DENIS TARASOV

Zurich, Switzerland

tarasov.denis.al@gmail.com **+41 078 311 77 01** **https://dt6a.github.io/** **linkedin.com/in/tarasovdeal/**

EDUCATION

ETH Zurich **February 2026 – Present**
PhD in Artificial Intelligence and Robotics
Supervised by Prof. Dr. Robert Katzschmann

ETH Zurich **September 2023 – March 2025**
Master of Artificial Intelligence, GPA 5.5/6

Constructor (ex Jacobs) University Bremen **September 2022 – June 2023**
Bachelor of Applied Mathematics and Computer Science

Higher School of Economics **September 2019 – November 2022**
Bachelor of Applied Mathematics and Computer Science, GPA 8.73/10

EXPERIENCE

Researcher at ETH Zurich Soft Robotics Lab **November 2025 – February 2026**
Research in RL for robotics.

Research assistant at EPFL CLAIRe lab **February 2024 – September 2024**
Research in RL exploration with VLM models and offline RL. Supervised by Caglar Gulcehre. *Skills*: Python, PyTorch, JAX, RL.

Research assistant at ETH Zurich Language, Reasoning and Education lab **September 2023 – February 2024**
Research in reasoning and RLAIF techniques. Supervised by Mrinmaya Sachan. Published paper at ICML 2024 workshops. *Skills*: Python, PyTorch, RLHF, RLAIF, NLP, LLMs.

Research intern at InstaDeep **June 2023 – September 2023**
Research in application of Offline RL techniques to the drug design. Published paper at NeurIPS 2023 workshop. *Skills*: Python, JAX, Molecular Biology, RLHF, NLP.

Research Scientist at Tinkoff AI **October 2021 – June 2023**
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`. *Skills*: Python, PyTorch, JAX, Flax, NLP, RL.

Research Engineering intern at Meta **July 2022 – September 2022**
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%. *Skills*: Python, PyTorch, Presto.

Research Engineering intern at Yandex **July 2021 – October 2021**
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. *Skills*: Python, PyTorch.

Research intern at JetBrains Research **July 2020 – August 2020**
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%. *Skills*: Python, PyTorch, TensorFlow 1, Molecular Biology.

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**

5th All-Russian Student Olympiad "Ya – professional" **2022**
Winner in the AI track among undergraduate students, 17th place out of 2000+ participants

"Metric Learning for Facial Descriptors" hackathon by HUAWEI Russian Research Institute at 5th RAAI Summer School **2019**
2nd place out of 29 participants

PUBLICATIONS

Sample-Efficient Real-World Dexterous Policy Fine-Tuning via Action-Chunked Critics and Normalizing Flows **Preprint 2026**
*Chenyu Yang, **Denis Tarasov**, Davide Liconti, Hehui Zheng, Robert K. Katzschmann*

NinA: Normalizing Flows in Action. Training VLA Models with Normalizing Flows **EWM Workshop NeurIPS 2025**
***Denis Tarasov**, Alexander Nikulin, Ilya Zisman, Albina Klepach, Nikita Lyubaykin, Andrei Polubarov, Alexander Derevyagin, Vladislav Kurenkov*

cadrille: Multi-modal CAD Reconstruction with Online Reinforcement Learning **ICLR 2026 (Oral)**
*Maksim Kolodiazhnyi, **Denis Tarasov**, Dmitrii Zhemchuzhnikov, Alexander Nikulin, Ilya Zisman, Anna Vorontsova, Anton Konushin, Vladislav Kurenkov, Danila Rukhovich*

Latent action learning requires supervision in the presence of distractors **ICML 2025**
*Alexander Nikulin, Ilya Zisman, **Denis Tarasov**, Nikita Lyubaykin, Andrei Polubarov, Igor Kiselev, Vladislav Kurenkov*

Vintix: Action Model via In-Context Reinforcement Learning **ICML 2025**
*Andrey Polubarov, Nikita Lyubaykin, Alexander Derevyagin, Ilya Zisman, **Denis Tarasov**, Alexander Nikulin, Vladislav Kurenkov*

Yes, Q-learning helps offline in-context RL **SSI-FM Workshop ICRL 2025**
***Denis Tarasov**, Alexander Nikulin, Ilya Zisman, Albina Klepach, Andrei Polubarov, Nikita Lyubaykin, Alexander Derevyagin, Igor Kiselev, Vladislav Kurenkov*

Object-Centric Latent Action Learning **Workshop on World Models ICLR 2025**
*Albina Klepach, Alexander Nikulin, Ilya Zisman, **Denis Tarasov**, Alexander Derevyagin, Andrei Polubarov, Nikita Lyubaykin, Vladislav Kurenkov*

N-gram induction heads for in-context rl: Improving stability and reducing data needs **SSI-FM Workshop ICRL 2025**
*Ilya Zisman, Alexander Nikulin, Viacheslav Sinii, **Denis Tarasov**, Nikita Lyubaykin, Andrei Polubarov, Igor Kiselev, Vladislav Kurenkov*

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 Mbou 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**