Claudio Enrique Bustamante Gallardo



AI Engineer · LLM & Multi-Agent Systems
Seville, Spain · Remote (EU) or hybrid
cebustama@gmail.com
linkedin.com/in/cebustama · claudiobustamante.com

PROFILE

AI Engineer with 10 years in software development, the last 3 building LLM systems. At Logicalis I build a multi-agent assistant for a major Spanish bank that creates change requests in its production ITSM platform. In my systems, the flow is fixed in code and every model output is validated before it is used, which keeps them predictable enough for a bank. M.Sc. in Artificial Intelligence (University of Seville) and former university lecturer.

CORE SKILLS

LLM engineering: Azure OpenAI, AWS Bedrock (Claude), OpenAI, Gemini, Ollama · structured output (Pydantic v2) · model routing · guardrails · prompt versioning
Agents: LangGraph, LangChain, multi-agent architectures · A2A · MCP
Evaluation: LLM-as-judge, rubric design, calibration against labelled data, regression sets
Retrieval: embeddings, vector search (Elasticsearch, Qdrant, FAISS, Chroma), RAG pipelines
Backend & infra: Python 3.12, FastAPI, asyncio, Redis, PostgreSQL, Docker, Kubernetes, Kafka, OAuth2
Practice: hexagonal architecture, pytest, mypy, pre-commit, GitFlow, Jira/Agile
Previously: Unity (C#), AR Foundation (ARCore/ARKit), mobile performance profiling

EXPERIENCE

- **AI Engineer (LLM Multi-Agent Systems), Logicalis Spain** *Apr 2026 – Present*
  Remote · Client: major Spanish bank
  Core developer of a multi-agent LLM assistant that helps the bank's infrastructure engineers write change requests. I built it from an empty skeleton, and it now creates changes in the bank's production ITSM platform.
  - Fixed agent sequence defined in code, with no LLM planner. Model output passes strict Pydantic schemas and a single write point, and anything invalid is discarded.
  - Turned the bank's change-rejection rules into tested checks, so users see what would be rejected while they are still filling in the form.
  - Added an LLM judge for text quality, with a filter that drops any suggestion containing invented data.
  - Long first messages go to a larger model that extracts every field at once. A regex-based detector decides when, and timeouts come from measured latencies.
  - Wrote 900+ tests that run without network, LLM or databases. The component started with none.
  - Ran demos for operations and management, estimated work in Jira, and wrote the handover docs a colleague used to run the component while I was away.
- **Lead Developer (Consultant), APEC Watch** *Oct 2025 – Mar 2026*
  Asia-Pacific Economic Cooperation (APEC) project, hired through a project advisor
  Early warning of aquaculture diseases. Built an LLM pipeline that reads technical documentation and turns it into validated, machine-readable rules for an on-device diagnostic engine, plus the Unity mobile app (with AR photo guides) that runs it.
- **AI Systems Architect (Freelance), BCS Studios SpA** *Sep 2023 – May 2026*
  Designed and built **Eón**, an LLM pipeline for interactive narrative. Several agents with narrow roles,

versioned prompts, an async client that works with any provider, and tooling to track failures and catch regressions.

- **R&D Developer, AI & Virtual Reality Lab, University of Seville** *May 2023 – Aug 2024*
  Worked on the lab's industry projects. The main one was a set of Unity AR apps for **Aertec** (aerospace engineering) to pass on technical knowledge with real-time 3D, tuned to run smoothly on phones.
- **Researcher, ASTER Research Project, University of Seville** *May 2022 – Dec 2022*
  Research project on collaboration between scientists and artists. Designed and built the platform where they co-created works based on research.
- **Founder & Lead Developer, Frisson Games (Chile)** *Oct 2016 – Aug 2024*
  Ran the studio and its development, from prototype to release. Shipped **ColoRise Challenge** (iOS/Android) and **Bird Boxing** (Steam Early Access), made science-communication games with the **ALMA Observatory** and other institutions, and won public grants from Corfo, the Ministry of Culture and Sercotec. Represented the studio at GDC and Gamescom.

PERSONAL PROJECTS

- **Event-driven LLM orchestration lab (in progress).** A deterministic planner decides the steps; LLM agents only analyse evidence and return schema-validated output, failing closed. Kafka event intake and one interface for several providers (Bedrock/Claude, OpenAI, Anthropic, Ollama).
- **Prompt evaluation toolkit.** Scores prompts against a 17-category rubric, with a linter for the objective categories and an LLM judge for the rest. Can run on a local model for confidential prompts.

EDUCATION

- **M.Sc. in Artificial Intelligence, University of Seville** 2021–2022
  Thesis: *Procedural Narratives*. Formal models of narrative and multi-agent simulation with reinforcement learning (PPO/SAC, Unity ML-Agents).
- **Doctoral research in Artificial Intelligence, University of Seville (discontinued)** 2023–2024
  Funded by my R&D work at the university's AI & VR Lab.
- **Computer Engineering, UTFSM (Chile)** 2009–2015
  Ingeniero Civil Informático, a 6-year professional degree from Universidad Técnica Federico Santa María.

TEACHING, LEADERSHIP & LANGUAGES

University lecturer at UTFSM (2020–2023) and Instituto Profesional Arcos (2019–2022), teaching AI for interactive systems, game systems and procedural generation.
Board member of **Video Games Chile**, the national game developers association (2018–2020), as Vice President and, in my last months, President. Represented the industry at GDC and Gamescom, sat on public funding committees, co-authored approved changes to Chile's public game-development fund (higher total budget and per-project cap) and coordinated the LATAM VGF diversity scholarships to GDC.
Languages: Spanish (native) · English (C1, TOEFL) · French (basic)