

Sergei Smirnov

Research Engineer, AI Safety, 5+ years of experience, PhD candidate in CS
mg1sergey@gmail.com ◇ [LinkedIn](#) ◇ [GitHub](#) ◇ Helsinki, Finland, GMT+3

ABOUT

Research Engineer with **5 years** of commercial and research experience, worked in teams from **1 founding engineer** to **15**. Built the **full development cycle**, from product owner and customers' requirements gathering to GPU cluster setup, ML architecture design, and product analytics. Contributed to open source distributed systems for ML data pipelines. **Research Fellow** at MARS and SPAR, ARENA bootcamp graduate. **Mathematics and programming olympiad winner**, Dell Technologies student grant awardee for excellence in math and IT. Mentored students for olympiad preparation and created an **ML course** with 200+ students and 4.8/5 feedback rating.

SKILLS

AI safety: TransformerLens, Inspect; capability evals, eval agents, safety benchmarks, reasoning reliability, calibration and uncertainty, jailbreak and injection defense, safety monitoring, alignment, ethics, RLHF, guardrails, interpretability, robustness.

Core ML: Python, C++, Java, Bash, SQL, Linux; Git, CI/CD, Docker, Kubernetes, Ray, AWS, W&B, MLflow; math, probability, statistics, calculus, linear algebra, optimization, algorithms and data structures; DL, LLMs, NLP, embeddings, RAG, agents and tool use, RecSys; PyTorch, TensorFlow, JAX, CUDA, NumPy, pandas, scikit-learn, HF, tokenizers, OpenAI Gym; fine tuning, adapters, RL, VAE, synthetic data, metrics and losses, model architectures.

EXPERIENCE

Research Fellow, SPAR 09/2025 – Present · 1+ mo

- Facilitate Moral and Knowledge Progress With AI. Mentored by [Tianyi Alex Qiu](#) and [Zhonghao He](#).

Research Fellow, MARS 07/2025 – 10/2025 · 3 mo

- Built an interactive demo on AI risks and misalignment for [CivAI](#). Mentored by [Jesse Gilbert](#).

Doctoral Researcher, University of Helsinki 09/2024 – Present · 1 yr 1+ mo

Research Engineer, SMEY.AI 06/2024 – Present · 1 yr 4+ mo

- Built predictive DNA sequence models leveraging a Neobank of over 1 000 yeast strains, cut R&D timelines from 2 years to 30 days, reduced manual screening effort 10x, and used model interpretability to guide gene edits that boosted oil production by 30%.

Research Engineer, LUT University 09/2023 – 05/2024 · 8 mo

- Developed and validated a noise-perturbation robustness [metric](#) framework for AI models evaluation.

Machine Learning Engineer, Epoch8 08/2022 – 08/2023 · 1 yr

- Parallelized the open source [datapipe](#) engine with Ray (5x throughput) for multi-threaded ML data processing graph execution, built a Rasa chatbot MVP that cut support response time by 60%, and engineered a search and recommendation system that boosted engagement by 14% at [Looky](#).

Machine Learning Engineer, Somin.ai 11/2020 – 07/2022 · 1 yr 9 mo

- Developed and optimized a Facebook ad-targeting model that boosted click-through rate by 15 percent and conversion rate by 10 percent on a \$50,000 monthly ad spend.

EDUCATION

FAEB, Technical AI Safety Bootcamp based on ARENA curriculum 06/2025 – 07/2025

Phd in Computer Science, University of Helsinki, Finland; Supervisor: [Prof. Giulio Jacucci](#) 2024 - Present

MSc in Technologies, LUT University, Finland; with honors, GPA 3.79/4.0, Supervisor: [Prof. Leonid Chechurin](#) 2023 - 2024

BSc in Computer Science, ITMO University, Russia; with honors, GPA 3.90/4.0 2019 - 2023