Adil Eddarif

Artificial Intelligence Engineer

Kenitra, Morocco | (+212) 605695329 | adil.eddarif@gmail.com
github.com/AdilEddarif adileddarif.com

Professional Profile

Artificial Intelligence Engineer with in-depth expertise in Data Science, Large Language Models (LLMs), machine learning, MLOps, computer vision, and deployment of AI solutions in the cloud. Proven experience in developing ML pipelines, optimizing models, and putting them into production.

Education

Master's Degree in Artificial Intelligence and Virtual Reality, Ibn Tofail University
Sep. 2023 – Jul. 2025

Bachelor's Degree in Mathematics and Computer Science, Ibn Tofail University
Sep. 2019 – Jul. 2023

Baccalaureate in Mathematical Sciences B, Alpha Educational Group High School
Sep. 2018 – Jul. 2019

Professional Experience

Artificial Intelligence Engineer, Cheikh Zaïd Foundation
May 2026 – Present

- Built an end-to-end medical-orientation **RAG assistant** for Cheikh Zaid Hospital in Rabat, using a two-stage hybrid retrieval architecture combining **FAISS**, **BM25**, and **cross-encoder reranking**, deployed through a **FastAPI** service with rule-based symptom routing and emergency red-flag detection.
- Designed and executed a rigorous evaluation framework based on a **1,000-question golden dataset** covering five clinical-risk categories, an independent **LLM judge** for faithfulness and hallucination assessment, and an automated CI-ready non-regression gate.
- Identified and corrected a critical safety-routing bug through systematic evaluation, increasing emergency recall from **50% to 100%**.
- Hardened the assistant for production safety by implementing prompt-injection filtering before every LLM call, deduplication and version-safe index rebuilding, and a strict advisory-only workflow that avoids diagnosis and systematically defers to physicians or emergency services.
- Validated the system's safety and reliability through a dedicated automated test suite developed with **pytest**.
- Built an end-to-end two-stage deep learning pipeline for chest X-ray triage, combining a **ResNet50** binary classifier for normal-versus-suspicious-opacity classification with a **YOLOv8** detector for bounding-box localization of radiological findings.
- Achieved **96.9% accuracy**, **97.9% F1-score**, **99.7% AUC**, and **98.2% sensitivity** on a held-out test set of 522 chest X-ray images.
- Designed a multi-step clinical workflow covering patient intake, AI-assisted image analysis, results review, and medical-report generation, supported by a doctor-in-the-loop feedback mechanism.
- Persisted physician corrections and validation decisions to structured CSV files for model monitoring and future retraining.

Freelance AI Engineer – Satellite Computer Vision, GeoAP
Dec. 2025 – May 2026 | Remote, Morocco

- **End-to-End Satellite CV Pipeline:** Architected and delivered a production-grade MLOps pipeline for detection and semantic segmentation of agricultural and urban objects from satellite imagery, processing 180 source images into 12,000+ training tiles via overlap-aware tiling with zero data leakage across train/val/test splits.
- **Multi-Model Architecture:** Trained three complementary models — **YOLO11m** object detection (mAP50 0.836), **YOLO11m** instance segmentation (mAP50-mask 0.737), and fine-tuned **SAM1 ViT-H** semantic segmentation (mIoU 0.804) — covering 10 object classes including buildings, trees, solar panels, swimming pools, cultivated parcels, rivers, and roads.
- **SAM1 Fine-Tuning:** Fine-tuned SAM1 ViT-H mask decoder (4M trainable params, 300M-param encoder frozen) with pre-computed embedding caching across 9,958 tiles, reducing epoch time from 88 min to 5 min on a T4 GPU. Implemented custom Dice + Focal loss and **torch.amp** mixed-precision training.
- **GIS-Ready Output:** Inference results exported as GeoJSON FeatureCollections with real-world WGS84 coordinates, enabling detected objects and segmentation masks to be visualized directly on real maps in Google Earth, QGIS, Leaflet, and any GIS-compatible tool.
- **SAHI Tiled Inference:** Integrated Slicing Aided Hyper Inference (**SAHI**) for detection on high-resolution satellite imagery, enabling robust small-object detection across arbitrary image resolutions.
- **FastAPI Inference Service:** Built a **FastAPI** REST API with four inference endpoints, lazy singleton model loading, and an interactive dark-theme frontend with live legend and GeoJSON download.
- **MLOps & CI/CD:** Configured a full **DVC** pipeline DAG (12 reproducible stages) with **Roboflow** dataset versioning. Implemented **GitHub Actions** CI/CD with linting, 25 unit tests, API smoke tests, and Docker build validation on every push.
- **Tech Stack:** PyTorch, YOLO11m, SAM1 ViT-H, SAHI, FastAPI, rasterio, Shapely, OpenCV, DVC, Roboflow, Lightning AI (T4 GPU), Docker, GitHub Actions.

