Nataliia P.

AI/ML Engineer /
Back-End Developer

Summary of Qualifications	I've been working with Python for 5.5 years and focus mainly on AI and LLM-based systems, and I'm crazy about coding. I build real, production-ready solutions using RAG, multi-agent workflows, and chat or voice AI. I work a lot with LangChain, LangGraph, and LlamaIndex, fine-tune and deploy models on AWS, and integrate AI into backend services so it can reason, use tools, and automate tasks. I also often take a lead role on AI projects, helping define the architecture and guiding how AI is used in the product. I have a strong backend background with FastAPI, Django, and Flask, understand cloud deployment and databases, and hold AWS Cloud Practitioner and AWS AI Practitioner certifications. I learn fast, communicate clearly, and enjoy building AI systems that actually work in real life. Also, previously I was working in travel industry, I'm crazy about traveling, so at the moment, I have visited 27 countries.

Skills

Programming Languages/Technologies

- SaaS/RIA Research & Development
- Asynchronous Programming
- Technical Documentation/Architecture Reports
- XML/JSON/YAML
- Python

Developer Productivity AI Tools

- ChatGPT
- GitHub Copilot

RDBMS

- MySQL
- PostgreSQL
- Alembic

NoSQL

- Redis

Virtualization Tools

- Docker / Docker Compose

Methodologies

- Agile, Scrum, Kanban
- Pair Programming

Operating Systems

- Microsoft Windows
- Debian/Ubuntu/ Arch/Linux

CI/CD

- GitHub CI/CD
- GitHub Actions

Frameworks/Libraries

- FastAPI
- Django/DRF
- SQLAlchemy
- Numpy/Pandas
- Requests
- Aiohttp
- Matplotlib
- Selenium
- BeautifulSoup
- LangChain
- LangGraph
- LlamaIndex
- Whisper
- Faster Whisper
- MCP Server Connectors
- ElevenLabs
- Telegram API
- Google AppScripts
- Starlette

Application/Web Servers

- Nginx

Cloud Providers

- AWS
- GCP

Development Tools

- Visual Studio
- PyCharm

Testing Tools

- Postman

	Versions Control ▪ Git ▪ GitHub ▪ GitLab	▪ Pytest ▪ Unitest, Mock

Experience

AI Knowledge Graph Platform (Ontologies) for Persistent Organizational Memory

Project Description: Designed and built an AI-powered knowledge platform that serves as a company's long-term memory by maintaining a persistent, evolving knowledge graph. The system enables teams to connect multiple internal and external data sources via MCP connectors, define domain-specific schemas, and continuously extract, structure, and link information into a unified graph.

Each project operates within its own isolated environment, maintaining a dedicated knowledge graph with clearly defined entities, relationships, and full source provenance. Every stored fact can be traced back to its original document, ensuring transparency, auditability, and trust in the generated insights.

The platform supports iterative, conversational search workflows where users can ask natural language questions, explore related entities, and build context across multiple queries. Search results are written back into the knowledge graph, allowing it to grow richer over time and continuously improve retrieval quality. The system is designed to scale to hundreds of thousands of entities per project while preserving performance and strict tenant-level data isolation.

Domain: AI & ML | Knowledge Graphs | Enterprise Search

Involvement Duration: 0.5 years

Project Role: AI & ML Engineer | Back-End Developer

Responsibilities:
- Implemented MCP-based connectors to ingest data from multiple internal and external sources;
- Developed schema definition and validation mechanisms tailored to different business domains;
- Built pipelines for continuous data extraction, indexing, entity resolution, and relationship linking;
- Implemented full source provenance tracking to ensure traceability of all facts back to original documents;
- Designed iterative search sessions supporting natural language queries and multi-step exploration of related entities;
- Optimized system performance to support large-scale graphs with hundreds of thousands of entities while enforcing strict tenant-level isolation.

