

Anton BELYY

anton.belyy@gmail.com · avbelyy.github.io

LinkedIn · GitHub · Scholar

EDUCATION

AUG '19 — MAY '22

M.Sc.Eng. in Computer Science (NLP), Johns Hopkins University, USA
GPA: 4.0. Advisors: Benjamin Van Durme (Primary), Vladimir Braverman (Secondary)
Research projects: knowledge graph completion, controlled text generation, schema induction
Research internship on **Conversational AI** at **Microsoft Research** in Summer '21: [\[paper\]](#)

SEP '13 — MAY '18

B.Sc. in Informatics and Applied Mathematics, ITMO University, Russia
GPA: 3.6. Advisors: Andrey Filchenkov (Primary), Konstantin Vorontsov (Secondary)
Research projects: plagiarism detection, hierarchical topic modeling
Thesis: [Construction and Quality Evaluation of Heterogeneous Hierarchical Topic Models](#)
• **Distinguished thesis award** (given to 2 out of approx. 50 undergraduates)

EMPLOYMENT

JUN '25 — PRESENT

Senior AI Software Engineer in AI Assistant, Bloomberg, USA

- **MCP Tools**: I build and maintain MCP tools powering Bloomberg AI Assistant, ensuring high availability and factuality of AI agents used by 100s of thousands of Terminal users.
- **MCP Resources**: I lead the adoption of MCP Resources as an agent-tool communication layer to enable higher factuality & lower latency / LLM token usage across all Bloomberg AI agents.

JUN '24 — MAY '25

Senior Machine Learning Engineer in Credit Risk Detection, Stripe, USA

- **Proactive LLM agents**: I built LLM agents to investigate merchants' online presence to detect early signs of bankruptcies, projected to save several million USD in credit losses yearly.
- **XGBoost → Transformer**: I led an experimental project migrating Stripe's credit scoring models from XGBoost to Transformer-based DNN, leading a team of five ML engineers.

JUN '22 — MAY '24

Machine Learning Engineer in Knowledge Graphs and LLM Datasets, Apple, USA

- **KG**: I am a founding engineer of the [ODKE](#) project, which transforms unstructured Web data into millions of facts every week to continuously improve coverage and freshness of Apple KG. I implemented scalable fact extractors using in-house and OpenAI LLMs, built fact verification model using Spark MLlib, built CI/CD pipelines and oversaw integrating ODKE with Apple KG.
- **LLM Datasets**: I used MapReduce, Spark and ML to turn petabytes of multimodal Web data into billion-scale, high-quality datasets used to train LLMs for Siri and Apple Intelligence.

SEP '17 — AUG '19

Machine Learning Engineer in Risks and ML Automation, Tochka Bank, Russia

- **Risk scoring**: using Hadoop and Spark, I enriched internal data for 200K+ bank clients and 50M+ transactions with Web-scraped public data to improve risk scoring accuracy by 10%.
- **ML automation**: built ML models to detect intent for the top-90% client questions in chat. Built systems to automate scheduling call center shifts while improving accuracy by 10%.

OCT '15 — SEP '16

Software Engineer in Ads, VK.com, Russia

- **URL fraud**: built service to periodically detect malicious URL redirect changes in VK ads.
- **Click fraud**: built ML models to detect users that generate fraudulent clicks in VK ad network. Model was deployed semi-automatically and helped recover up to 3% monthly ad revenue.
- **Ads search**: launched moderator search interface (incl. full-text search) over 30M+ VK ads.
- **Ads scoring**: implemented advertiser ranking for faster moderation of top-10% clients.

LANGUAGES AND TECHNOLOGIES

LANGUAGES

Python (advanced); Golang, Scala, JavaScript, SQL, C++, bash (proficient)

TECHNOLOGIES

Airflow, Spark, Iceberg (advanced); XGBoost, PyTorch, AWS, Docker, K8s (proficient)