

MUHAMMED BERK ÖNDER

AI/ML Engineer

☎ (+90) 545-528-3371 | ✉ berkmonder@gmail.com | 🌐 berkmonder | 📄 in/berkmonder

AI/ML Engineer with experience building, deploying, and scaling production ML and LLM systems in Telecommunications and Software/Mobile Applications industries. Specialized in Anomaly Detection, large-scale data processing, and LLM-powered systems and data pipelines, with strong expertise in Python, Docker, Git, Airflow, GCP, and modern LLM frameworks (LangChain, Ollama).

Experience

MLOps Engineer

Teknasyon

08/2024 - Present

- Designed and deployed a modular RAG-based LLM system using LangChain and Open WebUI, enabling high-precision retrieval from structured internal data via custom Python pipelines and tool integrations.
- Scaled a production spam detection pipeline of millions of weekly records by segmenting data processing, eliminating memory bottlenecks and improving stability.
- Built automated data pipelines for anomaly detection using Dataform and Airflow on GCP, enabling scalable ingestion of high-cardinality event data and reducing manual query overhead.
- Integrated MLflow into Airflow pipelines to standardize experiment tracking, improving model observability and reproducibility.
- Optimized fine-tuning of ~30B parameter Qwen models on RunPod, reducing training time and infrastructure costs through better GPU utilization.
- Implemented CI/CD pipelines with GitHub, Cloud Build, and Docker to produce versioned model artifacts and enable reproducible deployments.
- Processed high-cardinality event data across thousands of dimensions, improving pipeline scalability and reliability.
- Containerized ML and LLM services with Docker for consistent, portable deployments across environments.

Machine Learning Engineer

P.I. Works

07/2022 - 07/2024

- Built a telco LLM assistant using LangChain, enabling knowledge base retrieval, SQL-based data access, analysis, and visualization within a unified interface.
- Developed a heatmap forecasting system using regression models to predict network traffic from football match data, reducing aggregate error from 77% to 0.75%.
- Implemented multivariate anomaly detection on streaming time-series data using clustering techniques, improving detection across diverse scenarios.
- Designed a synthetic data generation framework using Monte Carlo methods, increasing data granularity by 60x.

Education

İstanbul University

MSc, High Energy and Plasma Physics GPA: 3.63

Graduation Date: 2024

BS, Physics GPA: 3.31

Graduation Date: 2021

Skills

Programming: 🐍 Python, 🗄️ SQL, 🐚 Bash, 🟢 VIM, 🖥️ Tmux

ML/LLM Frameworks: 🍌 Scikit-Learn, MLflow, 📈 TensorFlow, 🔥 PyTorch, ⭐ PySpark, 🤗 Huggingface, LangChain, LangGraph, 🦙 Ollama, LMStudio, vLLM, Unsloth

Python Packages: 📊 Pandas/Polars, 📦 NumPy, 📉 Matplotlib, 📊 Plotly, 🖨️ Typer, 🌐 NetworkX

MLOps & Infra: 📁 Git, 🐳 Docker, 🔄 CI/CD, ⚡ FastAPI, 🌪️ Airflow, 🌐 GCP, 🐧 Linux, 🟣 uv, ⚙️ Nix, Prefect, RunPod, 🐳 Podman

Languages: Turkish (native), English (proficient)