HOUSSEM EDDINE

AI Engineer · Medical Imaging · Agentic Systems · Computer Vision

Ariana, Tunisia
+216 46 210 801
houssemeddineds@gmail.com
linkedin.com/in/houssemeddineds
github.com/HoussemDs
https://houssemds.github.io



PERSONAL INFORMATION

Nationality: Tunisian

Languages: English (C1) | French (B2) | Arabic (C1) | German (B1)

Location: Ariana, Tunisia – open to relocation

PROFESSIONAL SUMMARY

Medical AI Engineer specializing in clinical imaging pipelines, LLM fine-tuning for healthcare, and production computer vision systems deployed in medical and industrial settings. Built and validated 3D cardiac segmentation under EU MDR/IVDR constraints alongside senior cardiologists at Alight Consulting Germany (Dice 0.88–0.91, 95%+ clinical acceptance rate). Currently available for remote roles; relocating to Europe from December 2026 for a 6-month placement, with full on-site availability post-graduation in 2027. Sponsorship required for on-site positions.

TECH STACK

Medical Imaging: MONAI 3D, 3D U-Net, DICOM (pydicom), SimpleITK, OHIF Viewer, 3D Slicer, ONNX, GradCAM/XAI

Deep Learning: PyTorch, TensorFlow, LLM Fine-tuning (SFT + DPO / RLHF), Transfer Learning, YOLO

Agentic AI: LangChain, LangGraph, CrewAI, MCP Server, Agent-to-Agent (A2A)

Deployment & MLOps: Docker, ONNX Runtime, REST APIs, Git, Linux

Metrics & Standards: Dice Score, HD95, IoU · EU MDR/IVDR documentation · GDPR-compliant data handling

Core: Python, SQL, NumPy, Pandas, Scikit-Learn, Matplotlib, Statistical Analysis

WORK EXPERIENCE

01/2026 – 06/2026

AI Engineer – Computer Vision (Contractor)

SBT Tunisia – Tunis, Tunisia – On-site

Project-based contract: deliver a production-ready visual QC system for automotive wiring harness manufacturing.

Tech stack: YOLOv8, Python, OpenCV, A2A (Agent-to-Agent protocol), MCP Server (Model Context Protocol)

- Engineered a real-time YOLO-based quality control system for automotive wiring harnesses, reducing client return rate from **6.8% to 0.2%** across exports to French and Italian clients, as tracked by post-shipment return data over the 30-day post-deployment period
- Deployed a mobile computer-vision pipeline validating cable routing, colour sequencing, and configuration compliance directly on the assembly line.
- Integrating Agent-to-Agent (A2A) communication and an MCP server for automated production reporting and database connectivity; architecture directly transferable to medical device QC and **ISO 13485** audit trail requirements (ongoing through Jun 2026).

06/2025 – 09/2025

Medical AI Engineer (Intern)

Alight Consulting – Germany – Remote ***Project:*** *cardiac-ai-pipeline*

Internship focused on clinical-grade cardiac MRI/CT segmentation and LLM-assisted reporting within an EU MDR/IVDR regulatory framework, working directly alongside clinical stakeholders and targeting CE-mark readiness.

Tech stack: MONAI 3D, PyTorch, LangChain, SFT + DPO, DICOM, SimpleITK, ONNX, 3D Slicer, OHIF Viewer

- Built a production-grade multi-class 3D cardiac segmentation pipeline (LV, RV, myocardium) on the ACDC dataset, achieving **Dice 0.88–0.91** and **HD95 3.0–5.0 mm,** an **18 % Dice improvement** over the 3D U-Net baseline (0.70).
- Resolved multi-scanner class imbalance via RandCropByPosNegLabel and isotropic voxel-spacing normalisation; exported to ONNX and validated outputs in OHIF Viewer and 3D Slicer for radiologist sign-off.
- Conducted a blinded preference study with 2 senior cardiologists across **40 cases**, model outputs rated clinically acceptable in **95 %+ of cases for all three cardiac structures (LV, RV, myocardium),** outperforming the U-Net baseline on visual quality and boundary delineation.
- Produced **EU MDR/IVDR-aligned** technical documentation including **GradCAM** explainability maps for regulatory traceability.
- Fine-tuned a domain-specific clinical LLM (SFT + DPO) on **200+ clinician-annotated cardiac reports**, reducing hallucination rate by **35 %** (evaluated on a 50-report held-out set scored by two independent cardiologists) and cutting per-study reporting time by **30 %** (from **~12 min to ~8.5 min** average) via an integrated LangChain agentic pipeline.

01/2025 – 06/2025

Geospatial AI Engineer (Contractor)

CMA CGM – France – Remote ***LinkedIn:*** *LinkedIn post*

Project-based contract: AI-driven satellite segmentation for environmental planning and reforestation site selection.

