Arthur Babkin

linkedin.com/in/arthur-babkin | artur.babkin.2015@gmail.com | github.com/ArthurBabkin | tg: @theother_archeee
Russian — native | English — C1

Education

Innopolis University
B.S., Faculty of Computer Science
- Major: Data Science and Artificial Intelligence. GPA: 4.5 / 5.0.

Publications

EMNLP 2026 (Industry Track) · From Production Traffic to Post-Training: Building a Self-Hosted LLM That Covers the Corporate Request Mix · co-author · ACL Anthology

Experience

Research Engineer (LLM Post-Training & Alignment) July 2025 — Present
T-Bank · Foundation Models R&D team, T-Tech
- Work across the full LLM training lifecycle: training data generation, supervised fine-tuning, reinforcement learning, distributed training, benchmark development and evaluation.
- Built large-scale multi-agent LLM pipelines for generating RL environments and training data: 4,000+ RL environments with verifiable rewards, designed to be diverse, solvable, challenging and unambiguous, and 150,000+ supervised fine-tuning samples, filtered with rejection sampling and other quality checks.
- Improved our models across instruction following (MultiChallenge, IFEval, internal benchmarks) and RAG, multi-turn and agentic banking support (Tau3 banking, internal benchmarks), bringing models with <30B active parameters to the level of state-of-the-art models on selected internal tasks.
- Achieved this through SFT on Megatron and RL, GRPO on verl with multi-component reward modeling designed against reward hacking, e.g. structured-output citation rewards that significantly reduced hallucinations, and RL environments with verifiable rewards (RLVR).
- Built, automated with Apache Airflow DAGs and maintained 15 internal and external benchmarks; worked on multilingual evaluation and helped adapt external benchmarks into Russian.

NLP Engineer November 2024 — July 2025
Croissan Studio · End-to-end AI product development

Telegram bot for company financial analysis
- Led a 3-person development team (NLP, backend) scaling the product from idea to MVP: a Telegram bot on aiogram with a FastAPI backend, integrations with international exchanges and market quotes, and PostgreSQL storage — investors realised 1000%+ ROI, confirmed by a follow-on investment cheque at a higher valuation.
- Built 3 LangGraph agents for real-time analysis of a company's financial health: by metrics and multiples (EPS, EBITDA, P/E, P/S, PEG and others); by news coverage — using Perplexity DeepResearch, RSS feed parsing and news deduplication by a Jaccard similarity threshold over embeddings; and by public filings (e.g. SEC filings).

Legal contract review system
- Parsed and built a vector index over 413K+ fragments of Russian legislation; benchmarked 10+ embedding models on quality, speed and dimensionality — search became 6x faster and the index 2.7x more compact.
- Migrated the index from IVFFlat to HNSW: throughput up 10–15x (to 760 rps), latency 12 ms → 1–2 ms at 98% recall — search was no longer the pipeline bottleneck.

NLP Engineer | Co-founder October 2023 — January 2025
Deeplit · AI math tutor platform
- Pitched the idea and the product to stakeholders in plain, non-technical language — raised $25 000 in cash (investment and grants) plus 2 NVIDIA A100 GPUs for 24/7 use.
- Hired and led a team of 7, including engineering: took the product to an MVP with 2000+ users and to the top-1 education innovation in Russia at the 2024 HSE University and Alfa-Bank award (skillbox.ru).
- Designed a 7-agent system orchestrated with LangChain/LangGraph on a FastAPI backend: each agent had its own toolset via tool calling (RAG, OCR, math calculators and verifiers) for different user support scenarios (solving, checking, guided solving, theory, etc.), with PostgreSQL storage — responses to the user in under 45 seconds.
- Average chat session time grew from 1 to 4.5 minutes (4.5x) and D30 retention from 3% to 14% (4.7x).
- Built an LLM-based math problem solver: prompt engineering and Tree of Thoughts with hypothesis generation and validation during solving, with RAG access — +30% on public 2024 Russian Advanced Maths State Exam papers; SFT+DPO fine-tuning via LoRA/QLoRA added +45%.