

Eugene Shevchuk – AI/ML Engineer, Serbia (remote / relocation)

Eugene is a Data Scientist with 6 years of experience in data science and software development, plus 10 years of prior experience in building teams and management. He is equally proficient in algorithmic development and software engineering. Eugene specializes in full-stack NLP and Computer Vision system design, building inference and AI/ML applications, as well as data, MLOps, QA, and DevOps pipelines. He applies an AI-driven spec-driven development approach, achieving 2-3x faster delivery while maintaining high code quality

EDUCATION

Harvard (HES) Data Science, Masters, Boston, US, 2026 (ongoing)

MIT x Statistics & Data Science, Micromasters, Boston, US, 2020

MIET Applied Mathematics, Bachelor, Moscow, RU, 2012

GPA **3.85/4.0**

Course AVG **94/100**

Course AVG **4.3/5**

EXPERIENCE

Technical Adviser / F15, USA, 12/2022 – Now

Provides strategy advising on AI transformation.

Leads AI/ML Team of 4 engineers.

Skills: Software architecture, QA, Data Analysis, Data Management

Data Science Team Lead / Renmoney, Nigeria, 03/2024 – 01/2025

Built and launched \$90.0000/year in-house serverless risk-engine (roles: Data Scientist/Engineer / Team Lead):

- Reduced annual cost from ~\$90 000 to \$10 000
- Improved deployment speed from 15 days to 3 hours
- Deployed an anti-fraud, and 4 credit scoring models with average GINI > 0.3, LLM-based SMS-parsing module

Responsibilities:

- Manage the team of 3 Data Scientists
- Design, implement, and maintain statistical and LLM ML solutions, ensuring efficient training, inference, and optimization
- Collaborate on project ideas and implementation with IT, Product, Risk Portfolio, Data Warehouse and BI teams

Publication on Risk.net "[How a serverless risk engine transformed a digital bank](#)"

Data Scientist/Engineer (Remote) / Integrated Biometrics, USA, 11/2020 – 12/2023

Built, AWS-based biometric sensor fleet analysis system, processing ~100TB of logs/year, including ML model for sensor wear-out and usage forecasting (roles: Data Scientist / Data Engineer)

Developed backend and mobile app of the Biometric data cross-platform distributed collection/storage/management service (Lead Developer):

Responsibilities:

- Design, test and deploy models and supporting data pipelines and back-end services
- Build high performance backend and front-end with C# (.Net 6, 7)
- Implement and maintain CI/CD pipelines, including unit-, integration and load-testing

Data Scientist / MIMA Time, RU, Moscow, 06/2020 - 11/2020

- **Created** Iron Ore quality prediction model for one of Global TOP 50 steel manufacturers (SciPy, Pandas, docker)
- **Developed** corporate react application for issue management/scheduling (React.js, MySQL, Linux)

Responsibilities - Build high performance backend system API for AI applications with Django

Chief Marketing Officer / Mobsted, RU, Moscow, 09/2018 - 10/2019

Digital Director / Breffi, RU, Moscow, 03/2017 - 08/2018

Global Digital Marketing Officer / FOREO, CN, Shanghai, 01/2015 - 03/2016

Sr. Digital Manager / Samsung Group, RU, Moscow, 03/2011 - 01/2015

APPENDIXES

PUBLICATIONS

02/07/2025 [How a serverless risk engine transformed a digital bank](#) at Risk.net. **Authors** E. Shevchuk, A. Katcalap, N. Dovgiy

TECHNICAL SKILLS

Programming Lang.

	Experience	Level
Python	4 years	Advanced
SQL	6 years	Advanced
C#	2 years	Intermediate
JavaScript/HTML	2 year	Intermediate
R	1 year	Basic
Terraform (HCL)	1 year	Basic

Databases

	Experience	Level
PostgreSQL	3 year	Intermediate
MySQL	2 years	Intermediate
SQLite	2 years	Intermediate

IDE/Development Tools

	Experience	Level
Visual Studio Code	1 year	Advanced
JetBrains EcoSystem	5 years	Advanced
Vim	5 years	Intermediate
MySQL Workbench / DBeaver	1 year	Intermediate

Development Technologies

	Experience	Level
Pandas	4 years	Advanced
Numpy	4 years	Advanced
AWS	4 years	Intermediate
DevOps: GitLab CI / Bitbucket CICD	5 years	Advanced
Bash	4 years	Intermediate
Sklearn	3 years	Advanced
TensorFlow	2 years	Intermediate
Keras	2 years	Intermediate

