

Rusu Dinu-Ştefan

rusudinustefan@gmail.com | rusudinu.ro | github.com/rusudinu | linkedin.com/in/rusudinu/
+40756478663 | Bucharest, Romania

Skills

Programming Languages: Java, Kotlin, Dart, TypeScript, JavaScript, Python, Solidity

Frameworks & Technologies: Spring Boot, Flutter, React, React Native, Angular, Liquibase, GraphQL, RabbitMQ, Redis, Spring MCP, FastMCP, FastAPI, LangChain, LangGraph

Tools & Platforms: AWS, Google Cloud, App Engine, VertexAI, OpenRouter, Kubernetes, ArgoCD, Docker, GitHub Actions, Grafana, Maven, Gradle, Confluence, Jira, Prompt Engineering, LMStudio, Ollama, Snowflake, Serverless, Feature Flags, Ethereum, Hyperledger Fabric, PostgreSQL, pgvector, Elasticsearch, MongoDB, SQL, NoSQL, Linux

Soft Skills: Teamwork, Communication, Leadership, Mentoring, Public Speaking, Time Management

Experience

Contracted Senior Engineer, Dell – Bucharest, Romania

Dec 2024 – Present

- Played key role in migrating users from authentication microservice to unified Dell Identity service with alerts and retry mechanisms
- Built suite of Kotlin libraries for handling throwing, retrying, and alerting on Teams for RabbitMQ listeners with complete JUnit5 and Mockito test suite
- Contributed to AI team tribe meetings, improved prompts for LangChain Agent for automated Jira ticket and code generation;
- Used Windsurf IDE AI for AI assisted coding

Senior Software Engineer, Tazz by eMAG (now Wolt) – Bucharest, Romania

Aug 2021 – Dec 2024

- Architected and built a **RAG-based AI chatbot** using **Python**, **FastAPI**, and **LangGraph** with advanced function calling capabilities, enabling natural-language querying of business data stored in **Snowflake**
- Designed the chatbot's agentic workflow with multi-step reasoning, automatic SQL generation, and context-aware retrieval to deliver accurate, real-time analytics to internal stakeholders
- Integrated tool-calling patterns allowing the LLM agent to autonomously select and execute Snowflake queries, aggregate results, and present human-readable insights without manual intervention
- Finished the Tazz Direct (Logistics as a Service) service in under three months
- Built Spring Boot microservice orchestrating automatic dispatching algorithm with bundled couriers and orders
- Developed Taxonomy system for dynamic category creation and custom product attributes for ML team
- Designed architecture for custom access control system with managed permissions elevation, coordinating 2 junior developers
- Engineered microservice tracking KPIs and triggering custom system alerts, preventing data loss across Tazz applications

Freelance, CodingShadows – Bucharest, Romania

Sep 2018 – Present

- Led and developed alongside teams of engineers comprehensive web and mobile applications, launched 45+ apps using Flutter across Android and iOS, ranging from Electronic Registers integrations to People Counting via Surveillance Cameras with real-time GPS tracking and data synchronization and built several decentralized apps (dApps) on Ethereum and Hyperledger Fabric blockchains
- Ensured scalability, performance and observability by implementing CI/CD pipelines with GitHub Actions, Docker, and Kubernetes on Hetzner Cloud, with monitoring via Prometheus, Grafana and Sentry
- Built 5+ ChatGpt apps, using **Python and FastMCP** (initially built with Spring MCP) as well as agentic chatbots using LangChain and LangGraph + Firebase AI with Tool Calling and Gemini Live API; used AI coding tools: **Junie**, **Claude Code**, and **Cursor**
- Set up local (private) LLMs using **LMStudio** and **Ollama**, connected to them using Python and **Spring AI**, Json Schema-enforced outputs, also used them for generating **embeddings**

Teaching Assistant & Research Technician, Politehnica University of Bucharest Sept 2024 – Present

- Held laboratories for Mobile and Embedded Computing, Software Development Methods, Object-Oriented Programming, and others
- "Tehnologii bazate pe inteligență artificială pentru soluții de încredere împotriva dezinformării" – implemented **Elasticsearch** vector search for semantic document retrieval

Contracted Trainer, Deutsche Bank Cloud School Mar – Sept 2022, 2023, 2024, 2025

- Taught and led Deutsche Bank juniors in Java and Spring Boot development
- Led full-stack project using Spring Boot and React for automated CV analysis and candidate reporting
- Created automatically graded assignments using GitHub Classroom with JUnit5 tests

