Jegor Kitskerkin

jegor.kitskerkin@gmail.com / SOFIA, BULGARIA

Results-driven professional with 5 years of experience (+2 years of part-time freelancing) in software development and computer science. Adept at leveraging Python for machine learning, deep learning, and data analysis, alongside extensive frontend and backend development expertise using modern frameworks and languages. Fan of Rust.

Work experience

Sr. Software Development Engineer 04/2024 - 05/2024

PowerByAI / Part-time / Remote

- Developed the frontend in collaboration with the UI/UX specialist for the financial LLM solution, including the chat interface, graphs, access control interface and analytics dashboards.
- Collaborated on the development of the backend for the LLM solution, driving the development of the access control system, authentication system, internal chatbot framework, LLM- and DB-integration layer and a microservice-based inference server.
- Architectured and implemented the deployment to the Azure cloud of the distributed backend component and the frontend via GitOps infrastructure-as-code using Terraform.

Software Development Engineer II 11/2023 - Present

LivePerson / Sofia, Bulgaria

- Reworked the inference architecture by replacing some of the custom software components to cloud-native (GCP) offerings for easier deployments and decreased maintenance requirements.
- Enhanced the automatic deployment of ML models, backed by Kubernetes and FluxCD. Lead to decreased error rate compared, offloaded the engineering team and improved time-to-market for our ML models.
- Developed an internal LLM benchmarking tool to promote widespread use and enhance the CI pipeline for data-centric model assessments.
- Redesigned the inference system to better parallelize preprocessing and request handling, resulting in 2-3x higher throughput under high load.

Machine Learning Engineer 03/2023 - 11/2023

Experian / Sofia, Bulgaria

- Streamlined the deployment process through DevOps and MLOps techniques (i.e. model serving optimization, asynchronous programming, model loading optimization, streamlining model CI in Jenkins)
- Created internal tools and dashboards for data annotators using TypeScript and Angular
- To promote scalability and fault tolerance, contributed to migrating a monolithic application to microservices backed by Kafka and NATS

Machine Learning Engineer 04/2020 - 02/2023

GeniusVoice / Tilburg, Netherlands

- Collaborated in a two-person team to develop a multi-component search system that combined ranking algorithms, cross- and bi-encoder Transformer models – enhancing top-1 accuracy by 20-30%.
- Designed interactive chatbots and orchestration solutions deployed across websites of multiple municipalities in the Netherlands, integrating the (previously mentioned) search system for enhanced functionality.
- Pioneered an approach for a span-extractive QA based on Transformer attention scores, achieving a 20x speed increase with only an 8% reduction in KPIs.

Fullstack Software Engineer 10/2018 - 08/2020

Freelance / Tallinn, Estonia

- Developed and deployed web applications using a tech stack of Vue.js, PostgreSQL, Python Django, and Spring Boot.

Education

Minor in Data Science, Eindhoven University 2021

BSc in Cognitive Science and Artificial Intelligence, Tilburg University 2019-2022

- Graduated Cum Laude (8.1/10 GPA)
- Highest scoring noise reduction model among the cohort for the Deep Learning class project
- Thesis covering Transformer-based video models for human action recognition

Personal achievements

- **Contributing Video Vision Transformer to HuggingFace.** After using the model for my Bachelor thesis, I have made an effort to add it to make it more accessible for others. View PR
- My fine-tuned Dutch language models have been used as part of a research and appeared in the Journal of Computer Assisted Learning (not my achievement per se, but I was pleased to find out) Read the paper here

Internships

Research Intern, Tilburg University 08/2021 - 01/2022

- To enhance existing dysgraphia detection ML solution, systematically compared and evaluated different representation techniques, applying advanced concepts including: interpolation, Bezier curve estimation, Fourier Transform, and the Ramer-Douglas-Peucker algorithm. These evaluations provided valuable insights into the most effective approaches for representing handwritten data.

Skills

Programming Languages	Python, TypeScript, TypeScript, SQL, Go, Rust
Data	PostgreSQL, ELK stack, Mongo, Redis, Kafka, Vector DBs
Frontend	React, Remix.js, Next.js
Dev/MLOps	Docker, Kubernetes, MLFlow, DVC, CI/CD, FluxCD
Cloud	GCP (1 year), AWS (1.5 years), Azure (6-9 months)