.Net 5, 6, 7	2 years	Intermediate
Docker	2 years	Intermediate
PowerShell	2 years	Intermediate
MS MAUI	1.5 years	Intermediate
CloudFormation	2.5 years	Advanced
Terraform	1.5 year	Intermediate
PyTorch	1 year	Basic
NLTK	1 years	Intermediate
React.js	1 year	Intermediate
TypeScript	½ year	Basic
SpaCy	½ year	Basic

Issue Tracking / Version Control System	Experience	Level
Jira	7 years	Advanced
Git	7 years	Intermediate
Bitbucket	2 year	Intermediate
GitHub	7 years	Intermediate
GitLab	5 years	Advanced

Development Methodologies	Experience	Level
Scrum	7 years	Advanced
Kanban	7 years	Advanced
Waterfall	6 years	Advanced

Other Skills	Experience	Level
Project Management	15 years	Advanced
Statistics	2 years	Intermediate
Time Series analysis	2 years	Intermediate
Natural Language Processing	2 years	Intermediate
Computer Vision	2 year	Intermediate

PROJECTS

ESG event Stock Impact Detection and analysis AI Platform MVP

Duration	January 2026 - February 2026
Application Type	Full Stack AI Agent Platform

Role	AI/ML Engineer
Description	Built a platform detecting environmental, social, and governance (ESG) violations, analyzing their impact on publicly traded company stock valuations for legal teams and compliance analysts. The system continuously monitors RSS feeds, scrapers target company and regulatory sources websites and news websites for violation signals, then uses NLP and ML to assess the materiality and market impact of detected violations on shareholder value. Built on a multi-hop data architecture with Bronze/Silver/Gold data layers processing through Delta Lake, with a Qdrant powered vector store semantic
Responsibilities	• Complete architecture and implementation: • AI Infrastructure & agents implementation • Backend/Front • Delta Lake • UI/UX
Tools, Technologies	Django, Django REST Framework, PostgreSQL, Prefect, LangChain, LangGraph, Qdrant, DuckDB, Delta Lake, PyArrow, Tavily Search, LangFuse, AWS

AI-Powered Profile Generation Service

Duration	November 2025 - February 2026
Application Type	Async backend API
Role	AI/ML Engineer
Description	Built an AI-driven profile generation service synthesizing detailed user profiles using search/scraping powered AI Agent. It leverages OpenAI's API with template based dynamic prompt construction to create real-data based configuration-driven user profiles. The flow and retry logic is managed with Django, using Celery and Redis for reliable async task processing. The solution includes embedding based semantic analysis benchmark assessing generation quality and text coherence.
Responsibilities	• Design infrastructure • Build AI Agent • Create admin UI • Backend QA
Tools, Technologies	Django, Django ORM, Python, Celery, Redis, OpenAI API, Jinja2, sentence-transformers, NLTK, PyTorch, PostgreSQL

Real-Time Localization in GPS-denied environments

Duration	June 2025 - February 2026
Application Type	Full-stack Discrete-Event Simulation System with Django Web Dashboard
Role	Data Scientist, Full-stack Engineer

Description	Built a discrete-event simulation system for Real-Time Localization in GPS-denied environments. The system models an interconnected network of wireless Tags. They autonomously determine absolute coordinates by exchanging relative positioning data and using Levenberg-Marquardt multilateration position inference from anchor broadcasts, achieving centimeter-level accuracy (<5cm). The simulation includes line-of-sight checks and propagation delay on a radio bus, with position accuracy tracked through coordinate exchange protocols and multi-stratum determination. The solution includes a full-stack analysis simulation results dashboard, stored in Delta Lake, served with Django. The backend provides REST endpoints (DRF Spectacular OpenAPI) for querying Delta tables with DuckDB for simulation filtering and step-level experiment reproduction. The frontend implements 3D visualization using Three.js, rendering tag trajectories, anchor locations, and environment geometry, with interactive camera controls and extended analysis sessions support.
Responsibilities	• Built discrete-event simulation engine • Designed quorum based multi-stratum position determination protocol • Engineered Django REST API with Delta Lake backend • Created analysis pipeline for simulation metrics (error MAE, coordinate success rate) and scenario-based benchmarking with YAML-driven configuration
Tools, Technologies	Python 3.13, salabim (discrete-event simulation), scipy (Levenberg-Marquardt multilateration), numpy, Django 6.0+, Django REST Framework, DuckDB, Delta Lake, pandas, Three.js, SQLite, Docker, docker-compose, pytest, GitLab CI, uv, "

