

Sergio Gómez Martín

AI / ML Engineer | Software Engineer

sergio.gomez.martin10@gmail.com | github.com/sergiogomezz | sergiogomezmartin.com | linkedin.com/in/sergio-gomez10
Barcelona, Spain

SUMMARY

AI / ML Engineer experienced in building applied AI systems, from LLM agents and enterprise AI applications to large-scale data pipelines. Strong background in Python, machine learning and software engineering, with hands-on experience translating real-world problems into production-oriented technical solutions. Interested in both applied AI engineering and the underlying models, particularly LLM evaluation, adaptation and inference.

PROFESSIONAL EXPERIENCE

AI Engineer Accenture AI	Jan 2025 – Present Barcelona, Spain
- AI technical lead for an enterprise agentic AI platform at a global automotive manufacturer, building tool-using LLM agents for historical data and internal document retrieval to support procurement decisions. - Designed an LLM evaluation framework combining LLM-as-a-judge with latency and cost metrics to benchmark model variants and agent performance. - Took technical ownership of a large-scale aerospace data pipeline processing 600M+ asset-tracking events, reducing Spark compute hours by 65% and improving performance of the critical processing node by 87%; the resulting application was designed for 300+ enterprise users. - Designed a LangGraph-based multi-agent system for a pharma client to support stock rebalancing across logistics centers through natural language, considering transport costs, resource idle costs, and operational constraints.	
Machine Learning Engineer Ikerlan	Aug 2023 – Jul 2024 San Sebastián, Spain
- Developed an end-to-end data platform on AWS using Apache Iceberg for large-scale industrial data storage and ML model serving in predictive maintenance, validated with 10GB+ of real-world data. - Designed a decision-support framework to assess the spin-off potential of internal research projects based on technical feasibility, market demand, and IP defensibility, contributing to venture-building initiatives that led to the creation of an AI startup.	

PROJECTS

Elvex: Autonomous Architecture for Generation and Coordination of LLM Agents Research Project	2025 – 2026
- Built an autonomous multi-agent architecture that dynamically decomposes complex requests, generates specialized LLM agents, and orchestrates their dependencies and tool use to solve general-purpose tasks. - Engineered the system for reliable execution across multiple LLM providers, with structured outputs, retry mechanisms, observability, FastAPI serving, Docker deployment, and automated testing.	

EDUCATION

MSc in Artificial Intelligence Universidad Internacional de La Rioja (UNIR)	2024 – 2025 Spain
GPA: 9.7 / 10.	
BSc in Computer Engineering Universidad del País Vasco (UPV/EHU)	2019 – 2024 Vitoria, Spain
- Final thesis grade: 9.6 / 10. - Exchange programs: Poznan University of Technology (Erasmus+), Universidad Politécnica de Cataluña (SICUE).	
BSc in Business Administration and Management Universidad del País Vasco (UPV/EHU)	2019 – 2024 Vitoria, Spain
Final thesis grade: 9.5 / 10.	

ADDITIONAL INFORMATION

Technical Skills: Python, Java, SQL, PyTorch, scikit-learn, LangGraph, Langfuse, FastAPI, Spark, PySpark, Pandas, Apache Iceberg, AWS, Docker, Git
AI: LLM Agents & Agentic AI, LLM Evals, RAG, NLP, Machine Learning, Deep Learning
Languages: Spanish (Native), English (C1), German (B1)