

Edouard Fossier

Applied AI Engineer · Full-Stack AI Systems · Prototyping

Paris, France · Open to relocate · Available immediately

edouardfossier@gmail.com · +33 6 79 54 65 33

github.com/edouardfossier · huggingface.co/edouardfossier · linkedin.com/in/edouardf · edouardfossier.com

PROFILE

Applied AI Engineer building full-stack AI systems from ambiguous product needs to working software. Production experience with RAG, agentic systems, voice AI, evaluation frameworks. Combines hands-on engineering, product/business judgment, and clear stakeholder communication.

2026 HIGHLIGHTS

- **1st place** · Voodoo × Anthropic Hackathon (Track 3, Market Intelligence) · [VoodRadar](#) - mobile-game ad pipeline
- **1st place** · GOSIM Agentic Hackathon (Z.AI Innovation) · [Xiexie](#) - voice-first macOS scam protection · solo
- **3rd place** · GPU Mode Paris (LLM pre-training) · [Custom GPT](#) on 32× NVIDIA B300 · 1 day
- **Built** · [Le Sphinx](#) - voice Akinator on Mistral Large + Small (Mistral hackathon, MCP server, 220 tests, prod)

EXPERIENCE

AI Engineer, Freelance · Mar 2026 - Present

Shipping applied AI products and prototypes for clients and real users.

- Built [Madeleine Agent](#) (LangGraph voice agent, 9 tools, OpenAI Realtime) for AI Sisters; business workflow apps (kombucha steering, beverage-shop inventory, real-estate search); [Les Reufs](#) WhatsApp voice-note product.
- Three hackathon podiums in 2026 (1st Voodoo & Anthropic, 1st GOSIM solo, 3rd GPU Mode). Multimodal pipelines on Mistral, Anthropic, OpenAI, Z.AI; FastAPI/Pydantic + React/Next.js full-stack delivery.
- Robotics: SO-101 pick-and-place datasets and 6+ LeRobot policy fine-tunes (ACT, SmoIVLA, Pi0, Diffusion, VQ-BET) on Hugging Face; Reachy Mini Wireless.

AI Engineer, beta.gouv / French State · Sep 2024 - Feb 2026

Technical lead for a sovereign RAG assistant for HR teams (~70 beta users).

- Designed multi-stage architecture: intent gating, retrieval over heterogeneous legal/admin docs, model-based document selector, cited generation. Built ingestion pipelines and evaluation framework.
- Built a 1,490-question goldset and ran 8,300+ evaluations across 10+ architecture variants (RAGAS, LLM-as-judge, user feedback). Diagnosed retrieval as primary failure mode.
- Improved global score 0.60 → 0.83 across three iterations. Stack: Python, FastAPI, Pydantic, PostgreSQL+pgvector, RAGAS, MLflow.

AI Engineer, The Artificial Business · Nov 2024 - Mar 2025

Contributed to a self-hosted AI agent platform under sovereignty/GDPR constraints.

General Manager, Le Wagon Paris · Mar 2016 - Aug 2022

Ran the largest Le Wagon campus. Scaled it from ~120 to 450 students/year and grew revenue from ~€540K to €3.1M ARR. Managed teams, teaching operations and ecosystem relationships.

Co-Founder & CEO, Vili & Ve · Sep 2014 - Dec 2015

Co-founded a UX/UI and web/mobile digital studio. Closed enterprise clients including CAC40 accounts; ran discovery, pre-sales and delivery.

EDUCATION

Le Wagon · Data Science & Data Engineering Certificate (2024)

ESCP Business School · Master in Management & Entrepreneurship (2013)

TECHNICAL SKILLS

- **AI/ML**: LLMs, RAG, agentic systems, MCP, voice agents, evaluation (RAGAS, LLM-as-judge), embeddings, reranking, fine-tuning, distributed pre-training (DDP, Muon, WSD).
- **Backend**: Python, FastAPI, Pydantic, SQL, ingestion pipelines, PostgreSQL + pgvector, Qdrant, Redis.
- **Frontend / Product**: TypeScript, React, Next.js, full-stack product building, user interviews, stakeholder alignment.
- **Infra**: Docker, Caddy, GCP, Cloudflare Workers, SLURM, multi-GPU clusters, CI/CD.
- **Models / Tools**: Mistral (Large + Small, in prod), Anthropic, OpenAI (incl. Realtime), Gemini, Z.AI GLM, Hugging Face, LeRobot, LangChain/LangGraph, Chainlit, MLflow, ElevenLabs, AssemblyAI.