

Aymen Fkir

✉ aymenfkir23@gmail.com | 📞 +216 29-581-238 | 🔗 [linkedin.com/in/aymen-fkir](https://www.linkedin.com/in/aymen-fkir) | 🐙 github.com/aymen-fkir

portfolio-coral-ten-79.vercel.app

Education

École Supérieure Privée d'Ingénierie et de Technologies, computer engineer August 2025 - June 2028
Higher Institute of computer science and Multimedia Gabes , Bachelor in Computer Science August 2022 - June 2025

Experience

Machine Learning Engineer, Evimo - Remote December 2025 - May 2026

- **Architected** an AI-driven Incident Management pipeline, refactoring 9+ microservices into an asynchronous system that utilized **asyncio.to_thread** to offload CPU-bound ML inference, preventing event-loop starvation.
- **Created a secure RS256-based authentication** service using FastAPI and Pydantic v2, implementing asymmetric key signatures to provide stateless, non-repudiable access for technician-facing interfaces.
- **Built** a type-safe data extraction layer leveraging LLM parsers (Llama/Docling) and Python Generics, ensuring 100% schema validation between unstructured logs and production data models.
- **Optimized** system reliability and performance by implementing **asyncio.Lock** for thread-safe data persistence and migrating to uv for deterministic, near-instantaneous dependency management .

Machine Learning Engineer, Omdena - Remote October 2025 - December 2025

- **Architected** an extensible Python synthetic data framework **using OOP principles and inheritance** to generate diverse, multi-domain logs while **ensuring seeded determinism for reproducible model validation** .
- **Co-led a data team of 10 collaborators**, managing task delegation and technical workflows to deliver high-quality, schema-aligned datasets for the classification pipeline .
- **orchestrated CLI tooling and Makefile integrations** to streamline the data generation process, enabling users to produce structured datasets via simple command-line interfaces .
- **Engineered robust CI and QA workflows** by implementing unit tests for data quality and determinism, while debugging pipeline errors to ensure stable project builds.

Data Scientist , Deepvolt - Tunis February 2025 - June 2025

- **Engineered** a scalable **etl** pipeline with **polars** and **ec2**, processing **20+ GB** of floating car Data **cutting costs by £40K**.
- **Led dataset design** for a robust traffic density predictive model achieving **90% of variance (R2)** with a **low error rate 35%** in traffic prediction using **polars ,python** and **sklearn** .
- **Boosted** deepVolt dlia model by more than **25% accuracy** through advanced traffic-based features using **python,scikit-learn** and **pandas** .

Data Scientist , Deepvolt - Tunis June 2024 - September 2024

- **Implemented** a comprehensive geographic analysis pipeline using **Python** and **GeoPandas** to identify optimal locations for electric vehicle charging stations, **covering 5,343.22 km²** and enabling data-driven infrastructure planning.
- **Developed and deployed** a real-time recommendation system **serving more than 5 clients** with an average **response time of 2.4 seconds**, utilizing **aws lambda** and **api gateway** with **Python**.

Projects

AI-Powered E-commerce Analytics Dashboard

[Github](#)

- **Engineered** an ETL pipeline to enrich mock e-commerce data by leveraging large language models (LLMs) for generating product reviews and categorizing items, resulting in a **92% reduction** in processing time and achieving a processing rate of **0.2 seconds per item** on consumer-grade hardware.
- **Optimized** LLM inference performance by **migrating from Ollama to Llama.cpp** and implementing a multi-stage Docker build, which **reduced image size by 80%** and **improved pipeline speed by 25%**.
- **Developed** a high-performance, concurrent client in **Go** using goroutines to handle parallel requests, **improving client-side throughput** and contributing to a **33% increase** in overall pipeline speed.
- **Designed and containerized** a multi-service application with Docker and Docker Compose to **ensure a portable and reproducible deployment environment**, incorporating health checks and service dependencies to guarantee a reliable system startup.

Smart Building Monitoring Project (Award-Winning)

[Github](#)

- Composed a sensor-based backend using **ESP32**, **FastAPI**, and **Docker** for real-time building monitoring with less than **1s** response. hosted on **Back4App**.
- Developed a **TensorFlow** model achieving **80%+ accuracy**; containerized with **Docker**, deployed to **Hugging Face**, and served via a custom **FastAPI** endpoint.
- Created an ETL pipeline using **Python**, **FastAPI**, and **Supabase** to convert sensor data into **BCF** files for a digital twin platform .

Skills

Programming : Python, SQL, Go
Cloud & MLOps: AWS, EC2, SageMaker, Docker, Git, Anaconda, Linux, Docker-Composer, FastApi
Tools: Databricks, Pandas, Sklearn, Seaborn, TensorFlow, Keras, NumPy, Polars, Spark, Kafka, Docling
Databases : PostgreSQL, MongoDB, Neo4j
Languages : English, French,Arabic

Certifications

Databricks fundamentals:	DATABRICKS ACADEMY	September 2025
Databricks Aws Platform Architect:	DATABRICKS ACADEMY	September 2025
Generative Ai with Diffusion model :	NVIDIA	November 2024

Leadership & Publications

- Project Manager CS Esprit SB**
- Lead Git and GitHub Fundamentals workshop
- Chair IAS ENIG SB 2022-2024**
- Led 5+ technical workshops on **AI**, **IoT**, and **digital twins**, engaging over than 30 students.
 - Won **6 out of 10** hackathons, including (iot spark,Dataquest and Eco gardien)
 - Secured sponsorship to attend the **IEEE IAS Annual Meeting in Nashville, USA**, representing ENIG on the international stage.
- Blogger On Medium**
- Authored and published multiple technical blogs on Medium to dissect complex project builds and share my approach.