

FLÁVIO PEREIRA

Data Scientist, AI Engineer

[linkedin.com/in/flaviofpereira](https://www.linkedin.com/in/flaviofpereira)

(+351) 917997317 ◇ engflaviofpereira@gmail.com

EDUCATION

University of Coimbra, Portugal

2015 - 2018

B.S. in Informatics Engineering

University of Coimbra, Portugal

2018 - 2020

M.Sc. in Informatics Engineering (AI)

EXPERIENCE

Hilti Group

November 2025 - Present

Freelancer AI Engineer

Paris, France

- Sole AI architect/engineer for a Synthetic Agent platform designed to simulate field-user behavior and validate corporate strategic initiatives using LangChain and LLMs.
- Bridged the technical-to-business gap by acting as the primary liaison for Product Owners, translating qualitative business requirements into scalable AI architectures and roadmaps.
- Engineered an Evaluation Framework by reverse-engineering stakeholder interview data into a Golden Dataset to ensure AI alignment with proprietary business doctrines (85% accuracy on real user preference).
- Autonomously built and demonstrated a an AI PoC within 4 weeks, securing immediate stakeholder buy-in for the project's next phase.

RedLight Software

October 2023 - October 2025

Data Scientist

Coimbra, Portugal

- As a research team member, contributed to developing a deep learning pipeline for medical autocoding, fine-tuning foundation large language models (LLMs) (e.g., PubmedBERT) to classify medical codes with a high hit rate.
- Leveraged WandB for experiment tracking, hyperparameter optimization, and maintaining a model registry for version control and reproducibility.
- Implemented logits calibration to produce reliable probabilistic confidence scores, improving the trustability of model predictions. Also, implemented and experimented with PEFT techniques (LoRA and QLoRA) to optimize model performance.
- Designed and deployed AWS SageMaker endpoints for real-time LLM inference, using S3 for scalable storage and ArangoDB for structured data management. Leveraged AWS Lambda and ECS for ML model productization.
- Built probabilistic ML models to evaluate the reliability of medical codes' LLM predictions assigned to medical notes.
- Leveraged Sentence Transformers for text similarity analysis and implemented a keyword detection system to automate semantic information extraction from medical reports.

A.P. Møller – Mærsk A/S

September 2021 - October 2023

Data Scientist

Porto, Portugal

- Collaborated with a customer-facing, agile Data Science team to deliver impactful business results.
- Built and optimized a time series forecasting pipeline (gap filling, outlier correction, cross-validation, and hyperparameter tuning).
- Developed an API for automated model training, validation, and testing, with PowerBI reports for drift analysis.
- Benchmarked models for time series forecasting, including ARIMA, XGBoost, LightGBM, and CatBoost.
- Managed data storage and retrieval using SQL and BigQuery for scalable ML operations.

- Contributed to a post-acquisition migration to tools like GitHub, Jenkins, and Kubernetes, enhancing project workflows.
- Used AWS EC2 for model deployment and scheduled runs, ensuring scalability and reliability.

HUUB

Data Scientist

June 2021 - September 2021

Porto, Portugal

- Joined as the first member of the AI team, which later expanded to five members.
- Maintained and improved Data Science projects deployed on AWS Machines.
- Built an end-to-end orders/products demand forecasting pipeline to streamline predictions and decision-making.
- Applied Scrum methodologies using Jira for project management, with GitLab/GitHub and Dropbox Papers for documentation.
- Worked with Python, Docker, Flask APIs, PostgreSQL, Apache Kafka, Jenkins, Linux Services, Sentry logging to Kibana, and basic Kubernetes for deployment and operations.

Prime IT

Data Scientist

January 2021 - May 2021

Braga, Portugal

- Built an end-to-end NLP pipeline including text preprocessing, vectorization, training, testing, and deployment.
- Conducted model hyperparameter tuning and containerized solutions (Docker) for scalable deployment.
- Applied bootstrap confidence intervals and threshold-moving to enhance results interpretability and performance evaluation.

Gramma

Data Science Intern

October 2019 - October 2020

Coimbra, Portugal

- Applied unsupervised learning to identify patterns in travel behavior and supervised learning to classify transport modes.
- Designed a serverless architecture on AWS using Lambda, S3, and Fargate to host and run Docker containers for asynchronous data transformation and loading into DynamoDB (spatiotemporal data on inhabitants, trips, and city infrastructures). Angular frontend with Chart.js for visualizations and ngx-leaflet maps.

Gramma

Part-time Software Engineer

October 2018 - March 2019

Coimbra, Portugal

- Focused primarily on front-end development using Angular.
- Conducted integration testing with Java, Cucumber, Jenkins, and Selenium, alongside Test-Driven Development (TDD) using Jasmine and Karma.

PROJECTS

Agentic System for Personal Portfolio Management (LangGraph)

Medical Autocoding with Large Language Models (GenAI)

Multi-Agent Task Management System

- Multi-agent system leveraging CrewAI to automate task retrieval, prioritization, and divide-and-conquer strategies.

Time Series Forecasting of Fashion Orders and Products

Medium Tech and Growth Writing

- Ongoing writer for Medium, focused on personal growth and technology topics.

Systematic Review Evaluator - NLP

Intelligent System for Planning Transport Network Infrastructures

TECHNICAL STRENGTHS

Languages	English (Proficient), Portuguese (Native)
ML/DL	Scikit-learn, PyTorch, TensorFlow, LangChain, LangGraph, CrewAI
Natural Language Processing	LLMs, Hugging Face, Sentence Transformers, Logits Calibration, Transformers
Data Manipulation	Pandas, Polars, NumPy, Matplotlib, Power BI
Programming Languages	Python, MATLAB, SQL, Solidity, C, Java
Cloud & DevOps	AWS (SageMaker, Lambda, S3, Fargate, DynamoDB, EC2), Docker, Kubernetes
Databases	PostgreSQL, BigQuery, DynamoDB, ArangoDB
Experiment Tracking	WandB, MLflow, ClearML
Tools	Git, Jupyter Notebooks, Tableau, Jira, DVC, Flask, FastAPI
Miscellaneous	Blockchain, Unix/Linux, LaTeX. Web Crawling/Scraping