

DENIS TARASOV

Zurich, Switzerland

tarasov.denis.al@gmail.com

+41 078 311 77 01

<https://dt6a.github.io/>

[linkedin.com/in/tarasovdeal/](https://www.linkedin.com/in/tarasovdeal/)

EDUCATION

ETH Zurich <i>Master of Artificial Intelligence, GPA 5.5/6</i>	September 2023 – Expected June 2025
Constructor (ex Jacobs) University Bremen <i>Bachelor of Applied Mathematics and Computer Science.</i>	September 2022 – June 2023
Higher School of Economics <i>Bachelor of Applied Mathematics and Computer Science, GPA 8.73/10</i>	September 2019 – November 2022

EXPERIENCE

Research assistant at EPFL CLAIRE lab Research in RL exploration with VLM models and offline RL. Supervised by Caglar Gulcehre. <i>Skills:</i> Python, PyTorch, JAX, RL.	February 2024 – September 2024
Research assistant at ETH Zurich Language, Reasoning and Education lab Research in reasoning and RLAIIF techniques. Supervised by Mrinmaya Sachan. Published paper at ICML 2024 workshops. <i>Skills:</i> Python, PyTorch, RLHF, RLAIIF, NLP, LLMs.	September 2023 – February 2024
Research intern at InstaDeep Research in application of Offline RL techniques to the drug design. Published paper at NeurIPS 2023 workshop. <i>Skills:</i> Python, JAX, Molecular Biology, RLHF, NLP.	June 2023 – September 2023
Research Scientist at Tinkoff AI Researched Offline RL problems and published papers at ICML and NeurIPS main conferences and multiple papers on workshops. Developed popular research-oriented Offline RL library https://github.com/tinkoff-ai/CORL . <i>Skills:</i> Python, PyTorch, JAX, Flax, NLP, RL.	October 2021 – June 2023
Research Engineering intern at Meta Worked as a part of the AI Applied Research Relevance team. Experimented with applying knowledge distillation to entity linking model in order to improve performance using additional signals from unlabeled data. Collected and processed data with signals, added support of newest model version in the debug tool, improved model performance by 4%. <i>Skills:</i> Python, PyTorch, Presto.	July 2022 – September 2022
Research Engineering intern at Yandex Researched neural networks distributed learning with slow network. Discovered weakness and purposed solution for existing implementations of 1-bit Adam and 1-bit LAMB to prevent divergence. Adapted 1-bit optimizers for production needs. Published results at SEIM 2022 conference. <i>Skills:</i> Python, PyTorch.	July 2021 – October 2021
Research intern at JetBrains Research Research work on predicting of antibodies CDR-H3 loops 3D structures from amino acid sequences with end-to-end Deep Learning approach. Recollected and reprocessed data because original data was incorrect. Modified baseline to make predicted structures physically valid. Applied transfer learning to protein folding network and beaten baseline by 20%. <i>Skills:</i> Python, PyTorch, TensorFlow 1, Molecular Biology.	July 2020 – August 2020

SUPERVISION

Bachelors diploma project supervisor at Constructor University Bremen	November 2023 – May 2024
--	---------------------------------

HONORS AND AWARDS

Best Student Paper Award at AI4D3 NeurIPS 2023 workshop	2023
NeurIPS 2023 Scholar Award Awarded with in-person registration and hotel stay to attend NeurIPS 2023	2023
5th All-Russian Student Olympiad "Ya – professional" Winner in the AI track among undergraduate students, 17th place out of 2000+ participants	2022
"Metric Learning for Facial Descriptors" hackathon by HUAWEI Russian Research Institute at 5th RAAI Summer School 2nd place out of 29 participants	2019