Tech stack: U-Net, PyTorch, CrewAI, Multimodal Satellite Imagery, Python, Docker

- Optimised U-Net on multimodal satellite imagery, achieving **92.4 % IoU** (8 % above baseline) and eliminating 70% of manual labelling workload, validated on 150+ reference tiles.
- Fine-tuned the model for desert terrain, achieving a **31 % improvement** in generalisation to unseen geographies.
- Architected a CrewAI multi-agent system automating analysis and reporting workflows end-to-end, reducing weekly analysis and reporting from **~10 hrs to ~20 min validation,** as reported by the CMA CGM operations team

06/2023 – 09/2023

Data Scientist (preceded by 1-month internship, continued as contractor)

Newrest – Tunis-Carthage International Airport – Tunis, Tunisia – On-site

Newrest operates catering and supply-chain logistics for Tunisia's largest international airport. Began with a 1-month internship, then continued through the summer as a contracted Data Scientist.

Tech stack: Python, Scikit-Learn, Prophet, Pandas, SQL, Matplotlib

- Built and deployed a **3-day** supply forecast model covering the full airport catering operation, predicting daily stock requirements across all product lines for Tunisia's busiest airport.
- Reduced food waste by **55 %** and eliminated stock shortages within **3 months** by replacing manual inventory estimates with ML-driven forecasts.
- Presented findings directly to operations management and iterated the model based on real-world feedback from logistics staff.

05/2023 – Present

Founder & AI Engineer – Retail Intelligence

Self-employed – Tunisia

Designed and deployed production AI systems on a real business I own, forecasting, computer vision, and LLM reasoning applied to live transactions, real customers, and live inventory since 2023.

Tech stack: Python, PyTorch, Prophet, Scikit-Learn, LangChain, VLM, OpenCV, Facebook Graph API

- Built a demand forecasting pipeline on 90-day rolling sales data and Facebook trend signals, improving forecast accuracy from 52% to **88% MAPE** and reducing end-of-season unsold inventory by **42%** through improved purchasing decisions across collections.
- Developed a VLM-based pipeline to analyse customer preferences and product performance, directly informing store layout and restocking strategy, improving sales conversion rate by **+20%** following layout optimisation.
- Implemented Tunisia-specific social media scraping to detect emerging fashion trends **10–12 days ahead of peak demand**, giving a measurable sourcing lead-time advantage and reducing stockout incidents from **~8 to ~2 per month**.
- Achieved sufficient business performance to repay a **24-month** bank loan in full within **6 months** of deployment, with documented financial records confirming repayment

EDUCATION

06/2022 – 06/2027
(Currently Year 4 of 5)

Diplôme d'Ingénieur in Computer Science – AI Specialisation

ESPRIT University, Tunis, Tunisia

GPA: 15/20 (French scale) ≈ 3.7/4.0 GPA | Honours standing

RESEARCH & LEADERSHIP

09/2025 – 12/2025

Research : Integrating Vision-Language Models and Multi-Agent LLM Systems for Clinical Reasoning in Brain Tumor Analysis

ESPRIT University × Research Department – Tunisia ***Project:*** *NeuroAI-Pipeline*

Academic research sprint delivered to an industry partner in 3 months. ***Preprint:*** *OSF Preprint*

Tech stack: PyTorch, MONAI, MedViT, U-Net/ResNet, FAISS, MongoDB, GradCAM/XAI, LangGraph, Gemini, Whisper, RAG

- Designed a multimodal brain MRI analysis pipeline combining MedViT embeddings, FAISS-based case retrieval, and a U-Net/ResNet segmentation agent with GradCAM explainability, evaluated on the BraTS/Medical Segmentation Decathlon dataset, achieving 3.2× processing throughput over a single-modal baseline and sub-300 ms inference latency.
- Orchestrated a three-LLM agentic architecture (Gemini for image reasoning, LLaMA 8B for XAI interpretation, LLaMA 70B for report structuring) producing structured JSON and exportable PDF clinical reports with explainable segmentation.
- Built a virtual consultation system accepting text, audio (Whisper), and video input, routing through a five-agent RAG pipeline for symptom normalization, scoring, risk estimation, and report generation, delivering a complete case assessment in under 2 minutes end-to-end.

Nov 2025

1st Prize, Hackathon Talents Tunisie – FMIA 2025

Cité de la Culture, Tunis

1st place, EdTech category (5,000 TND prize) ***LinkedIn:*** *LinkedIn post*

- Built **Eduverse AI**: platform converting any lesson (PDF, video, web) into interactive RPG-style learning games via AI — no coding required.

06/2023 – present

Tech Manager – AI & Healthcare Education

DeepFlow – Tunisia (Multi-University)

Community-led AI education initiative operating across multiple universities in Tunisia.

- Designed and delivered 20+ AI/ML workshops on medical imaging and healthcare AI to 200+ students across ESPRIT, INSAT, and SUP'COM universities — covering topics from MRI analysis to agentic diagnostic systems.
- Built and maintained curriculum covering deep learning, computer vision, and clinical AI deployment; coordinated speakers and facilitated hands-on sessions over 3 years. Workshop materials and session recaps available on LinkedIn.