

Joaquin Perez

Design Engineer

+1 347-682-0896 | joaquin@hypnolabs.io | New York, NY
www.jperez.io | github.com/bitbrujo | linkedin.com/in/perezjoaquin

SUMMARY

I design and build AI-native products — owning both the interface and the implementation. My work spans UX/UI, coded prototypes, frontend architecture, and production AI workflows across SaaS, adtech, and agentic systems. I close the gap between design intent and shipped software: comfortable in Figma, in the codebase, and in the ambiguous space between them. Recent work includes AI campaign tools, agent dashboards, MCP clients, and LLM-backed generation interfaces.

EXPERIENCE

Blockboard

New York, NY | Jun 2024 – Present

B2B advertising platform using AI for campaign transparency and efficiency.

Design Engineer

- 2025 Digiday Technology Award: Best Programmatic Platform
- Owned product design direction across a 10-person cross-functional team, translating ambiguous AI/adtech workflows into shipped campaign creation, creative generation, and monitoring interfaces
- Designed and built BlockVantage end-to-end — a self-serve CTV product that turns a brand name and URL into a launched, optimized campaign; reduced time-to-first-campaign by 60%+
- Designed BlockVantage's AI decision surface: model-recommended audience segments and estimated reach/impressions presented with real-time confidence caveats
- Designed and built Creative Studio: UI layer over a backend API stack for generating CTV and social video ads, including URL-to-video and AI avatar commercials — turned raw model capability into a polished production interface
- Built the BlockVantage interactive AI demo: 35% demo-to-trial conversion, 60% reduction in LLM API calls via Cloudflare caching, 75% reduction in malicious traffic via Workers-based security
- Designed product flows and prototype architecture for a zero-PII Finix payment integration, aligning UX, compliance constraints, and engineering before buildout

Technical Product UX

- Reduced design-dev cycle time by 33% through workflow redesign and living documentation
- Redesigned BlockCONNECT dashboard: restructured hierarchy, labeling, and monitoring flows — reduced pixel-related support tickets by 70%
- Built reusable component library integrated with RESTful APIs, establishing consistent interaction patterns across AI-backed campaign and reporting interfaces

Motley

New York, NY | 2025 – Present

Family intelligence platform for NYC — neighborhood, transit, safety, and school fit weighted against family priorities.

Founder

- Solo-built and shipped Motley end-to-end — a NYC family-intelligence platform (Next.js 16, React 19, Drizzle/Neon PostGIS, Inngest, Stripe): 63 routes, a 68-table geospatial schema, and 45 background jobs
- Built an agentic in-app assistant — an 18-tool versioned registry with confirm-before-act mutations, tier gating, prompt-injection hardening, and an LLM-as-judge eval harness
- Architected the AI layer: ~20 versioned LLM flows on a multi-provider stack (MiniMax-M3 default, Vercel AI SDK) with per-call cost/latency logging

- Built the data and geospatial pipeline — Exa charter-network scrapers, MySchools admissions monitoring, PostGIS proximity — feeding a MapLibre + deck.gl explorer on a self-hosted PMTiles tileset, plus a standalone MCP server (14 tools) exposing it to agents

Hypno Labs

New York, NY | Jun 2023 – Present

Product studio building for the agentic web — composable interfaces for humans and machine-readable APIs for autonomous systems.

Owner, AI Product Engineer

- Rigger: Open-source (MIT) visual dashboard for Anthropic's Claude Agent SDK — 10-tab configuration surface, real-time SSE streaming with live cost tracking, sandboxed execution
- LoopCraft: Production MCP client with visual UI builder on the MCP-UI protocol — built while the spec was still being drafted with no reference implementations. 5 action types, 3 rendering modes, bidirectional server communication
- Agent Church: Identity and memory infrastructure for AI agents — dual payment protocol (x402/L402) with 7 autonomous completions, A2A protocol with 12 registered skills

Axiom

New York, NY | Jul 2023 – May 2024

Product Designer (Contract)

- Owned end-to-end design for iOS and SaaS products across health-tech and consumer verticals, from discovery through interactive prototyping
- Conducted competitive analysis, user research, and usability testing across multiple engagements

Encompass Digital Media

Stamford, CT | 2021 – 2023

Post-Production Engineer and Cloud Support

- Automated transcoding and QC workflows for 200+ media assets/month — 60% reduction in manual processing
- Reduced content delivery failures by 40% through AWS implementation; containerized applications with Docker
- Built internal Airtable tooling used by 25+ team members for real-time asset tracking

HBO, VICE, A&E, Sony, PBS

New York, NY | 2008 – 2021

Documentary and Television Editor (Freelance)

- Edited 100+ documentary projects for HBO, VICE, PBS, A&E, and Sony; delivered for audiences of 1M+ viewers
- Managed production timelines across 13 years of freelance work, collaborating with 50+ directors and producers

SKILLS

AI and Agentic Systems: Claude Code, Model Context Protocol (MCP), MCP-UI, A2A Protocol, Vercel AI SDK, LangChain, Honcho, Ollama, ComfyUI, Local LLMs

Languages: TypeScript, JavaScript, Python, Golang, HTML, CSS

Frameworks and Libraries: React, Next.js, Node.js, Vue, Tailwind CSS, Radix UI, Framer Motion, Storybook, Playwright

Design: Figma, Framer, Webflow, Design Systems, Interaction Design, After Effects

Infrastructure and Cloud: Docker, AWS, Cloudflare (Workers, KV, DNS), Vercel, Git, REST APIs

Databases: PostgreSQL (Neon), Redis, Supabase, Drizzle ORM, Inngest

Monitoring and Observability: Grafana, Sentry, custom frameworks

EDUCATION

University of Florida, B.S. Telecommunications

Certifications: AI Product Design (ELVTR), Product Design (Flatiron), Google Data Analytics, Google Cybersecurity, CompTIA Network+