

Meriem Melki Computer Science Engineering student

meriem.melki@insat.ucar.tn | +21696805806 | Tunisia | linkedin.com/in/meriem-melki | github.com/MelkiMeriem

SUMMARY

ML/NLP Scientist | AI Engineer

Master 2 Computer Science Engineering student (graduating Jan 2027) seeking a Ai / ML/NLP Scientist internship. Expert in identifying business trends, customer segmentation, and predictive modeling. Solid background in Natural Language Processing, Information Retrieval, and Large Language Models (LLMs), with hands-on experience building production ML systems. strong analytical mindset with ability to gather business requirements with attention to detail. Self-driven with excellent communication skills, comfortable translating applied research into working prototypes using Python, deep learning, and agentic frameworks.

PROFESSIONAL EXPERIENCE

AI intern, QCMed.tn

06/2026 – 08/2026 | Tunisia (Remote)

- Optimized a production RAG pipeline for a medical exam-prep chatbot delivering real social impact for students preparing for national exams, lifting answer faithfulness from 0.73 to 0.92 and answer correctness from 0.38 to 0.82. Analyzed user feedback and identified trends in query patterns to improve system performance across multiple exam subjects.
- Cut hallucinations with a cross-encoder re-ranker and query router, improving information retrieval quality and hardening the system against prompt-injection attacks, demonstrating attention to detail and strong collaboration within a cross-functional team.
- Boosted context recall from 0.68 to 0.96 .
- **Technologies:** Python, RAG, OpenAI, Qdrant, RAGAS, FastAPI , Prompt Engineering

Data Science Intern, Acrelec (KFC Data Project)

09/2025 – 10/2025 | France (Remote)

- Built ML models predicting KFC order waiting times across four channels, reaching ~2.3 min MAE (18% below baseline) on 50K+ orders, improving sales forecasting accuracy and identifying trends in customer ordering patterns.
- Technologies: Python, Pandas, Scikit-learn

Software Intern, Euro-Information (International Information Development)

07/2025 – 08/2025 | Tunisia

- Designed a SQL batch pipeline to automatically detect anomalies in employee badge records, replacing a manual review process and improving detection consistency.
- Technologies: SQL, DevBooster framework.

EDUCATION

Engineering student (Bac+5 / Master 2 equivalent level), INSAT

09/2022 – 01/2027

Relevant coursework: Machine Learning, Deep Learning, Data Analysis, Statistics, Object-Oriented Programming, Databases, Big Data, Business Intelligence , Web Frameworks , Cloud .

TECHNICAL SKILLS

AI & Machine Learning: Neural Networks | Deep Learning | PyTorch | TensorFlow | Supervised Learning | Unsupervised Learning | Model Evaluation | Computer Vision | Object Recognition | Perception | Decision-Making | Reinforcement Learning | Transfer Learning | Fine-tuning

Programming: Python | C++ | SQL | Java | JavaScript

Data Science: Customer Segmentation | Sales Forecasting | Trend Analysis | Data Cleaning & Preprocessing | Feature Engineering | Exploratory Data Analysis (EDA) | Applied Mathematics

Generative AI: LLMs | RAG | LangGraph | ChromaDB | Embeddings | APIs

Cloud & AI: AWS | Azure | Amazon Bedrock | SageMaker | Cloud Deployment

Web Frameworks: React | Angular | NestJS | FastAPI

Tools: Linux | Unix / Bash | Git | GitHub | Docker

Software Engineering: Prototyping | REST APIs | Version Control (Git/GitHub) | Agile

Big Data: Apache Spark | Hadoop | Distributed Computing | ETL

Databases: SQL | NoSQL

Vision: Computer Vision | Pose Estimation | Object Detection | YOLOv8 | Real-time Vision

LANGUAGES

Arabic | English | French | Spanish

CERTIFICATES

AWS Certified AI Practitioner

ACADEMIC RESEARCH & APPLIED PROJECTS

Respiratory Sound Classification

- Conducted academic research comparing PANNs+BiGRU, EfficientNet-Bo, and an Audio Spectrogram Transformer for 4-class respiratory sound classification on ICBHI 2017.
- Best AST config scored 0.6835, surpassing the published paper baseline (0.6810), balanced 0.68 sensitivity/specificity
- Technologies: Deep Learning, Audio Processing, Transformers, Python

LangGraph Blog Agent

- LangGraph DAG agent turning a topic, YouTube URL, or uploaded audio/video into a translated Markdown blog post, with Whisper transcription
- Built an end-to-end system with a FastAPI backend and React front-end, taking the prototype to a polished, functional application.
- Technologies: LangGraph DAG, Groq LLM

AI-Powered EDR Agent

- Agentic system combining Sysmon event collection with an ML layer to classify alerts, reduce false positives, and trigger automated responses.
- Achieved 84.6% detection recall (11/13 MITRE ATT&CK techniques) with ~25s mean time-to-detect and sub-second automated response.
- Technologies: LangGraph, ML, Python, Sysmon, SIGMA, ELK, Automation

Pose Estimation & Behavior Classification System

- Fine-tuned a COCO-pretrained YOLOv8s-pose model via transfer learning (96.9% mAP@50) and benchmarked Random Forest, Gradient Boosting, and LSTM classifiers for 5-class worm behavior recognition (best: 0.686 accuracy, 0.632 macro-F1) •
- Technologies: Deep Learning, Computer Vision, Transfer Learning, PyTorch, YOLOv8, scikit-learn, Python

CDPH Spark Analytics — Big Data Pipeline

- Distributed ETL pipeline in Apache Spark (SQL + Streaming) and Java processing San Francisco complaint data, with multi-line CSV parsing, batch/streaming analysis, and results persisted to HDFS.
- **Technologies:** Apache Spark, Spark SQL, Java, HDFS/Hadoop, Maven