SUYASH KUMAR SINGH

Frontend + AI Engineer — React / Next.js, LLM Product UI, Agentic Systems
Bangalore, India | [redacted] | [redacted]
github.com/blinderchief | [redacted] | x.com/blinderchief_

Data Science & AI student at IIT Guwahati who has shipped two products as co-founder and engineer, building both the **React / Next.js frontends** and the **Python LLM backends** behind them. Builds the surface where a person works with an AI: live-streamed agent runs, step-by-step traces, source-cited answers, and review steps for output the model is unsure about. Uses **Claude Code and Codex** in daily coding, talks to users, and helps decide what gets built.

EDUCATION

Indian Institute of Technology Guwahati — BSc (Honors), Data Science & AI *2023 – 2027*
CGPA: 8.0 / 10. Coursework: Machine Learning, Deep Learning, NLP, Statistics, Data Structures & Algorithms, Optimization.

EXPERIENCE

Co-founder & Engineer, SnowMind (snowmind.xyz · docs · github) *Feb 2026 – Aug 2026*

- **Sole engineer on a live product** (announcement) with **30+ users and $1,000+ in deposits**, backed by Avalanche & Team1. Built the **Next.js 15 web app in TypeScript**: deposit and wallet-signed withdrawal flows, and a dashboard tracking funded accounts, deposits, and 290+ autonomous rebalances.
- Built the optimization engine (APY polling, 15-min TWAP smoothing, safety-gated rebalancer) and ERC-4337 session keys scoped to supply/withdraw at the EVM level, so a compromised backend cannot move user funds.
- Talked to users directly, found they used the product in situations not planned for, and reshaped it around them. **Stack:** TypeScript, Next.js, Node.js, ZeroDev, Supabase.

Co-founder & Engineer, Reimburse AI (reimburseai.app), Reimburse Inc. — US-incorporated *Dec 2025 – Feb 2026*

- Built a document-AI pipeline using **GPT-4o Vision** to read receipts and validate each expense against company policy, with a **human-review step for low-confidence cases** and eval sets for ambiguous, low-quality, and adversarial inputs.
- Signed **9 letters of intent** from businesses before launch; built the Next.js app and instant non-custodial USDC payouts across four chains. **Stack:** Python, FastAPI, OpenAI, TypeScript, Next.js, Postgres.

PROJECTS

Harness — Self-Hostable Agent Platform with Verifiable Answers (github) *2026*

- **Next.js UI over a FastAPI agent runtime**: chat, agents, knowledge, trace, and settings screens. Runs **stream live over SSE**, and the trace view shows every step (LLM call, tool call, verification) with its token cost.
- **Human control over the agent**: tool permissions (allow / ask / deny), per-run cost budgets, a checkpoint at every step (kill and resume mid-run), and claim verification that labels each answer grounded / stale / unsupported. Any provider via **MCP**. 62 backend tests. **Stack:** Next.js, TypeScript, Python, FastAPI, Postgres.

Warrant — Repair Layer for Production AI Agents (github) *2026*

- Built an **11-screen operator console in React 19**, React Router 7, Vite, and Tailwind v4 on a **hand-built design system** (no component library): incidents open a signed evidence chain, a kill switch can be tripped and reset, and every page reads live from the Python engine. Wrote and revised the PRD from v0.1 to v0.6.

Memora — RAG Knowledge Assistant (github · demo · award) *2nd Prize — Hackathon*

- **Next.js 15 App Router frontend** (TypeScript, Tailwind, shadcn/ui, Framer Motion): chat with source citations, filtered hybrid search, and a knowledge-graph view, over a RAG backend (Qdrant, dense + BM25, reranking; FastAPI, Gemini).

ExpenseLM — Fine-Tuning a 4B Model for Structured Extraction (github) *2026*

- Fine-tuned **Qwen3-4B** (SFT with QLoRA, then DPO on harvested failures) and built a 6-system, 5-metric eval harness; **SFT more than doubled policy-reasoning accuracy** (27% → 55%), which prompting alone could not teach.

ACHIEVEMENTS

- **Avalanche Team1 Hackathon, Bengaluru — Winner** (result). Built SpectraQ, an AI prediction-market agent on MCP (github).
- **AMD AI Sprint Hackathon, IISc Bangalore** (github · report). Served Llama-3.1-8B with vLLM on AMD MI300 GPUs.
- **Aya AI Hackathon, IIIT Bangalore — Runner-up** (announcement). CodeChef 3★ (1855) · CodeForces Pupil (1280).

TECHNICAL SKILLS

Frontend: TypeScript / JavaScript, React, Next.js (App Router), Vite, Tailwind CSS, shadcn/ui, Framer Motion, SSE streaming UI.
AI & Agents: Claude, GPT / Codex, Gemini — tool calling, agent loops, multi-step workflows, MCP servers, LangGraph, RAG, evals.
LLM Training: SFT, DPO, QLoRA fine-tuning (Unsloth, TRL), eval-harness design, vLLM inference serving.
Backend & Tools: Python, FastAPI, Node.js, PostgreSQL, Supabase, Qdrant, Docker, Git; Claude Code and Codex used daily.
Spoken: English (fluent), Hindi (native).