Adrià Flores Albó

adriaflores.com | adriaflores | adriaflores | adriafloresalbo@gmail.com | +34 684 277 913 |

Research-oriented AI engineer working across LLMs, Computer Vision, and production AI systems, with experience building and pretraining language models from scratch, fine-tuning, OCR/VLMs, document AI and evaluation.

Education

Universitat Politècnica de Catalunya (UPC). BSc in Artificial Intelligence — Barcelona 2022 - Jun 2026

- **Academic performance: 8.94**/10 (Year 3) | **9.65**/10 (final semester)
- **Relevant Coursework:** Large Language Models (**Honors**), Advanced Reinforcement Learning, Unsupervised Learning, Neural Networks & Deep Learning, Computer Vision, Voice & Dialogue, Advanced Databases

Experience

Medidedalia — AI Research Engineer Intern (30h/week) *Barcelona (Hybrid)* Oct 2025 – Jun 2026

- Completed BSc thesis on SchemaRouter, a budget-aware router for schema-conditioned document extraction that selects the most suitable RAG-based strategy under quality/cost/latency constraints.
- Built and benchmarked a privacy-preserving anonymization service for clinical and legal documents, combining deterministic rules, fine-tuned NER models, and LLM-based support for ambiguous cases.
- Researched, evaluated, and integrated a production-ready OCR service for sensitive medical-legal PDFs, benchmarking specialized OCR/VLM models and enabling spatial text localization for source citation.

Projects

Kimi K3 from Scratch [] [🔗]

Independent Research — Jul–Oct 2026

Methods & Technologies: *LLM Pretraining, Scaling Experiments, Hybrid Attention, Sparse MoE, GPU Optimization, PyTorch/CUDA*

- Reimplemented the Kimi K3 language architecture from scratch at 295M parameters, integrating hybrid attention (KDA + Gated MLA), Attention Residuals, Stable LatentMoE.
- Conducted controlled learning-rate and batch-size experiments across three model scales, optimized the training pipeline for GPU throughput and checkpoint recovery, and completed pretraining on 6B FineWeb-Edu tokens.

Uncensoring LLMs with Abliteration []

UPC — Research Project — May 2026 ***Tools:*** *PyTorch, Transformers, Hugging Face*

- Implemented abliteration, a training-free model steering technique that identifies a refusal direction in the residual stream and removes it through weight orthogonalization.
- Evaluated the abliterated Llama-3-8B model with automated refusal and capability benchmarks, reducing refusal behavior while preserving general performance.

RAG System Optimization with Graph Neural Networks []

UPC & Nuclia — Feb 2025 – May 2025 ***Tools:*** *PyTorch, GNNs, KGs, LangChain, Sentence-Transformers, NetworkX, Streamlit*

- Drove a **7x** increase in retrieval accuracy (6.2% to 42.3%) for a commercial RAG system by replacing an inefficient KG traversal with a novel GNN strategy.
- Researched and benchmarked three retrieval models for optimal performance and scalability.

Explainable AI for Clinical Decision Support []

UPC — Mar 2025 – May 2025 ***Tools:*** *PyTorch, SHAP, SpaCy, Gemini, Scikit-learn, Pandas, Streamlit, ChromaDB*

- Architected a Seq2Seq Autoencoder to predict 30-day hospital readmission for 100k diabetes patients.
- Built an XAI agent integrating SHAP with a RAG-enhanced Gemini LLM, supported by a custom feature pipeline using SpaCy embeddings and TSNE to generate human-readable clinical justifications.

Technical Skills

AI/ML: PyTorch, Transformers, Hugging Face, PEFT/LoRA, Computer Vision, GNNs, Reinforcement Learning, World Models, Generative Models, Diffusion Models

LLM & Vision Systems: LLM evaluation, structured information extraction, mechanistic interpretability, model steering, quantization, RAG, LangChain, NER, OCR/VLMs, vLLM, Ollama

AI Infrastructure: Python, C/C++, SQL, FastAPI, Docker, Kubernetes, Apache Spark, Vector DBs, Graph DBs

Awards & Competitions

- **Winner** — More than a Hack, Barça Innovation Hub (2026): AI football-analysis chatbot integrating tracking data, inferred events, and live commentary audio. [🔗] *3,500+ applicants*
- **Winner** — Datathon FME (2025), Smadex Challenge: Predicted seven-day in-app purchase revenue using LightGBM and neural networks with learned embeddings. []
- **Winner** — BitsxLaMarató (2025), BSC Challenge: Benchmarked bioinformatics models to learn a continuous representation bridging two medulloblastoma patient groups. []

Languages: Spanish (Native), Catalan (Native), English (Fluent, C1)