Arina Puchkova

Munich, Germany | arina.pchkva@gmail.com | LinkedIn | GitHub | Google Scholar

Profile

ML researcher and engineer focused on RL, code models and post-training infrastructure. Built Kubernetes reward systems that cut training step time by 15%, deployed models up to 70B, and co-authored NeurIPS 2026 research.

Experience

JetBrains Aug. 2024 – Present

Machine Learning Researcher *Munich, Germany*

- Lead and contribute to research on reasoning efficiency, task interference and on-policy distillation, from hypothesis formation and experimental design to analysis and iteration.
- Design and run distributed reinforcement learning experiments for code generation and other tasks using NeMo-RL, Megatron-Core and VeRL.
- Developed Kubernetes-based reward infrastructure for distributed RL training, managing containers that apply model-generated code changes and run tests. Implemented container queuing, health checks and failure recovery; reduced training step time by 15%.
- Co-authored a NeurIPS 2026 paper on evaluating general coding capabilities.

Research Intern

- Conducted multi-step reinforcement learning experiments using execution feedback.
- Filtered code completion training data and achieved a 12% improvement on a Kotlin completion benchmark after fine-tuning.

Ex-Human, Inc. Aug. 2021 – Oct. 2023

Machine Learning Engineer *San Francisco, CA (remote)*

- Trained generative models up to 34B parameters using parameter-efficient fine-tuning; used TPUs to halve training time for GPT-J and T5.
- Deployed and benchmarked models up to 70B using vLLM, TGI and NVIDIA Triton-LLM; configured one-line deployment and maintained latency below 1.5 seconds per model instance.
- Led development of safeguards for open-domain conversation, building safety classifiers and aligning the dialogue model to reduce unsafe outputs by 52%.
- Prepared training datasets from open-source and proprietary data for instruction tuning.

Leroy Merlin May 2021 – Oct. 2021

Data Analyst Intern *Moscow, Russia*

- Analyzed search quality metrics and A/B tests, conducted customer interviews, and supported language and recommender models; helped develop keyword extraction that improved search accuracy by 11%.

Education

Ludwig Maximilian University of Munich Oct. 2023 – Apr. 2026

MSc Data Science, with honours *Munich, Germany*

- Grade: 1.3 (German grading system). Thesis: Multi-turn Reinforcement Learning for Code Generation (grade: 1.0).

Higher School of Economics Sept. 2018 – June 2022

BA Political Science, with honors; minor in Data Analysis *Moscow, Russia*

- GPA: 9.36/10 (major), 10/10 (minor); graduated in the top three of the class.

Publications

- Co-author, "Don't Claim Benchmark-Oriented Optimization Improves General Coding Capability — Diverse Evaluation Is Required." *NeurIPS 2026.*
- Co-author, "Fine-Tuning GPT-3 for Russian Text Summarization." *CoMeSySo 2021*, Springer.

Technical Skills

Programming: Python, SQL, R
ML and training: PyTorch, Hugging Face, NeMo-RL, VeRL, TRL, Ray, DeepSpeed
Inference and infrastructure: vLLM, TGI, Triton Inference Server, Kubernetes, Docker, GCP
Other tools: Git, Weights & Biases, Flask, Qdrant, Grafana
Languages: English (C1), Russian (native), German (B1)