

Sebastian O Rodriguez

sebastian@guavaai.ai · github.com/Sebastian-O-Rodriguez · linkedin.com/in/sebastianorodriguez

Summary

Full-stack engineer shipping production AI systems end-to-end: LLM integration, golden-set evals, RAG with vector search, observability, and bounded agentic loops, under strict security (RLS) and testing discipline (80%+ coverage). Founder of Guava AI (agent orchestration + client BI); previously Scale AI and Salesforce.

Technical Skills

Languages: TypeScript, Python, SQL, JavaScript

Frontend: React, Next.js, Expo / React Native, Tailwind, shadcn/ui, Astro, Svelte

AI / LLM: OpenRouter, LLM evals (LLM-as-judge, golden sets, regression floors), LangGraph, RAG + pgvector embeddings, MCP tool routing, agentic loops, Langfuse / LangSmith observability, Supabase RLS, Claude Code / Cursor

Backend / Infra: Node, NestJS, FastAPI, SQLAlchemy, Pydantic, Polars, PostgreSQL, Prisma, Azure + AWS, Docker, Kubernetes, GitHub Actions

Experience

Guava AI · Founder & Full-Stack Engineer 2025–Present · Remote

- Ship the full Guava AI product stack: agent orchestration (guava-os) via dependency-graph governed execution, worktree isolation, MCP tool routing, review gates, and an eval harness with golden sets + LLM-as-judge (258 tests); client BI analytics (guavabi) with a deterministic engine behind an LLM trust boundary (514 tests).
- Own the build through delivery: scope requirements, run diagnostics, and ship client dashboards and automations in TypeScript, Python, FastAPI, Next.js, and PostgreSQL; set the coverage and eval-release bar for the team.

Scale AI · Strategic Projects 2022–2024 · Remote

- Led text and multimodal data-labeling operations: raised queue quality 25% → 75% in 4 days, scaled daily production 250 → 25,000+, and built an onboarding funnel producing 250+ qualified hires while managing 50+ operators.

Salesforce · Technical Writer 2020–2022 · Remote

- Launched an elastic services product with a 3-person team; owned product and API reference documentation plus the internal knowledge base.

Projects

guava-os · Agent orchestration platform

- Governed AI-agent execution end-to-end: dependency-graph orchestration, isolated worktrees, MCP tool routing, human review gates; golden-set evals with LLM-as-judge. TypeScript.

guavabi · Multi-tenant analytics (BI)

- Deterministic analytics engine behind an LLM trust boundary; converts LLM-constructed queries into deterministic, auditable outputs across tenant-isolated dashboards.

RoutineMe · AI health tracker

- LLM classifier → RAG-grounded nutrition estimator → typed action executor; golden-set evals at 100% on a live OpenRouter run; per-call AI observability; bounded 3-step agentic loop; RLS. Expo, Supabase, OpenRouter.

PMLaD · Multi-tenant property SaaS

- 4-layer RLS tenant isolation (FORCE RLS, boot-time verification); staged Azure CD gated on 80%+ coverage and Playwright E2E. Next.js, NestJS, Prisma, PG16, Azure.

Education

Carnegie Mellon University · BS Technology & Sonic Product Innovation, Minor Business Administration 2020