

Igor Ignashin

Email: ignashin.in@phystech.edu

GitHub: github.com/ThunderstormXX

Profile:

- Research engineer focused on LLM post-training, SFT, inference efficiency, and optimization. I build reproducible training and evaluation pipelines for teacher-generated SFT data, reward parsing, RLVR/GRPO-style math training, benchmark reporting, and vLLM-based early-exit/adaptive decoding.
- Strong mathematical background in optimization and stochastic processes, with research experience connecting theory, experiments, and practical large-model training workflows.

Technical Skills:

- Languages: Python, C++, SQL
- LLM Training/Post-Training: SFT, teacher distillation, RLVR/GRPO-style training, reward functions and graders, synthetic and teacher-generated data, evaluation pipelines
- ML/Optimization: PyTorch, JAX, vLLM, TRL, Hugging Face workflows, RL, LLMs, convex and nonconvex optimization, game theory
- Mathematics: stochastic processes, optimization, minimax problems, probability theory
- Data/Infra: multi-GPU server workflows, JSONL data pipelines, YQL, DataLens, Nirvana, Linux, Git

Work Experience:

- Visiting Research Student, MBZUAI (under Eduard Gorbunov) Jan 2026 – Mar 2026
 - Conducted theoretical research on LMO-based optimization methods and acceleration schemes, including analyses of the Muon optimizer and Lookahead-type techniques.
- Laboratory of Mathematical Optimization Methods, MIPT / ISP RAS / BRAIn Lab 2024–Present
 - Developing LLM post-training and inference-efficiency projects, including a Qwen/DeepScaleR math pipeline with teacher-generated SFT data, reward parsing, RLVR/GRPO-style training, and benchmark reporting, and a vLLM-based early-exit/adaptive-decoding evaluation pipeline.
 - Working on a Huawei project on neural network-based schedulers, with a focus on reducing the performance gap between classical and learned approaches.
 - Proved local convergence of the ALSO method for nonconvex and μ -strongly concave minimax problems in “*Aligning Distributionally Robust Optimization with Practical Deep Learning Needs*”.
 - Leading student research projects at YSDA and MIPT AI360 on LLM inference acceleration, distillation, multi-agent RL, game-theoretic LLM training, SGD dynamics, and applied optimization, with an emphasis on theory, experiments, and publication-quality research.
 - Collaborating with Andrei Leonidov’s group on SGD analysis and multi-agent reinforcement learning, including SGD noise analysis, theoretical frameworks for nonconvex optimization dynamics, and joint research on multi-agent LLM systems in game-theoretic scenarios and social dilemmas.
 - Researched equilibrium traffic assignment problems and Frank-Wolfe algorithm modifications, resulting in two journal publications and two conference presentations. Currently developing stochastic extensions of the LMO oracle and exploring applications to game-theoretic formulations.
- Yandex, Data Analyst Intern June 2022 – September 2022
 - Automated a daily pipeline for ticket distribution and closure analytics for B2B contractors, improving transparency and monitoring efficiency.
 - Built data processing workflows in YQL and designed dashboards in DataLens for operational insights.
 - Worked with Yandex internal tools, including Nirvana, for orchestrating data workflows.

Selected LLM Projects:

- **Qwen math SFT/RLVR pipeline:** set up a post-training workflow from base evaluation to SFT, RLVR/GRPO-style training, final evaluation, and Base vs SFT vs SFT+RL reporting for Qwen models.
- **DeepScaleR teacher-SFT distillation:** built a pipeline for generating teacher solutions, scoring them with local math reward/parsing code, training a student model on SFT JSONL, and evaluating on AIME 2024, MATH 500, AMC 2023, Minerva Math, and OlympiadBench.

- **vLLM early-exit/adaptive decoding:** built a reproducible offline inference and evaluation contour with deterministic sharding, resumable JSONL outputs, merge/progress scripts, metrics, and AdaDecode baseline integration on multi-GPU servers.
- **LLM pruning and healing:** worked on layer pruning, LoRA healing, decoder search, and attention-head merging experiments for LLM compression.

Selected Publications:

My Google Scholar

- Aligning Distributionally Robust Optimization with Practical Deep Learning Needs. Accepted to a NeurIPS 2025 workshop.
- NeurIPS 2026 submission on stochastic dynamics of SGD, proposing a new perspective beyond Brownian-motion assumptions (*Why SGD is not Brownian Motion: A New Perspective on Stochastic Dynamics*).
- The Ridiculous Redundancy of LLMs: Which Layers and Heads Actually Matter. Work on LLM layer pruning, LoRA healing, and attention-head merging for decoder compression.
- Publication in *Computer Research and Modeling*: Frank-Wolfe Modifications for Equilibrium Traffic Assignment.
- Preprint: Stochastic Origin Frank-Wolfe for Traffic Assignment. Presented at TFN-2025. Accepted to the Journal of Mathematical Sciences, Series B.
- Preprint: Modeling Skier Flows via Wardrop Equilibrium in Closed Capacitated Networks. Presented at TFN-2025. Accepted to the Journal of Mathematical Sciences, Series B.
- *Conjugate Frank-Wolfe in Machine Learning*. Presented at OPTIMA and accepted to CCIS.

Education:

- B.S. in Applied Mathematics and Physics, Moscow Institute of Physics and Technology (MIPT), Faculty of Applied Mathematics and Informatics.
- M.S. candidate, MIPT, Faculty of Applied Mathematics and Informatics, Department of Intelligent Data Analysis.

Summer Schools and Hackathons:

- AIRI Institute Summer School on Artificial Intelligence 2025
- Summer School “Control, Information and Optimization” in memory of B. T. Polyak 2025
- Summer School and Hackathon in Financial Mathematics, Quantathon 2024
- 5th and 6th International Summer Schools on Financial Mathematics (2024–2025)

School Olympiad Achievements:

- Winner, Phystech Mathematics Olympiad.
- Winner, Step into the Future Olympiads in Mathematics and Physics.
- Winner, KFU Interregional Mathematics Olympiad.
- Winner, municipal stage of the All-Russian School Olympiad in Mathematics.