

Amr Ahmed

Cairo, Egypt

+20 1552366372 | amrahmedgohar205@gmail.com | GitHub: Amr2054 | LinkedIn: amrahmed205

Engineering Profile

AI Engineer experienced in building AI systems, including research-grade deep learning models and production-grade RAG pipelines with FastAPI, Docker and vector databases. Proficient in Python and PyTorch, with applied work spanning computer vision, LLM/agent pipelines, and sports analytics. Additional depth in medical imaging and clinical data from a biomedical engineering background.

Technical Skills

- **Languages:** Python (primary), C, SQL, Java
- **Deep Learning:** PyTorch, CNNs, RNNs/LSTMs, Transformers, Vision Transformers, Diffusion Models, LoRA/PEFT, Contrastive & Self-Supervised Learning
- **Computer Vision & ML:** OpenCV, Scikit learn, image classification & segmentation, feature engineering, model evaluation & interpretability
- **LLMs, RAG & Agents:** LangChain, CrewAI, OpenAI/Anthropic APIs, Ollama, Qdrant, FAISS, pgvector, prompt engineering, multi-agent orchestration
- **Backend & Infrastructure:** FastAPI, Node.js, PostgreSQL, REST APIs, Docker, Git/GitHub, Linux

Experience

FlyRank AI

Machine Learning Intern

Remote

Jun 2026 – Aug 2026

- Built reproducible ML workflows in Python and scikit-learn over large-scale search intelligence datasets, covering feature engineering, model training, and evaluation.
- Analyzed search performance signals to produce data-driven content optimization recommendations delivered to the product team.

Digital Egypt Pioneers Initiative

Data Science Trainee

Hybrid

Jul 2025 – Dec 2025

- Completed an intensive applied data science program (Python, SQL, statistics, ML), delivering end-to-end predictive modeling projects on real-world datasets from data cleaning and feature engineering through model selection and evaluation.

Education

Cairo University, Faculty of Engineering

B.Sc. in Systems and Biomedical Engineering

Giza, Egypt

2022 – 2027

- **Focus:** Deep Learning, Artificial Intelligence, Computer Vision
- **GPA:** 3.88/4.00

Projects & Research

Hierarchical Group Activity Recognition in Volleyball – PyTorch, ResNet50, LSTM

- Reimplemented the CVPR 2016 published paper in modern PyTorch, replacing the original AlexNet/Caffe stack with a ResNet50 + two-stage Person LSTM & Group BiLSTM pipeline to model both individual player actions and team-level dynamics from volleyball footage.
- Surpassed the original paper's benchmarks across all 8 baselines, achieving a +9.95% improvement over the published CVPR result.

LoRA-Based Domain Adaptation of UNETR++ Vision Transformer – PyTorch, Transformers, LoRA

- Engineered a custom Low-Rank Adaptation (LoRA) framework from scratch to transfer the deep spatial features of UNETR++ (a 42M-parameter Vision Transformer) to a data-scarce 3D domain.
- Reduced trainable parameters by 99.5% (to just 135k) while achieving an elite 0.94 Mean Dice score—nearly matching the 0.95 full fine-tuning baseline—for precise volumetric segmentation, bypassing massive computational and GPU memory bottlenecks.

DocuQuery: Multi-LLM RAG Backend for Bilingual Document Intelligence – FastAPI, Qdrant, Postgres, LangChain, Docker

- Architected an asynchronous Retrieval-Augmented Generation (RAG) backend supporting bilingual (Arabic/English) document ingestion, semantic search, and grounded question answering.
- Integrated multi-provider vector database support (Qdrant, pgvector) and multi-provider LLM support (OpenAI, Cohere, Ollama) to retrieve relevant text and generate grounded, context-aware answers.

CancerScope: Explainable ML Pipeline for Metastatic Tissue Detection – Python, OpenCV, Scikit-Learn

- Implemented an automated ML pipeline bypassing deep learning black-boxes by extracting morphological, color, and textural features from histopathologic scans to identify metastatic tissue.
- Implemented computer vision algorithms (GLCM, LBP, SFTA) to train a Random Forest model, achieving a 0.9542 ROC AUC.

Football Tactical Retrieval Engine – Research Collaboration, PyTorch, Transformers, Contrastive Learning

Ongoing

- Designing a hierarchical Transformer to map multi-agent tracking data from the FIFA World Cup 2022 into tactical embeddings for similarity retrieval over plays.
- Training with supervised contrastive learning on weakly-derived labels; ongoing research collaboration with teaching assistants at a German university.

Certifications & Coursework

Diffusion Models Course – Hugging Face

2026

Pytorch for Deep Learning – DeepLearning.AI

2026

Machine Learning Specialist – CSkilled Academy

2025