Project Team Size: 2-6 team members

Tools & Technologies: FastAPI, Python, Youtu GraphRag, PostgreSQL, MindsDB, MCP, NetworkX, Neo4j.

Multi-Channel AI Voice & Chat Sales System (Web Voice, Chatbot, Outbound, Cold Calling)

Project Description: Built a unified AI-driven voice and chat system consisting of four specialized bots, all operating on a shared backend and fully synchronized with a single CRM. The system is designed to engage prospects across multiple touchpoints – website, outbound calls, instant follow-ups, and in-browser voice conversations – while maintaining a consistent context and lead history.

	Each bot handles a specific stage of the customer journey, but all share the same conversational intelligence, real-time decision-making, and CRM automation layer. The result is a scalable, always-on sales assistant that behaves like a coordinated team rather than isolated bots.
Domain:	Conversational AI | Sales Automation | Voice AI
Involvement Duration:	1 year
Project Role:	AI Engineer | Back-End Developer
Responsibilities:	▪ Designed and implemented a multi-bot architecture with a shared backend and centralized CRM integration; ▪ Developed a website chatbot that acts as a virtual assistant, answering service-related questions, engaging visitors in natural dialogue, collecting lead data, and qualifying users in real time; ▪ Built an autonomous cold-calling voice bot that pulls contacts from GoHighLevel, places outbound calls, explains offers, handles objections, and adapts the conversation flow dynamically like a human sales representative; ▪ Implemented an outbound follow-up voice bot triggered by website form submissions, automatically calling users, continuing the conversation by voice, qualifying leads, and advancing them through the sales pipeline; ▪ Developed a low-latency web-based voice bot using WebSocket and WebRTC, enabling real-time voice conversations directly from the browser; ▪ Integrated OpenAI's Realtime model across all bots for natural, live voice interactions; ▪ Implemented deep GoHighLevel integration for contact creation and updates, appointment scheduling based on availability, automated email meeting invites, and storage of full call transcripts, summaries, and audio recordings; ▪ Integrated ElevenLabs as an optional TTS engine to provide more natural, human-like voice output.
Project Team Size:	5-8 team members
Tools & Technologies:	OpenAI RealTime, LlamaIndex, ChromaDB, GoHighLevel, Twilio, ElevenLabs, AWS S3, Redis, Python, FastAPI, Docker.
	AI-Powered Real-Time Voice Assistant with CRM Automation (AI Receptionist)
Project Description:	Designed and implemented a real-time AI voice assistant powered by OpenAI's latest Realtime model, capable of conducting natural, human-like voice conversations. The assistant answers questions about company services, qualifies leads, collects and stores customer data, and performs automated follow-ups. The system integrates tightly with CRM platforms to streamline sales and support workflows, while maintaining full call transcripts and summaries for future reference.
Domain:	AI & ML | Conversational AI | Sales & Marketing Automation
Involvement Duration:	0.5 years
Project Role:	AI Engineer | Back-End Developer
Responsibilities:	▪ Developed a real-time voice conversational assistant using OpenAI Realtime models; ▪ Implemented CRM integrations (Airtable, GoHighLevel) for automatic contact creation and updates; ▪ Built logic for lead qualification and conditional workflow execution; ▪ Integrated SMS and email delivery with secure file upload links, enabling the AI to read and reference user-provided documents during calls; ▪ Implemented appointment scheduling directly in GoHighLevel based on real-time availability;

	▪ Designed post-call processing to store full call transcripts and AI-generated summaries in the CRM, including caller phone numbers; ▪ Prepared optional ElevenLabs TTS integration for more natural, human-like speech output.
Project Team Size:	1-3 team members
Tools & Technologies:	Airtable, N8N, AWS S3, OpenAI RealTime, FastAPI, GoHighLevel, Twilio, ElevenLabs, Python, Outlook, Docker, FAISS.