Computer Vision model detecting 3D face coordinates relative to a camera

Duration	April 2025 - July 2025
Application Type	Computer Vision model API
Role	Data Scientist / ML engineer
Description	Developed a model to detect facial points of interest (eyes, nose, and mouth corners), and used them to determine a face's position relative to the camera generating the video feed. The solution uses a single RGB camera and estimates 3D faces locations by RANSAC-based Perspective-n-Point method, matching detected points against a 3D face model. As a result, the system identifies faces and outputs their 3D coordinates to the controller, which adjusts the operator's equipment position relative to their faces in real time. The retrained YOLO11-nano model achieved 0.3–0.4 inch error at 16 FPS without a GPU, and 150 FPS on a Tesla T4 for faces within 5 feet from the camera per one face.
Responsibilities	Finding training dataset, designing and training the model
Tools, Technologies	SkLearn, YOLO11, MTCNN, OpenCV, AWS SageMaker, Docker

Bank risk engine ([featured in risk.net magazine](#) in Jan 2025)

Duration	May 2024 -now
Application Type	Serverless AWS back-end application
Role	ML-Engineer / Team Lead

Description	Built a back-end service to predict default probability on loan applications. It uses about 15 machine learning models, a number of external services, and a set of the bank policy rules. The case was featured in risk.net magazine in January 2025 as an effective and innovative example of digital transformation.
Results:	• Cut infrastructure costs from \$90,000 to \$10,000. • Reduced deployment speed from 15 days to 3 hours. • Improved API performance from 50,000 to 500,000 requests per day. • Reduced downtime from 5–20 hours per week to seconds.
Responsibilities	• Got CEO buy-in • Designed key architecture and ML Ops elements • Wrote ~40% of the core code base. • Restructured the DS team and external stakeholder processes, including teaching basic level python to one of the stakeholder teams (5 people).
Tools, Technologies	Stack: AWS (multiple services), Cloud Formation, Git, CI/CD ML: LogReg, CatBoost, Custom Statistical Models

Fraud detection score model

Duration	Sep 2024 - May 2025
Application Type	Bank risk engine model
Role	Data Scientist / ML-engineer
Description	Used external fraud prediction score and user feature data to build an in-house scoring model. The goal was to reduce the number of paid requests to the vendor. Built a CatBoost model using about 300 features. Achieved an initial 0.85 AUC .
Tools, Technologies	Stack: SkLearn Pipelines, AWS SageMaker, CloudFormation ML: CatBoost

LLM-based SMS Parsing solution

Duration	Sep 2024 - June 2025
Application Type	Bank risk engine AWS Lambda
Role	ML-Engineer
Description	Developed a solution that parses transactional SMS from client phones to predict their wealth level. It uses the OpenAI API to create regular expressions for extracting transaction amounts from SMS. Most clients receive messages from fewer than 15 banks, and these banks change their SMS format less than twice a year. This method allowed us to process millions of SMS with only a few hundred paid API requests, using the generated regular expressions and calling the paid API automatically only when a bank changed its SMS format.
Tools, Technologies	Stack: AWS Lambda, DynamoDB, RedShift, Open AI API

LLM-based JIRA taks timeline bias prediction model

Duration	Sep 2024 - Dec 2024
Application Type	API

Role	Data Scientist / ML-engineer
Description	Build a model to correct JIRA task estimation bias. Used a dataset of tasks with planned and actual durations, and context text features like comments, priority, and participants. Achieved MAE 0.3
Tools, Technologies	Stack: PyTorch, SkLearn Pipelines, Docker ML: HuggingFace, CatBoost, TF-IDF, NLTK

Distributed biometric data collection & management service + mobile app

Duration	July 2022 - Dec 2023
Application Type	Distributed back-end + Cross-platform mobile app
Role	Full-stack Engineer / DevOps
Description	The product is a cross-platform library for Android, Windows, and Linux with wrapping eco-system: service to run standalone and wrapping mobile app, to use the UI it provides. It stores fingerprint data and syncs this data between all devices with the library installed. The library also includes a fingerprint matcher, firmware, and a UI for capturing fingerprints when a sensor is connected via USB to the device. Typical use-case. A group of operators goes to a refugee camp to identify people there. When they return to the field office and connect to Wi-Fi, the data from their mobile devices syncs to the workstation. Workstations from different locations sync data to a central cluster in the capital. Each mobile device, workstation, and server runs the same library, built from a single code base. This library and engulfing eco-system is what the project delivered.
Tools, Technologies	.Net 6, 7, 8, MAUI, SQLite, MySQL, AWS EC2, GitLab CI, WiX Toolset, unit/integration tests & CI/CD, Windows services