PUBLICATIONS

- [The Role of Deep Learning Regularizations on Actors in Offline RL](#) Preprint 2024
Denis Tarasov, Anja Surina, Caglar Gulcehre
- [Is Value Functions Estimation with Classification Plug-and-play for Offline Reinforcement Learning?](#) TMLR 2024
Denis Tarasov, Kirill Brilliantov, Dmitrii Kharlapenko
- [Distilling LLMs' Decomposition Abilities into Compact Language Models](#) ICML 2024 Workshops AI4MATH and AutoRL
Denis Tarasov, Kumar Shridhar
- [Offline RL for generative design of protein binders](#) NeurIPS 2023 Workshop AI4D3 (Oral)
Denis Tarasov, Ulrich Mbow Sob, Miguel Arbesú, Nima H. Siboni, Sebastien Boyer, Andries Smit, Oliver Bent, Arnu Pretorius, Marcin Skwark
- [Revisiting the Minimalist Approach to Offline Reinforcement Learning](#) NeurIPS 2023
Denis Tarasov, Vladislav Kurenkov, Alexander Nikulin, Sergey Kolesnikov
- [CORL: Research-oriented Deep Offline RL Library.](#) NeurIPS 2023
Denis Tarasov, Alexander Nikulin, Dmitriy Akimov, Vladislav Kurenkov, Sergey Kolesnikov
- [Katakomba: Tools and Benchmarks for Data-Driven NetHack](#) NeurIPS 2023
Vladislav Kurenkov, Alexander Nikulin, Denis Tarasov, Sergey Kolesnikov
- [Anti-Exploration by Random Network Distillation.](#) ICML 2023
Alexander Nikulin, Vladislav Kurenkov, Denis Tarasov, Sergey Kolesnikov
- [Q-Ensemble for Offline RL: Don't Scale the Ensemble, Scale the Batch Size.](#) NeurIPS 2022 Offline RL Workshop
Alexander Nikulin, Vladislav Kurenkov, Denis Tarasov, Dmitriy Akimov, Sergey Kolesnikov
- [Let Offline RL Flow: Training Conservative Agents in the Latent Space of Normalizing Flows.](#) NeurIPS 2022 Offline RL Workshop
Dmitriy Akimov, Vladislav Kurenkov, Alexander Nikulin, Denis Tarasov, Sergey Kolesnikov
- [Prompts and Pre-Trained Language Models for Offline Reinforcement Learning.](#) ICLR 2022 Workshop GPL
Denis Tarasov, Vladislav Kurenkov, Sergey Kolesnikov
- [Fixing 1-bit Adam and 1-bit LAMB algorithms..](#) SEIM 2022 (Oral)
Denis Tarasov, Vasily Ershov
- [Predicting ethnicity with data on personal names in Russia.](#) Journal of Computational Social Science 2023
Alexey Bessudnov, Denis Tarasov, Viacheslav Panasovets, Veronica Kostenko, Ivan Smirnov, Vladimir Uspenskiy

REVIEW EXPERIENCE

- [ICML 2025](#) March 2025
- [Scaling Self-Improving Foundation Models Workshop at ICLR 2025](#) February 2025
- [ICLR 2025](#) October 2024
- [AAAI 2025](#) September 2024
- [NeurIPS 2024](#) June 2024
- [ICML 2024](#) March 2024
- [ICLR 2024](#) October 2023
- [NeurIPS 2023](#) June 2023
- [Workshop on Reincarnating Reinforcement Learning at ICLR 2023](#) February 2023

SIDE PROJECTS

[Multi-agent hide and seek](#) March 2021 – June 2021
Project at JetBrains Research. The goal of project was to use Unity Engine to reproduce results from "[Emergent Tool Use From Multi-Agent Autocurricula](#)" paper by OpenAI where reinforcement learning agents learned to use different tools to maximize their reward while playing competitive game. One of the environments was fully replicated but desired behavior was obtained only in simplified setting because more complicated setups require vast computations. *Skills:* C#, Unity, Unity ML agents.

[IntelliJ IDEA team statistics plugin](#) April 2021 – June 2021
The goal of the project was to create IntelliJ IDEA plugin and web-application which allow to track user statistics in working groups via different (including defined by users) metrics. My part of the project was to implement server and web-application. *Skills:* Python, Django.

[Ethnic classification of names and surnames](#) October 2020 – April 2021
Research project at [Yandex School of Data Analysis](#) on Russian names ethnic classification problem. The goal was to create machine learning algorithm which predicts person's ethnicity by name. My job was to find the best fitting algorithm for the this problem. *Skills:* Python, PyTorch, scikit-learn, pandas, XGBoost, CatBoost, FastText.

[Mithril Project](#) February 2020 – June 2020
Project on creating smart T-shirt which allows to track whether the posture of the wearer is correct. I did the following things in the project: programmed sensors, created Android application to interact with T-shirt by Bluetooth, used machine learning algorithms to track posture correctness. *Skills:* C, C++, STM32, Kotlin, Python, scikit-learn, pandas.