Modular AI Platform for Data, Media & Intelligent Services

Project Description:	A modular AI-driven ecosystem that unifies data monetization, conversational orchestration, agent-based automation, and analytical platforms into a single scalable infrastructure. The system enables users to transform activity data, documents, and digital media into structured assets that can be analyzed, monetized, or consumed by AI agents. The platform combines conversational intent routing across multiple services, tools for building single- and multi-agent workflows, graph-based legal and compliance analytics, and OCR-driven financial document analysis. All components share a service-oriented architecture focused on extensibility, data integrity, and high-performance processing at scale.
Domain:	AI & ML | Data Platforms | Conversational AI
Involvement Duration:	1.5 years
Project Role:	Back-End Engineer | AI Engineer
Responsibilities:	▪ Participated in implementation of a modular, service-oriented architecture supporting multiple AI-driven products within a single ecosystem; ▪ Built data ingestion and asset-generation pipelines for transforming user activity, documents, and media into structured, monetizable assets; ▪ Implemented conversational orchestration logic to route user intent across heterogeneous services through a unified dialogue interface; ▪ Developed AI agent infrastructure supporting single-agent and multi-agent workflows, including hierarchical orchestration and visual workflow modeling; ▪ Designed graph-based data models and analytics for entity resolution, relationship discovery, and conflict-of-interest detection; ▪ Implemented OCR and domain-specific information extraction pipelines for large-scale financial document analysis; ▪ Ensured scalability, performance, and isolation across services handling large datasets and concurrent real-time interactions.
Project Team Size:	10-15 team members
Tools & Technologies:	Python, FastAPI, PostgreSQL, Redis, Celery, RabbitMQ, SocketIO, WebSockets, NetworkX, FalkorDB, LlamaIndex, OpenAI, LiteLLM, MCP, ParadeDB, BM25, AWS Textract, Mistral LLM.

Mental health chatbot

Project Description:	Chatbot is designed to helped people with their mental problems. It consists of well-trained LLM model that is totally secure. The bot can interact with users verbally and by text messages in a real time. It has memory with previous conversations. It can simulate natural conversation and offer guided therapeutic exercises.
Domain:	AI & ML | Mental health
Involvement Duration:	1 year
Project Role:	AI & ML Engineer | Back-End Developer
Responsibilities:	▪ Back-end development; ▪ Prompt engineering; ▪ Development of streaming text-to-speech functionality;

- Real-time speech-to-text pipeline implementation;
- Development of streaming text-to-text functionality;
- Fine-tuning LLM model;
- Attach memory to the LLM;
- Development a simple front-end for the chatbot.

Project Team Size: 3-4 team members

Tools & Technologies: Python, FastAPI, Whisper, React, ElevenLabs, Llama, Langchain, AWS Bedrock, AWS Sagemaker, PostgreSQL, Docker.

Sales bot

Project Description: The project is a software tool that, using a filter system, selects applications that do not need to be processed by sales managers. This system uses several criteria such as customer interests, geographic location, technical skills, and other factors to accurately select potential customers.

Domain: Automation | Sales | AI&ML

Involvement Duration: 1 year

Project Role: Back-End Developer

Responsibilities:

- Back-End Development;
- Prompt engineering;
- Refactoring exciting code;
- Scrapping and parsing;
- Fixing validation errors;
- Optimized work by dividing sync and async code to different threads;
- Deployment on AWS;
- Improving matching system;
- Moved message chats and customized them at different platforms through GCP;

Was working with app scripts to increase functionality to message chats.

Project Team Size: 4-5 team members

Tools & Technologies: Python, Starlette, Selenium, AWS, GCP, BS4, Google AppScripts, async programming, Pydentic, SQL Alchemy, PostgreSQL, Django, Docker.

Certificates

AWS AI Practitioner
AWS Certified

AWS Cloud Practitioner
AWS Certified

Python Course
Mate academy

Education

Master's degree in international economy
National university "Lviv Polytechnic"

Languages

English – Advanced
Estonian – Intermediate
Ukrainian – Native
Russian – Native