

CHRISTOPHER MANZANO VIMOS

Data Scientist & AI Engineer

Madrid, Spain · +34 682 201 537 · christopher.manzano.vimos@gmail.com

github.com/chriss1245 · linkedin.com/in/christopher-manzano-vimos

PROFILE

Data scientist and AI engineer with end-to-end ownership of production AI systems: architecture design, key technical decisions, deployment, and operation. Nearly four years shipping ML, NLP, and Generative AI solutions for clients in energy, insurance, SaaS, legal tech, and medical imaging (US DoD-funded research). Experienced deploying open-source LLMs on private infrastructure and building agentic pipelines with LangChain and LangGraph. Outside client work, I design, build, and operate my own AI products end to end at manapple.dev. I mentor other engineers, define quality standards, and work directly with business stakeholders to turn ambiguous problems into systems that deliver measurable value.

INDEPENDENT PROJECTS: MANAPPLE.DEV (LIVE IN PRODUCTION)

Warren: AI investment analyst (warren.manapple.dev)

- Multi-agent system built with LangGraph on Gemini models that combines quantitative market data, technical analysis, competitor research, and intrinsic-value estimation to generate buy/sell investment theses for portfolio positions.
- Full product ownership: architecture, deployment, public user registration, and SSO with a guest demo mode for read-only exploration.

Gollum: Auction tracker with automated alerting (gollum.manapple.dev)

- Playwright-based scraping pipeline that tracks auctions and triggers notifications through a self-hosted email delivery service I deployed and operate.

EXPERIENCE

AI & Machine Learning Engineer and AI Consultant, WhiteBoxML

Madrid · Sept 2022 – Present

End-to-end technical ownership of production AI systems for enterprise clients: architecture, delivery, and functional leadership. Coordinated and mentored other engineers, defining quality standards and design criteria for ML, NLP, and GenAI systems focused on scalability, maintainability, and business value.

Iberdrola: Renewable energy generation forecasting (AWS)

- Designed and built scalable ML pipelines with distributed processing to forecast energy output for 200+ renewable plants across the US, Spain, and the UK, ingesting over 2 TB of data daily.
- Developed ingestion and normalization of meteorological data from multiple providers (ECMWF, GFS, WRF).
- Engineered for reproducibility, traceability, and robustness in production environments.

Intermundial: LLM-based intelligent claims automation

- Architected NLP/GenAI pipelines with LangChain and LangGraph for hierarchical document classification (F1-evaluated) and information extraction in insurance claims, achieving over 90% validation accuracy on identity fields (DNI, NIE, passport).
- Built an LLM-as-judge evaluation framework for open-ended fields such as cancellation conditions, rating extractions as correct, partially correct, or incorrect across heterogeneous providers.
- Designed a per-claim document-requirements matrix mapping each claim type to its mandatory documentation, and deployed open-source LLMs on private infrastructure, tuning smaller models to production quality under data-privacy constraints.

Nexodus: Predictive models & GenAI for a leading flex-space SaaS platform

- Designed and trained churn-prediction models to drive customer retention.
- Built occupancy-forecasting models used for pricing and capacity optimization.
- Led the Generative AI workflows behind automated content creation for the Flexspace Observatory blog.
- Collaborated directly with the client on defining, validating, and iterating solutions.

Bid Genius: Conversational assistant for public tendering (LLM product)

- Early-stage development of a proprietary LLM tool for assisted generation of tender documentation.
- Designed workflows for structured generation of legal and administrative documents, solving challenges around long documents, global coherence, and context management.
- Owned key technical components while working closely with legal and business profiles.

JOT Internet Media: Hierarchical keyword classification with NLP

- Built a multi-stage hierarchical classification system based on RoBERTa for Google Ads keyword ranking.
- Fine-tuned pre-trained language models (Hugging Face) on digital-marketing data; defined taxonomies, text preprocessing, and evaluation aligned with advertising business metrics.

Fen (US Department of Defense): medical AI research applying deep learning to radiology

- R&D project supporting radiologists in the evaluation and diagnosis of low back pain.
- Designed, trained, and evaluated computer-vision models for hernia detection, vertebral segmentation, and detection of scoliosis and lordosis, with rigorous experimental methodology.
- Delivered a web application to drive clinical adoption of the system.

EDUCATION

M.Sc. in Language Technologies (NLP), UNED

Sept 2025 – Present

B.Sc. in Data Science & Engineering, Universidad Carlos III de Madrid

2019 – 2023

Thesis: Reinforcement Learning applied to financial portfolio management.

42 Madrid, Fundación Telefónica

2021 – 2022

Intensive project-based software engineering program focused on autonomy and technical judgment.

SKILLS

GenAI & NLP: LLMs (open-source, private deployment), LangGraph, agentic workflows, Hugging Face, fine-tuning, Agentic RAGs

ML & Deep Learning: predictive modeling, computer vision, time-series forecasting, reinforcement learning, experiment design

MLOps & Cloud: AWS, BigQuery, Snowflake, Databricks, MLflow, distributed data processing, reproducible production pipelines

Leadership: technical mentoring, architecture decision-making, project management, stakeholder communication

LANGUAGES

Spanish: native. **English:** full professional proficiency (Trinity ISE II); daily work with UK and US clients. **French:** basic.