Biometric sensor fleet management log analysis solution

Duration	Nov 2020 - Jun 2022
Application Type	API
Role	Data Scientist / Data Engineer
Description	The solution collects log data from individual fingerprint sensors using an installed connector. It sends the collected data through a chain of intermediate servers to a central server. It also provides about 40 dashboards showing the current state of the sensor fleet. These dashboards display logs gathered at each node. Users can see usage statistics, device wear indicators, and predictions for wear-out. The solution is implemented with Python, TensorFlow, and PostgreSQL.
Tools, Technologies	Python, Grafana, Postgres SQL, Docker, AWS EC2, ECR, ECS, Unit/integration/load tests & CI/CD, Terraform

Iron Ore quality prediction model

Duration	Jun 2020 - Nov 2020
Application Type	API
Role	Data Scientist / ML-engineer

Iron Ore quality prediction model

Description	1-st place competition winning model using data from 60 temperature sensor array to replace expensive lab test iron ore quality prediction.
Tools, Technologies	SciPy, Pandas, Docker

Corporate task manager web App

Duration	Jun 2020 - Oct 2020
Application Type	React.js App + Django Back-end
Role	Full-stack developer
Description	Corporate task manager with 12 screens UI and a backend
Tools, Technologies	React.js, Django, AWS EC2

EDUCATION & COURSES

Harvard University, Data Science Master's degree 09/2023 – Now

Computer Vision (CSCI E-25)
Natural Language Processing (CSCI E-89B)
Advanced Topics in Data Science (CSCI-109B)
Foundations of Data Science and Engineering (CSCI-101)
Machine Learning with Python: Linear to Deep Learning (6.431x)
Fundamentals of Statistics (18.6501x)
Data Analysis in Social Science (14.310Fx)
Probability - The Science of Uncertainty and Data (6.431x)

08/2023, Udemy: The Complete ASP.NET MVC 5
08/2023, Udemy: Multithreading and Parallel Programming in C#
08/2023, Udemy: C# Intermediate: Classes, Interfaces and OOP
08/2023, Udemy: C# Advanced Topics

07/2021, Udemy: Build a Backend REST API with Django - Advanced

06/2021, LondonApp Developer: DevOps Deployment Automation with Terraform, AWS and Docker

05/2020, MIT x: Statistics & Data Science MicroMasters program

Python Programming, Machine Learning, Data Analysis, Statistical Modeling, Probability, Statistical Inference, Deep Learning, Natural Language Processing, Computer Vision, Data Visualization, Hypothesis Testing, Bayesian Inference, Linear Regression, Supervised Learning, Unsupervised Learning, Neural Networks, Data Preprocessing, Feature Engineering, Model Evaluation, Clustering, Random Variables, Conditional Probability, Central Limit Theorem, Markov Chains, Text Processing, Image Processing, Dimensionality Reduction, Time Series Analysis, Causal Inference, Expectation Maximization, R Programming, Monte Carlo Simulations methods

C# Programming, Object-Oriented Programming, ASP.NET MVC, Web Development, Backend Development, Multithreading, Parallel Programming, Entity Framework, LINQ, Asynchronous Programming, REST API Development, Authentication, Authorization, Database Management, Software Architecture, Software Development, Polymorphism, Inheritance, Encapsulation, Delegates, Lambda Expressions, Generics, Task Parallel Library, Concurrent Collections, Data Validation, Routing, Controllers, Views, Models, Extension Methods

Python · Django · Django REST Framework · Docker · GitLab CI/CD · PostgreSQL · Test Driven Development · REST API Development · Backend Development · Web Development · Authentication · Database Management · Continuous Integration · Deployment · Software Testing

Docker · AWS ECS · AWS RDS · AWS Application Balancer · AWS VPC · AWS Route 53 · AWS IAM · AWS Certificate Manager · AWS S3 · CICD · Infrastructure as code (IaC)

Data Visualization · NumPy · PCA · Git · Pandas · Analytical Skills · Data Science · python · Scikit-Learn · Machine Learning · R · PyTorch

05/2020, Microsoft: Certificate: Querying Data with Transact-SQL

Transact-SQL, SQL, Microsoft SQL Server, Azure SQL Database, Data Querying, Data Modification, SELECT Statements, JOIN Operations, Set Operators, Aggregate Functions, Subqueries, Table Expressions, Data Pivoting, Stored Procedures, Error Handling, Transaction Management, Database Management