

# João Nobre

Fortaleza, Brazil

contact@jnobre.com | github.com/joaobnobre

*Builder. Started coding at 7, sold a company at 17, two exits before 21.*

---

## EXPERIENCE

---

### Rummo.ai

Founder

2025 — Present

Fortaleza, Brazil

- Solo-building an AI financial assistant that lets users connect bank accounts and talk to their money in natural language.
- Built the full stack alone: LLM pipeline, backend, frontend. Pivoted from Brazil to the US market after validating product-market fit.
- Own everything: product, infrastructure, distribution, and growth. No co-founder, no team.

### Hypebud

Technical Lead

2023 — 2025

Brazil

- Sole technical hire at Brazil's largest influencer marketing agency. Owned all infrastructure from day one.
- Built live streaming infrastructure (Nginx-RTMP, dedicated ingest servers) delivering sub-2s latency for large-scale outdoor productions with tens of thousands of concurrent viewers.
- Ran production Linux environments: load balancing, CDN integration, multi-platform video distribution, and failover for high-stakes live events.
- Shipped AI-powered video processing and search tools that became core infrastructure across the company.
- In the room for contract negotiations with major names in Brazilian entertainment and sports, including Neymar Jr.

### Pag.Ninja

Founding Engineer | Early equity, successful exit

2022 — 2024

São Paulo, Brazil

- First engineer at a payments infrastructure company that built acquiring and processing rails for merchants. Exited equity stake post-acquisition.
- Built the transaction engine from scratch in Node.js and Go: authorization, capture, refund, chargeback, and settlement flows across multiple acquirers and banking APIs.
- Designed the smart routing layer: dynamically picked the best acquirer per transaction based on approval history, cost, and latency. Lifted overall approval rates.
- Built the fraud and risk stack: real-time scoring, velocity checks, and rule-based automation running on every transaction before authorization.
- Owned the entire infrastructure on AWS (ECS, RDS, ElastiCache) with Kubernetes, zero-downtime deployments, and PCI-DSS compliant tokenization.

### Vigia de Preço

Founding Engineer

2018 — 2021

Brazil

- Joined at 14 as a core developer on Brazil's equivalent of Honey, a price tracking browser extension.
- Grew the product from early stage to acquisition by Mosaico/Buscape in 2021 (~\$1.4M).
- Drove business decisions including onboarding Felipe Neto (one of Brazil's biggest YouTubers) as a distribution partner.

## TECHNICAL PROJECTS

---

### Terminal for Coding Agents

- Built a desktop terminal designed for running AI coding agents in parallel.
- Each task spins up its own git worktree and terminal session, so multiple agents work simultaneously without stepping on each other.
- Live terminal streaming, built-in diff viewer, and real-time agent monitoring from a single interface.
- Works with any CLI-based agent: Claude Code, Codex, Gemini CLI, Copilot.

**Automated Clip Generator for Short-Form Video** | Node.js, AWS Lambda, FFmpeg, Whisper, OpenAI, TikTok/Reels/Shorts APIs

- End-to-end pipeline that watches live streams on Twitch, Kick, and YouTube and auto-identifies highlights using layered AI.
- Layer 1: real-time transcription via Whisper. Layer 2: context scoring combining transcript, chat velocity, sentiment, and humor detection.
- Auto-cuts clips, adds animated captions, and distributes to TikTok, Reels, and Shorts. Zero manual intervention.
- Handles multiple concurrent streams via SQS + Lambda. Under 3 minutes from stream moment to published post.

**Video Search Infrastructure** | TypeScript, Cloudflare Workers, Vectorize, Whisper, OpenCV, DeepFace

- Distributed transcription pipeline with Whisper + semantic chunking for full-text search across a large video library.
- Computer vision layer for face recognition, cataloguing who appears where, when, and in what context across thousands of hours.
- Vector indexing via Cloudflare Vectorize for semantic search combining visual and transcript signals.
- RAG layer for contextual Q&A with streaming responses and source citations.

**TECHNICAL SKILLS**

---

**AI/ML:** LLM integration (Anthropic, OpenAI, Bedrock), RAG pipelines, vector search (SQLite-vec, Vectorize), Whisper, computer vision (OpenCV, DeepFace), AI agents (Mastra, MCP), semantic chunking

**Backend:** TypeScript, Node.js, Go, Rust, Python, Hono, tRPC, NestJS, Effect.js, WebSockets

**Frontend:** React, TanStack (Start, Router, Query), Next.js, Three.js, GSAP, Tailwind CSS, Svelte

**Desktop & Mobile:** Electron, React Native, Swift (iOS/macOS), Kotlin (Android)

**Data:** PostgreSQL, SQLite/Turso, Drizzle ORM, Redis, MongoDB, Convex

**Infra:** AWS (Lambda, ECS, RDS, SQS, ElastiCache, Bedrock), Cloudflare Workers, Docker, Kubernetes, Terraform, Bun, Turborepo

**EDUCATION**

---

**University of Fortaleza (UNIFOR)**

*Systems Analysis and Development | Left to go full-time on Pag.Ninja*

**2023 — 2024**

*Fortaleza, Brazil*

**ACHIEVEMENTS**

---

Math and science olympiad medals in high school: Canguru (math), OBMEP (math), OBQ (chemistry).