Data Scientist, Attijariwafa Bank
Feb. 2025 – Mar. 2026

- Analysis of large financial datasets from bank branches.
- Construction of supervised predictive models to estimate annual business objectives.
- Deployment of a complete MLOps pipeline with MLflow: tracking, validation (MAE, RMSE, R^2), and model export (.pkl).
- Replacement of heuristic rules with predictions based on historical data (Y-2 to Y).
- Achieved an R^2 score of 0.9, indicating strong predictive accuracy.
- Design and deployment of an internal chatbot to extract information using semantic search.

AI / Computer Vision Intern, Arteka Company
Aug. 2024 – Sep. 2024

- Development of image recognition algorithms for medical imaging using OpenCV and GANs.
- Generation of synthetic cancer cell datasets to improve model robustness.
- Contribution to the design of a proof of concept for an AI-assisted diagnostic system.

Academic Projects

Optimized AI Research Assistant (RAG & TTS)

- Development of a local RAG (Retrieval-Augmented Generation) system using **Llama-3** and **Marker** for deep analysis of research papers.
- Implementation of a real-time audio narration pipeline via **Kokoro TTS**, with asynchronous pre-generation to eliminate GPU latency on **NVIDIA RTX 3050**.
- Automated and structured extraction of complex bibliographies (40+ references) through layout segmentation and semantic filtering.
- Optimization of VRAM management through a "Selective Decoding" system and chunking for the analysis of dense documents.

Masked Multimodal Autoencoders for Medical Applications (ViT + BERT)

- Integration of a Vision Transformer (ViT) for image processing and BERT for medical text processing.
- Pre-training on the ROCO dataset (81,000 annotated medical images) for multimodal learning.

AI Chatbot with Retrieval-Augmented Generation (RAG)

- Design of an intelligent chatbot integrating document synthesis, question-answering, and semantic search.
- Technologies used: Flask (REST API), React (user interface), HuggingFace, ChromaDB, LangChain.

OCR for the Tifinagh Alphabet

- Development of a CNN-based OCR model for recognizing handwritten Tifinagh characters.
- Application of data augmentation and preprocessing techniques; achieved detection accuracy of 93%.

Autonomous Drone Delivery System

- Design of a 3D simulation in Gazebo for a drone delivery system, enabling the delivery of items on command.
- Improvement of simulation efficiency through the implementation of modular communication protocols between nodes.

Technical Skills

Languages: Python, C++, Java, C#, SQL, HTML, PHP
Vision & AI: PyTorch, TensorFlow, Scikit-learn, YOLO11, SAM (Segment Anything), SAHI, OpenCV, rasterio, Shapely, Pandas
LLM & NLP: HuggingFace Transformers, LangChain, RAG, BERT, PyMuPDF, FAISS, BM25
MLOps & Cloud: Docker, DVC, MLflow, Roboflow, GitHub Actions, Git, AWS, ChromaDB
Tools & Backend: FastAPI, Flask, React, Linux, Jupyter, NoSQL, PostgreSQL

Languages

French: fluent | English: professional | Arabic: native | German: beginner