

VLADISLAV YAROMIY

Data Scientist / AI Engineer

Paris, France | vlad.yaromiy@gmail.com | +33 7 55 98 76 06
linkedin.com/in/yaromiy | github.com/vladndd | Born 29/07/2002

SUMMARY

Data Scientist and AI Engineer with 4 years of experience building and shipping production ML and Generative-AI systems end to end — from data pipelines and model training to LLM/RAG applications, evaluation and monitoring. Founder of Umbra (umbrafeed.com), the infrastructure layer for live tokenized-asset market events. CS background (MIAGE, Université Paris Dauphine-PSL) and a keen interest in financial markets.

SKILLS

Languages	Python, Java, JavaScript, SQL
Machine Learning	scikit-learn, pandas, NumPy, XGBoost, LightGBM, clustering (K-means), classification, NLP, SHAP
Generative AI	LLMs, RAG, prompt engineering, embeddings & vector search, document extraction, LLM evaluation
AI Engineering / MLOps	Model deployment & serving, evaluation & benchmarking, drift monitoring, MLflow, Airflow orchestration, Docker
Cloud & Data	Google Vertex AI, GCP, BigQuery, PostgreSQL, MongoDB
Statistics	Statistical modeling, A/B testing, scoring, risk models
Tools	Git, Unix, Claude, Codex, Jupyter, Matplotlib

EXPERIENCE

Meilleurtaux	2024 – 2026
Data Scientist / AI Engineer	Paris, France
- Designed and shipped an internal RAG assistant over company documentation (chunking, embeddings, vector search, prompt engineering), cutting internal-search time for employees. - Built and deployed a document classification and information-extraction pipeline for client files on GCP / Vertex AI, combining ML classifiers with LLM-based extraction to automate manual processing. - Built AI evaluation and benchmarking frameworks (offline metrics, backtesting, model comparison) and production drift monitoring with MLflow, validating LLM/ML quality before and after release. - Developed an NLP client-sentiment model on customer interactions (classification) to track satisfaction and flag at-risk cases in production. - Engineered ML models for customer segmentation, credit scoring, conversion prediction and anomaly detection (clustering/K-means, XGBoost, LightGBM, scikit-learn, SHAP), served on Vertex AI over BigQuery feature pipelines orchestrated with Airflow and containerized with Docker.	
Product-live	Sept. 2023 – Aug. 2024
Full-Stack Developer	Paris, France
- Designed and optimized <i>Data Hub</i>, a product-data centralization platform (TypeScript, Nest.js, PostgreSQL), reducing time-to-market. - Automated data-quality checks, containerized with Docker.	
Voypost	May 2021 – Jul. 2021
Full-Stack Developer	Remote
- Designed an ERP system (React, Node.js) to optimize workflows and collect key operational data. - Image-loading acceleration (lazy loading, Redis caching), improving system efficiency.	
Yva Int.	Aug. 2019 – May 2021
Full-Stack Developer	Remote
- Built an e-commerce website (React, Node.js, REST API), from data modeling to API implementation. - Integrated the Nova Poshta API to optimize delivery processes.	

EDUCATION

M.Sc. MIAGE — Business Intelligence & Decision Support	2024 – 2026
Université Paris Dauphine-PSL — ML, AI & reasoning, distributed systems, financial analysis	Paris
B.Sc. MIAGE	2023 – 2024
Université Paris Dauphine-PSL — graphs, IS engineering, databases, linear programming	Paris
Université de Lille — B.Sc. Computer Science, Years 1 & 2 (research-intensive track)	2021 – 2023
Scientific Baccalaureate (French High School Diploma, Science)	2019

PROJECTS

- Umbra** — umbrafeed.com (Founder) — **Infrastructure layer** for live tokenized-asset market events: sub-second price spreads, funding-rate shifts and price spikes across **12+ exchanges** and **4 chains** (1,000+ instruments), with human workspaces and an **AI-agent layer**. Built the low-latency ingestion pipeline and a **REST + WebSocket** market-data API (Python, JavaScript).
- PACE 2023 National Challenge** — Solved an algorithmic problem (*twin-width*) in Python. Score: **96.03/100**.

LANGUAGES

French – fluent • English – fluent • Ukrainian – native • Russian – bilingual