

Adrián Hernández Padrón

ML Engineer | End-to-End ML, LLM and Agent Systems in Production

📍 Tenerife, Spain ✉ adrianhdezp10@gmail.com ☎ +34 611 45 14 44 🔗 adrianhernandezpadron 📱 Adrihp06

Summary

ML Engineer with four years of experience taking models through the full lifecycle: problem definition, modeling, deployment, and the monitoring that keeps them trustworthy once real users depend on them. My work spans classical ML and statistics, LLM and agent systems, and the MLOps that holds both together — versioning, promotion across environments and CI/CD. A physicist by training, I came to data science through environmental research, and I still care most about projects where the numbers change a decision. Spanish native, English at professional working proficiency.

Skills

MLOps and Platform: Databricks (Asset Bundles, Unity Catalog model registry and lineage), MLflow tracking and model promotion across dev, test and production, Docker, Kubernetes, ArgoCD, Airflow, CI/CD with Azure DevOps and GitHub Actions, model and data drift monitoring, Azure, Microsoft Fabric, Apache Spark, SQL

Machine Learning and Statistics: Python, pandas, NumPy, scikit-learn, XGBoost, PyTorch, Matplotlib, statistical modeling, forecasting, clustering, anomaly detection, optimization, embeddings

LLMs and Agents: Prompt and retrieval design, RAG with hybrid search and reranking, vector databases (Qdrant), tool use via MCP, multi-agent systems, evaluation and guardrails, OpenAI, Anthropic, Azure AI Foundry, Vertex AI

Engineering: FastAPI, REST APIs, Git, pytest, code review

Experience

Intelequia, Senior Data Scientist

Jan 2026 to present

- Lead Intelequia's Data Science team, owning the MLOps practice across client projects: model versioning in Unity Catalog, promotion through dev, test and production, and the monitoring that follows a model once it is live on Azure and Databricks.
- Build LLM agents and data agents on Azure AI Foundry, grounding them on each client's own data and gating every change behind an evaluation set before it reaches users — the same release discipline applied to non-deterministic systems.
- Built a material allocation tool combining statistics, classical ML and embeddings, served behind an API with a UI for the client team.

DISA Corporación Petrolífera S.A., Data Scientist

Sept 2023 to Jan 2026

- Moved ML from ad hoc runs into repeatable, versioned training and inference workflows with Docker, Kubernetes, ArgoCD and MLflow on Databricks, promoted across dev, test and production behind either manual approval or automatic CI/CD depending on the model, with drift monitoring and rollback once live.
- Worked on DISA's route solvers, adding cluster constraints to improve fleet efficiency and lower logistics costs across the energy distribution network.
- Delivered DISA's first GenAI use cases: sentiment analysis, a RAG assistant over internal documents and LLM workflows that replaced manual review steps.
- Led a joint project with the University of La Laguna on anomalies in authentication logs, comparing embeddings, sequential models and autoencoders.

Loro Parque Fundación, Research Data Scientist

Jan 2023 to Aug 2023

- Measured the environmental impact of anthropogenic underwater noise: built the pipeline turning large volumes of marine audio into research-ready data, then trained the sound classification and event detection models that made the analysis possible.

Turismo de Tenerife, Research Department Coordinator

July 2022 to Jan 2023

Maintained the tourism Data Warehouse and turned its data into statistical analyses and Power BI reports for the research team.

Selected Projects _____

Travelruter

A travel planner I have designed, built and run in production for close to a year, to find out how much value an agent adds when it has real trip data.

- Built the FastAPI backend with async SQLAlchemy 2.0 and PostgreSQL + PostGIS, structured as DDD bounded contexts — trip planning, routing, maps, identity and billing — with Alembic migrations and integrations to OpenRouteService, Google Maps and Amadeus for routing, geocoding and hotel search.
- Built a PydanticAI agent that streams responses over WebSocket and reaches each user’s trips through an MCP tool layer of 18+ typed travel tools — search destinations, manage POIs and accommodations, calculate routes and transit, generate and optimize daily schedules, track budgets — so it reads and edits live data instead of answering from context alone.
- Designed the agent around composable capabilities — thinking, web search, tool approval, history processing — over a multi-provider model layer that resolves Anthropic, OpenAI, Google and local Ollama models at runtime, with per-session message history and tool-result truncation keeping context and token cost bounded.
- Exposed the same trips through a remote MCP server with streamable-http transport so ChatGPT and Claude Desktop can act on real trip data, which forced the tool surface to stay small, typed and cheap to call.

Certifications _____

- [Microsoft Certified: Machine Learning Operations Engineer Associate \(AI-300\)](#)
- [Microsoft Certified: Fabric Analytics Engineer Associate \(DP-600\)](#)

Education _____

MSc	Universidad Internacional de Valencia (VIU) , Astronomy and Astrophysics	June 2025
MSc	Universidad Internacional de Valencia (VIU) , Data Science and Big Data	June 2023
BSc	Universidad de La Laguna (ULL) , Physics	Sept 2021