Mehdi *Mehdiyev*

DATA SCIENTIST @ YELO BANK

I build production ML for lending — and make it fast, monitored and reliable enough to run every day.

Baku, Azerbaijan · linkedin.com/in/mekhty-mekhtyev · github.com/Mekhty111 · mekhty111.github.io

NOW	EXPERIENCE	FOCUS	CORE STACK	LANGUAGES
Data Scientist · Yelo Bank	1 yr 5 mos in banking ML	Credit risk · production ML · LLM / RAG	Python · SQL · LightGBM · FastAPI · Docker	RU · AZ · EN

About

Data Scientist at Yelo Bank in Baku. I work on production ML systems the bank uses daily — personalized credit offers and micro-business pre-offer scoring — from the data access layer to the API and monitoring.

My strongest ground is tabular ML for finance — credit scoring, probability-of-default models, interpretable WoE scorecards next to gradient boosting benchmarks. I'm now extending into LLM systems and retrieval-augmented generation.

Russian Native · **Azerbaijani** Full professional · **English** Advanced

Experience

Feb 2026 — Now · 10 mos

Data Scientist

Yelo Bank · Baku

- Full-time role after the internship, continuing work on the bank's production ML systems.

`Python` `SQL` `FastAPI` `Docker` `GitLab` `Grafana`

Jul 2025 — Feb 2026 · 8 mos

Intern Data Scientist

Yelo Bank · Baku

- **Retail Personalized Offer** — model that computes a maximum personalized credit offer per client and predicts whether they can handle it. Rewrote the DB layer, moved to parallel processing with Python multiprocessing (17–20× speedup), added decorator-based logging and automatic result persistence: 5,000 clients went from 14.5 hours to 45 minutes.
- **Micro-scoring Pre-Offer** — engine that finds micro-business clients with unused limits, runs Risk Rules / Stop Factors / Business Rules chains and auto-creates loan leads (refinancing or parallel top-up). Restored it end-to-end: fixed deprecated integrations, rewrote data access, vectorized and parallelized inference.
- Built a FastAPI interface with dynamic toggles for Risk Rules and lead creation. Both systems shipped containerized with Docker, versioned on GitLab and monitored in Grafana.

`Python` `multiprocessing` `SQL` `FastAPI` `Docker` `GitLab` `Grafana`

Sep 2025 — Dec 2025 · 4 mos

Youth STEM Mentor

Algorithmics Global · National School Hackathon

- Mentored five teams at the 48-hour National School Hackathon (Ministry of Science and Education × Algorithmics Global) — 300 finalists selected from 350,000+ participants.
- Three of my teams reached the Top 16; one took 3rd place overall.

`Mentoring` `Hackathon`

Education

2024 — 2026

Data Science & ML Engineer program

Skillfactory

- Applied ML track: classical ML, gradient boosting, neural networks, MLOps basics.
- Project-based: HH.ru vacancy analysis, Booking.com rating prediction, bank marketing, NYC taxi trip duration.

2022 — 2026

B.Sc. Computer Science

Azerbaijan State Oil and Industrial University

- Algorithms and data structures, databases, software engineering.

Stack

scikit-learn · LightGBM · XGBoost · CatBoost · Logistic Regression / Scorecards · Optuna · SHAP · Clustering · K-means · DBSCAN · Anomaly detection · PyTorch · TensorFlow · Keras · LLMs · RAG · Embeddings / Vector search · Topic modeling (LDA) · Python · Pandas · NumPy · SQL / PostgreSQL · Feature engineering · WoE / IV · PCA · t-SNE · UMAP · Matplotlib · Plotly · FastAPI · Flask · Docker · Grafana · Prometheus · CI/CD · GitLab · Python multiprocessing · Jinja2 · Git · Dataiku · Jupyter

Projects

2026 · LLM · Retrieval-Augmented Generation

Tax_RAG

Question answering over tax documentation: documents are chunked and embedded, the most relevant passages are retrieved for each question, and an LLM answers grounded in those sources.

2026 · Credit risk · PD model

Home Credit Default Risk

End-to-end probability-of-default pipeline on 7 Kaggle tables: EDA → WoE/IV feature selection → logistic-regression scorecard with a LightGBM benchmark → out-of-time validation, served through FastAPI with a Jinja2 UI.

0.762	0.524	0.399	0.0001
AUC OOT	GINI	KS	PSI

Certificates

- Machine Learning Specialization — Jul 2025
- OCI 2025 Certified AI Foundations Associate — Jun 2025
- Supervised Machine Learning: Regression and Classification — Jul 2025
- AI & ML — AI Unleashed: From Basics to Neural Brilliance — Jul 2025
- SQL (Advanced) — Jul 2024
- Software Testing — 256 h, 94/100
- Machine Learning in Production — Jul 2025
- Advanced Learning Algorithms — Jul 2025
- Unsupervised Learning, Recommenders, Reinforcement Learning — Jul 2025
- Dataiku Core Designer — Jul 2025
- SQL (Intermediate) — Jul 2024
- Data Science & ML Engineer — In progress — 2026