Carlos André Ferreira

AI Engineering · Agent Orchestration & Evaluation

Viseu, Portugal · carlosandreferreiraei@gmail.com · linkedin.com/in/carlos-andré-ferreira-8367a912a · PT / EN

PROFILE

Senior software engineer and PhD researcher with ~8 years of industry experience, specialised in multi-agent systems: orchestration (LangGraph, Pydantic AI, MCP), agent evaluation (pass^k, Reliability@Cost, failure taxonomies) and observability (OpenTelemetry GenAI conventions). Working thesis: autonomy without measurement is a liability — I build the measurement layer that makes agent systems shippable.

EXPERIENCE

Senior Software Engineer — Critical Software
Large-scale automotive software programme. Client work under NDA.

Web Platform Engineering — European tourism sector
Designed and delivered several multilingual tourism websites and platforms for European markets: content-heavy products where performance, SEO and editorial workflows drive bookings.

SELECTED PROJECTS & RESEARCH

EvalGraph — author
Agent evaluation & self-improvement framework: pass^k, Reliability@Cost, MAST failure taxonomy, Wilson confidence intervals, honesty gate; couples to LangGraph/Temporal via OpenTelemetry GenAI conventions.

CIRA-MAS — first author, accepted at CENTERIS 2026
Multi-agent architecture for cybersecurity incident response (LangGraph orchestration + LightRAG retrieval) with explicit human approval gates.

Agent Skill Library — author
Transferable reasoning skills for coding agents (plan-then-execute, verification gates, debugging methodology, stop-and-escalate, context discipline, failure modes), grounded in empirical transfer research.

Orchestration reference architecture
Three-tier hub-and-spoke design for fleets of coding agents: AGENTS.md as single source of truth, MCP for shared tools, CI gates as the standardisation mechanism.

EDUCATION

PhD candidate, Data Science & Web Technologies — UTAD / Universidade Aberta (ongoing)
Research: multi-agent orchestration, retrieval architectures, agent evaluation.

MSc, Informatics Engineering & Web Technology — UTAD (2025) · final grade 17/20

CERTIFICATIONS

ISTQB® Certified Tester, Foundation Level v4.0 — PSTQB (2025) · CTFL4-2025-00297-PSTQB
Certified Trainer, e-learning specialisation (CCP Formador a Distância + Digital) — IEFP (2025) · FEe-f+769105/2025

TECHNICAL STACK

Agents	LangGraph · Pydantic AI · MCP · OpenTelemetry GenAI
Backend	Python · FastAPI · PostgreSQL/pgvector · Neo4j · Graphiti
Models	vLLM · llama.cpp · Unsloth · PEFT · quantised local inference
Practice	Evals as CI · append-only audit logs · row-level security · HITL gates