

Jorge González Conde

AI Agent & Automation Developer

✉ jorgegoco70@gmail.com 🌐 Portfolio 🐙 GitHub 🔗 LinkedIn 📍 Santander, Spain

SUMMARY

AI Agent & Automation Developer who takes AI from prototype to deployed, containerized services. I design and ship AI agents, RAG pipelines and document-AI automation in Python and FastAPI — working day-to-day with Claude, OpenAI, the Model Context Protocol (both client and server), vector databases and LandingAI's document extraction. I write my own agent orchestration rather than using a framework, and I set my own timeouts, retries and validation rather than relying on library defaults. Fifteen years running industrial operations mean I understand the business bottleneck, and the failure modes, before I automate anything. Available for contract work and open to remote AI/Python roles.

TECHNICAL SKILLS

AI & Agents: Anthropic Claude API, OpenAI API, agent orchestration and tool-calling loops, Model Context Protocol (client and server, stdio and SSE), RAG, vector databases (ChromaDB), document AI via LandingAI ADE (OCR / VLM extraction), Ollama (local inference), prompt engineering, LLM output validation

Backend & Deployment: Python, FastAPI, REST API design, async background jobs, Docker, VPS deployment (Contabo / Easypanel), AWS S3 (application level), Linux

Reliability: explicit timeouts and deadlines, retries with exponential backoff, schema validation, structured error handling

Data: PostgreSQL, ChromaDB

Frontend: TypeScript, React, Next.js, Tailwind CSS

EXPERIENCE

AI & Automation Developer — Freelance

Remote | Sep 2022 – Present

- Built a document-intelligence suite on LandingAI's Agentic Document Extraction (ADE), including an end-to-end loan/KYC intake pipeline that classifies financial documents, extracts type-specific fields and **reconciles those fields across documents** to flag contradictions between a payslip, a bank statement and an ID. Also: schema-based extraction with visual grounding on handwriting, stamps and broken tables; and RAG document Q&A with **bounding-box citations cropped from the source PDF**, so a citation is an image of the evidence rather than a page number to trust. Python/FastAPI back ends, React and vanilla-JS front ends, containerized and deployed.
- Built **both sides of the MCP tool loop**: a client (~250 LOC) that connects to any number of MCP servers declared in config and discovers their tools dynamically — turning an M×N integration problem into a clean M+N pattern — and my own MCP server on FastMCP, deployed over SSE, exposing all three primitives (tools, resources and prompts).
- Built a four-agent research system (planner → researcher → writer → editor) that autonomously produces cited academic reports from arXiv and live web search. **Hand-written orchestration over the aisuite SDK — no framework** — with native tool-calling across three search tools, a turn cap so a confused agent can't loop, and the planner's output validated against a required-step contract so a bad plan degrades instead of breaking the run.
- Engineered the reliability layer that AI prototypes usually skip: **explicit 60 s / 90 s timeouts, five retries with exponential backoff restricted to 429 and 5xx responses, and connection pooling**. FastAPI accepts a job and returns a task ID immediately; work runs in the background with real-time progress tracking; results persist to PostgreSQL.
- Built the application layer and bidirectional client actions for a voice agent on the **Vocal Bridge SDK** — the agent drives an on-screen UI (advancing steps, highlighting items, starting timers) while the UI feeds context back to the agent.
- Built an AI job-application tool that analyzes job postings and generates tailored CVs and cover letters, with automatic language detection and PDF export.
- Deployed containerized services on Contabo VPS via Easypanel with custom domains and SSL.

Stack: Python, FastAPI, Claude API, OpenAI API, MCP (client and server), RAG, ChromaDB, LandingAI ADE, Docker, PostgreSQL · [Live demos](#) + [source](#)

Full-Stack Developer — The Ticket Merchant

Remote | Oct 2023 – Sep 2024

- Built an internal web application (Next.js, Node.js, SQL Server) on an API and microservices architecture.
- Ran code reviews and resolved complex integration problems with a multidisciplinary team.

Stack: Next.js, Node.js, SQL Server, REST APIs, Git

Technical Mentor (Volunteer) — Microverse

Remote | Nov 2022 – Present

- Mentor junior developers through detailed code reviews and architecture guidance, focused on code quality and shipping complete full-stack projects.

Stack: React, Redux, JavaScript, Ruby on Rails, Git

Production Manager — MAI

Ourense, Spain | 2002 – 2017

- Ran operations and a 20+ person team: defining what "correct" looks like for a process, measuring when it drifts, and tracing failures to root cause — the process discipline I now bring to automation work. Representative result: ~25% productivity improvement.

EDUCATION & CERTIFICATIONS

Full-Stack Web Development — Microverse | 2022

1300+ hours: algorithms, data structures, full-stack development (Ruby/Rails, JavaScript, React, Redux), pair programming, Git workflows.

Machine Learning, AI & Mathematics Specializations — Coursera, edX, MITx, Stanford, DeepLearning.AI | 2016 – 2021

30+ certificates, including Machine Learning (Stanford – Andrew Ng), Deep Learning Specialization, NLP Specialization, Machine Learning with Python (MITx), Mathematics for Machine Learning (Imperial College London).

Industrial Electronics & Automation Engineering — Universidad de Vigo | 1998 – 2002

Electronics, automation, robotics, control systems, microcontrollers, sensors.

LANGUAGES

Spanish (native) · English (professional working proficiency)