post-training benchmark; benchmarked DPO/PPO and implemented evaluation metrics; published in ICML 2025 workshop.
Exploring a teacher–student framework to train an LLM using off-policy RL (TOPR); work in progress.

02.2023–09.2023 Research Assistant (remote), **Cambridge University**.
Worked on model-based offline RL, contributed to planning project, published in NeurIPS2023.

06.2021–11.2022 Research Scientist, **Yandex**.
Developed a framework to learn image representations object-wise for unsupervised object detection. Led a team of two students. Worked on an SSL method to learn representations from long videos (e.g. EPIC Kitchen).

03.2020–01.2022 Junior Research Scientist, **Skoltech**.

Investigated properties of the ensemble methods with reduced memory consumption at inference and high accuracy (e.g. BatchEnsemble). Published in AISTATS2022. Investigated loss surface of neural networks. Published in ICML2020.

04.2019–06.2021 Deep Learning Engineer, **Samsung AI Center**.

Developed a generator architecture for GANs without any pixel-interaction after style (noise) conditioning. Published in CVPR2021. Developed generative image-to-image model with style (day-time of an image) and content (landscape) disentanglement. Published in CVPR2020.
Worked on robot navigation in an indoor environment. Added physics into the environment simulator; debugged a robot in a real environment.

12.2016–08.2018 Lead Analyst (ML), **Tinkoff Bank**.

Developed, implemented and tuned Slot-filling and Information Retrieval (IR) transformer-based neural network architectures for the chatbot. Slot-filling and IR models were deployed.

Technical skills

Proficient in: Python, Pytorch, git, Latex
Familiar with: Ray, vLLM, Deepspeed, Jax, Keras, Tensorflow, SQL, Java, Matlab, C++

Awards

2026 University of Montreal AI scholarship, $10,000 CAD

2017-2019 Skoltech Merit scholarship, 56,000 RUB/month

Service and Teaching

Reviewer ICML21-22, ICLR22-23, NEURIPS22-25

Lecturer Invited lecturer on deep learning, GoTo summer camp for high-schoolers, 2020 & 2021 multi-day sessions.