Education

Master of Software Engineering – Politehnica University of Bucharest Oct 2024 – Jul 2026

Bachelor of Computer Engineering – Politehnica University of Bucharest Oct 2020 – Jul 2024

Honors & Awards

- **Speaker at Devtalks** – "Scaling Flutter to 100k MAU as a solo developer" in 2025
- **Speaker at DevCon** – "Building for Fun, Scaling for Profit: Lessons from a Decade of Indie App Development" in 2025
- **8 published papers** on Blockchain technology integrated with LLMs
- Invited to **PRO TV** (ILikeIT) in 2019, 2022, and 2024; **Euronews** in 2024
- SCSS - UPB, **1st place** in 2022 and **2nd place** in 2021
- "InfoEducatie" National Olympiad, **2nd place** in 2020, and **top 15%** in 2019

Volunteering

- Member of the Scientific Committee (Subcomisia Științifică) and Technical Committee (Subcomisia Tehnică) for the National Romanian AI Olympiad – solely architected and maintained the competition infrastructure, provisioning bare metal Kubernetes clusters on VPS and also building a GPU cluster with NVIDIA GPU Operator and JupyterHub, achieving **zero downtime** throughout all rounds
- Judge & Mentor at the "InfoEducatie" National Olympiad in 2021 and 2022
- Organizer & Judge at the "Hardcore Entrepreneur" UPB contest in 2021, 2022, 2023, 2024, and 2025

Other Work

EHR Blockchain & LLM-based Platform

- Created web platform for storing Electronic Health Records using Hyperledger Fabric blockchain and Llama3 LLM chatbot, developed using Angular frontend with two NestJS microservices connecting to PostgreSQL and MongoDB databases
- Integrated **LLM with function calling** to read blockchain records, enabling users to view and query their health data through a conversational chat-based UI
- Implemented RBAC by using Keycloak authentication with different user roles (doctor, patient, pharmacist) and adaptive UI

Scanner Alimente - scanneralimente.ro

- Offline mobile app scanning food ingredients and scoring nutritional value, reaching #1 in Health & Fitness
- Built Python pipeline for ingredient matching (name to barcode), ingredient title healing, and additive detection from labels
- Integrated **Gemini Live API** for real-time voice conversations about additives and implemented chat mode with **function calling powered by Google VertexAI**; backend LLM routing via **OpenRouter**

- Implemented recipe search using **vector search** via **pgvector** for semantic ingredient and recipe matching
- Peaked at 100k monthly active users using Flutter, Spring Boot, Liquibase, PostgreSQL on Kubernetes

Divigo

- Oversaw product development with a customer-oriented mindset, driving features from ideation to delivery
- Built cross-platform mobile app in **React Native** with electronic register integration for receipt emission and **Viva Wallet** card payment processing
- Developed web platform connecting to electronic registers for ticket sales with automated receipt generation
- Implemented real-time bus tracking on map with zone-in/zone-out detection for bus positioning
- Integrated wirelessly connected surveillance cameras (RTSP protocol) to capture images on zone events
- Built **Python** people counting pipeline using **OpenRouter** with two different models for error balancing, comparing passenger counts against tickets sold

Comentarii BAC Romana si eseuri audio

- Top 5 trending educational apps on Google Play (2021) and top 10 on App Store (2024), cross-platform educational app with essays for national exam, 25,000+ total downloads, 3,000 monthly active users

Kubernetes Cluster using Raspberry Pi 4s & Dell PowerEdge R610

- Established high-performance Kubernetes cluster with three Raspberry Pi 4s for auto-scaling microservices
- Set up RHEL8, NGINX, Prometheus, Grafana, MongoDB, Elasticsearch with custom fan controls and auto-restart, now migrated to Hetzner Cloud VPS.

Gomoku MiniMax

- Gomoku game in Scala with MiniMax, Alpha-Beta pruning, and memoization
- Played against colleagues and teacher's AI, resulting in draws and wins for the implementation

Game Assembler & Emulator - "Fantasy Game Console"

"Infoeducatie" Hackathon 2020 - 24h

- Engineered an Assembler and Emulator with a teammate, written in C++, for packing and playing games on an emulated old console which had 128 bytes of RAM, 4000 bytes of ROM, 3840 VRAM, and a resolution of 40 x 192 pixels with 6507 instructions per frame
- 6 general-purpose registers and 4 special-purpose registers for program counter, stack indicator, input data, and utility