Ahmed Sta

Agentic AI Engineer | Multi-Agent Systems, Tool Calling & LLMs | France

France (mobile) 07 59 56 17 78 sta.ahmed09@gmail.com LinkedIn GitHub Portfolio

Professional Profile

Fresh MSc in AI graduate with **more than one year of experience in production AI and agentic projects**, building tools used daily by doctors and science teams. I design **multi-agent architectures** (LangGraph): tool calling, routing between agents, state management and RAG. I run **experiments to evaluate agent behavior** (LLM-as-a-judge, benchmarks, error analysis) and make agents more **robust** with guardrails, grounding checks and structured outputs. Solid **software engineering** habits: Python, tests, CI/CD, code review, Git, Docker. Strong ML / deep learning base (PyTorch, fine-tuning, VLMs).
Seeking a full-time role as Agentic AI Software Engineer in France – Available immediately.

Technical Skills

Agentic AI	LangGraph, Multi-Agent Systems, Tool Calling, Function Calling, Routing, State Management, MCP
LLMs	OpenAI, Claude, Gemini, Gemma / Ollama (local), Prompt Engineering, LoRA / PEFT
RAG & Retrieval	RAG, Agentic RAG, Embeddings, Reranking, Vector DBs, LlamaIndex, Elasticsearch
Eval & Robustness	LLM-as-a-Judge, Agent Evaluation, Benchmarks, Error Analysis, Guardrails, Structured Outputs
ML & DL	PyTorch, TensorFlow, Scikit-learn, XGBoost, Hugging Face, VLMs, OCR, Speech-to-Text
Software Eng.	Python, Pytest, Code Review, Git, Linux, Docker, CI/CD (GitHub Actions), FastAPI, SQL
Cloud & MLOps	AWS, Azure, GCP, MLflow, Airflow, Langfuse

Education

MSc in Artificial Intelligence – *Paris School of Technology & Business*, Paris **2025 – 2026**
Engineering Degree in Computer Science (Data Science) – *ENSI*, Tunis **2021 – 2024**

Professional Experience

Agentic AI Engineer | Cure51 – *France* – *Internship* **Mar 2026 – Sep 2026**

- **Designed and built a full multi-agent system with LangGraph, end to end**: orchestration (nodes, routing between agents, shared state), the agents and their tools, guardrails and evaluation.
- **Internal agent**: SQL agent and retrieval / Q&A over reports, with query validation and structured outputs.
- **External research agent**: answers scientific questions by calling research-paper and clinical-trial APIs and scientific websites (tool calling).
- **Ran experiments to evaluate agent behavior**: automated LLM-as-a-judge evaluation comparing models and prompts on answer quality, latency and cost; continuous improvement with new nodes and use cases.
- **Owned a production LLM extraction pipeline on AWS** (OCR, translation, VLMs) turning multilingual medical reports into structured entities; built evaluation sets from real documents to benchmark and choose models.
- **Technologies:** Python, LangGraph, LangChain, LLM APIs, Tool Calling, Structured Outputs, LLM-as-a-Judge, AWS, SQL, OCR, VLMs

AI Engineer – Healthcare | WIC Doctor – *Tunisia* – *CDI* **Nov 2024 – Feb 2026**

- **Multi-agent assistant for doctors**: a router delegates each question to specialized medical agents that use tools to search patient records and medical knowledge (Agentic RAG), with grounded, cited answers.
- **Patient chatbot agent with tools** (function calling): books appointments, collects patient information and calls backend services on its own.
- **VLM medical image analysis pipeline**: retrieval of relevant medical context to build the prompt, then VLM analysis of the image, with clinical findings as structured output.
- **Real-time medical speech pipeline**: Speech-to-Text, medical term correction, then LLM-generated structured reports, with grounding checks and confidence scores against hallucinations.
- **Shipped to production**: tools used daily by doctors, improved release after release.
- **Technologies:** Python, LLM APIs, Multi-Agent Systems, Agentic RAG, Tool Calling, Embeddings, VLMs, Whisper

Data Scientist | PRIME Lab – *Canada* – *Internship* **Feb 2024 – Aug 2024**

- **Benchmarked many ML models** (ensembles, XGBoost) for breach risk prediction, on precision, recall and F1.
- Full data science pipeline with Pandas (cleaning, feature engineering), exploratory and error analysis.

Data Scientist | Fysali SAS – *France* – *Internship* **Jun 2023 – Aug 2023**

- Audio violence detection pipeline: Speech-to-Text, then **fine-tuned Transformer classification** (Hugging Face), evaluated with precision, recall and F1.

Projects & Competitions

Talentia – Agentic recruiting platform (live): multi-step LLM pipeline and a voice agent that runs candidate pre-interviews. Demo – GitHub

Unifi PDF Lifting (GenAI, Zindi 6/73): RAG pipeline extracting ESG metrics from multi-page PDFs (chunking, embeddings, LLM extraction).

Data science competitions (Zindi): 3 gold medals; Landslide Prevention 2/172, IndabaX Tunisia 5/115.