Yuvraj Singh Shekhawat

workwithyuvraj0712@gmail.com | Portfolio | GitHub | LinkedIn

Summary

AI Product Engineer with 2+ years of product design experience and full-stack engineering expertise. Specialises in multi-agent AI systems, LLM orchestration, real-time operational platforms, and financial intelligence tooling. Ships production-grade systems end-to-end independently, from UI design through distributed backend architecture. Proficient in Next.js, TypeScript, Supabase, and REST API development with hands-on experience integrating Gemini, OpenAI, and Anthropic APIs.

Projects

- **HiveState** | *Multi-Agent AI Orchestration & Observability Platform* | Next.js · TypeScript · React Flow · Supabase · Gemini API
  - Architected a visual DAG-based pipeline builder enabling multi-agent LLM workflows with dependency-aware parallel execution across 4+ agent nodes, reducing manual orchestration overhead to zero.
  - Engineered a production-grade resilience layer with idempotent execution fingerprinting, circuit breakers auto-pausing after 3 consecutive API failures, and PostgreSQL state persistence enabling mid-pipeline resume without re-billing completed steps.
  - Built X-Ray observability trace exposing per-node token split, micro-cost calculation, and raw LLM payloads per step, delivering transparency that LangGraph and n8n do not natively provide.
- **ResolveOS** | *Intelligent Support Operations Platform* | Next.js · TypeScript · Supabase · OpenAI API
  - Architected a dual-surface support platform with capacity-aware priority routing, replacing chronological queues with weighted SLA-driven execution (5:1 prioritization).
  - Implemented idempotent webhook ingestion via PostgreSQL UNIQUE constraints, SLA state machine enforcement through delayed job queues, and optimistic concurrency control — guaranteeing zero duplicate tickets under high-volume retry storms.
  - Designed graceful degradation architecture with circuit breakers that detect AI latency or rate limits, routing tickets as AI_PENDING and preserving full SLA tracking without dependency on LLM availability.
- **BurnSight** | *Forensic FinOps & SaaS Spend Intelligence Engine* | Next.js · TypeScript · Supabase · Gemini API
  - Engineered a full ETL pipeline normalising chaotic vendor strings via fuzzy matching and a persistent canonicalization engine, feeding deterministic forensic waste detectors.
  - Built Ghost Employee License detection, Free Trial Trap flagging, and Silent Cost Drift analysis using relational joins and time-series logic with full auditability.

Experience

- **Product Designer** 2025 – Present
  Zenia Mobile *Remote*
  - Own end-to-end product design across mobile and web surfaces, covering user research, wireframing, high-fidelity Figma prototypes, and engineering handoff on fast startup timelines.
- **Freelance UI/UX Designer** 2023 – 2025
  Independent Clients *Remote*
  - Delivered production-ready UI/UX for SaaS dashboards, mobile apps, and B2B platforms across multiple clients, building scalable design systems and component libraries.

Technical Skills

Languages: TypeScript, JavaScript, SQL, HTML, CSS

Frameworks & Libraries: Next.js 14 (App Router), React, React Flow, React Query, Framer Motion, Tailwind CSS, Zod, PapaParse

Backend & Data: Supabase, PostgreSQL, REST API design, Webhook processing, ETL pipelines, State machine architecture

AI & LLM Integration: Gemini API, OpenAI API, Anthropic API, Multi-agent orchestration, Prompt engineering, Token cost optimisation

Architecture Patterns: Idempotent systems, Circuit breakers, DAG execution, Optimistic concurrency control, Queue-based SLA enforcement

Design: Figma, Design systems, Component libraries, Prototyping, Responsive UI, Developer handoff

Tooling: Git, GitHub, Vercel, Claude Code, Cursor

Education

- **B.Tech, Computer Science Engineering** Vellore, India
  Vellore Institute of Technology *2022 – 2026*