

DIOGO BRANCO

AI Engineer | Agentic AI & RAG Systems

Lisbon, Portugal • +351 914 825 299 • diogompcbranco2@gmail.com

[linkedin.com/in/diogo-preto-branco](https://www.linkedin.com/in/diogo-preto-branco) • github.com/DiogoBranco

PROFESSIONAL SUMMARY

AI Engineer who has shipped a production RAG system to the call-centre and retail-store operations of Portugal's largest telco, raising answer accuracy from 40% to over 85% and serving thousands of frontline staff. Deep, hands-on experience across the agentic AI stack — Microsoft Agent Framework, Semantic Kernel, LangChain, LangGraph and Model Context Protocol (MCPs) — combined with production cloud deployment on Azure (AI Foundry, AI Search, Container Apps) and data pipelines on Microsoft Fabric. Part of the Agentic squad at MEO's AI Centre, responsible for the company's agentic AI solutions, and collaborating daily with Microsoft and McKinsey on strategic AI initiatives. Master's in Computer Science from Instituto Superior Técnico (specialisations in AI and Distributed Systems), with a thesis on fine-tuning LLMs for clinical trials.

WORK EXPERIENCE

AI Engineer *Sep 2025 – Present*

MEO (Altice Portugal) — Centro de Inteligência Artificial

- Spearheaded the development and production deployment of an enterprise RAG system for internal knowledge management, raising answer accuracy from ~40% to over 85% since joining the team and rapidly advancing to technical leadership responsibilities within the project.
- System serves the call-centre and retail-store operations of Portugal's largest telco, indexing 3,000+ complex, heterogeneously structured internal documents split into ~50,000 chunks, answering queries in under 30 seconds and saving several minutes of manual search time per query for frontline staff.
- Member of the Agentic squad at MEO's AI Centre — the squad responsible for the company's agentic AI solutions — contributing to architecture and technical decisions alongside the tech lead, and staying on top of the rapidly evolving agentic AI space (models, frameworks, workflows) to guide tooling choices for the team.
- Builds production-grade agentic AI solutions using multiple frameworks — Microsoft Agent Framework, Semantic Kernel, LangChain — and integrates with external tools and data sources via Model Context Protocol (MCPs), selecting the right approach per use case.
- Owns releases of the RAG system end-to-end: containerising services with Docker, deploying to Azure Container Apps through CI/CD pipelines, and running A/B tests and RAGAS-based evaluation to validate changes before rollout.
- Uses Azure AI Search as the main vector database, with custom indexing strategies for the complex internal document structures.
- Contributed to the cloud architecture of MEO's AI Factory — the internal AI platform — working alongside cloud and security teams to provision and govern Azure infrastructure within strict corporate connectivity and compliance constraints.
- Designed and implemented data ingestion pipelines feeding curated datasets into Microsoft Fabric (Lakehouse), enabling downstream AI and analytics use cases.

- Collaborates extensively with cross-functional teams from Microsoft and McKinsey on a daily basis, contributing to strategic AI initiatives and ensuring alignment between technical implementation and business objectives.

Junior Software Engineer *Feb 2023 – Jun 2025*

DataMorfose

- Automated internal company workflows to optimize operational productivity.
- Conducted advanced data analytics and developed automation pipelines to improve data insights.

PROJECT EXPERIENCE

AI Legal Team Platform *2025 – Present*

- Building an AI-powered legaltech platform targeting independent lawyers and small firms in Portugal, focused on automating the reception and triage of judicial notifications.
- Integrating with the Citius judicial platform and the Ordem dos Advogados email infrastructure (IMAP/SMTP) to track incoming notifications, classify them, and auto-draft context-aware responses.
- Architected a multi-agent backend with specialised agents for legal research, contract analysis, compliance, and strategy, combined with a vector-database-backed semantic search layer over Portuguese legal documents.

Agentic Copywriting Automation — Social Media Agency *2026 – Present*

- Building an agentic AI system to automate the copywriting workflow of a social media company, replacing manual content production with an orchestrated team of specialised LLM agents.
- Designing agents for brief interpretation, draft generation, brand-voice review, and final content packaging, with human-in-the-loop checkpoints.

Clinical Trials Natural Language Inference System *2024 – 2025*

- Developed a system for natural language inference on clinical trial reports using the Mistral-7B model, as part of a high-achieving Master's thesis (17/20).
- Fine-tuned Mistral-7B on the clinical trials domain and combined it with an innovative few-shot learning approach using chain-of-thought reasoning to improve robustness on complex medical text.
- Created data augmentation pipelines using GPT-4o to enhance model performance and achieved competitive results on the SemEval-2024 NLI4CT task with strong faithfulness metrics.

EDUCATION

Master's in Computer Science *2021 – 2024*

Instituto Superior Técnico — Specialisations: Artificial Intelligence & Distributed Systems

- Thesis: “Adapting Large Language Models for Clinical Trials: A Comprehensive Approach with Fine-Tuning and Prompt Engineering” — Grade: 17/20.

Bachelor's Degree in Computer Science *2018 – 2022*

Instituto Superior Técnico

Advanced Training Course in Blockchain & Smart Contracts *2022*

Instituto Superior Técnico & Universidade Católica Portuguesa

High School *2004 – 2018*

Colégio Manuel Bernardes

TECHNICAL SKILLS

AI & ML: LLMs, RAG, Agentic AI, Multi-Agent Systems, Microsoft Agent Framework, Semantic Kernel, LangChain, LangGraph, Model Context Protocol (MCPs), Fine-Tuning, Prompt Engineering, Chain-of-Thought, Few-Shot Learning, Embeddings

Evaluation & MLOps: RAGAS, LLM-as-Judge, A/B Testing, CI/CD, Docker, Azure Container Apps, Model Observability

Cloud & Data: Microsoft Azure (AI Foundry, OpenAI, AI Search, Container Apps, Blob Storage, Functions), Microsoft Fabric (Lakehouse, Pipelines), Data Ingestion Pipelines

Vector Databases: Azure AI Search (primary), general experience with vector DB indexing and hybrid search

Languages: Python, TypeScript, JavaScript, SQL

Frontend & Backend: Next.js, React, FastAPI, Node.js

Tools & Methods: Git, GitHub, GitHub Actions, Claude Code, Agile/SCRUM

LANGUAGES

Portuguese (native) • English (fluent)