

Aman Jha

9971604279 | email | linkedin | github | portfolio

EDUCATION

Technocrats Institute of Technology
Bachelor of Technology, Artificial Intelligence

Bhopal, India
Nov. 2022 – June 2026

EXPERIENCE

thehodl.fun

Remote

Founder & Lead Engineer

Sept. 2025 – Feb. 2026  

- Built and shipped a **cross-chain token launchpad** on Push Chain from zero to production, owning architecture, smart contracts, backend, and frontend **end to end**
- Scaled to **800K+ transactions** and **500K+ token launches** on testnet; **Runner-Up**, Project G.U.D hackathon (Push Chain)
- Engineered a modular, upgradeable **bonding-curve engine** with constant-product AMM math and virtual reserves, integrated **Uniswap V3** for automatic DEX listing, and hardened contracts to **100% test coverage**

PROJECTS

MetaMemory

Nov. 2025 – Present

TypeScript, PostgreSQL (pgvector), Redis, Neo4j, OpenAI/Gemini 

- Building a persistent memory infrastructure for AI agents scoring **89% task-averaged accuracy** (84.4% overall) on LongMemEval-S (500-question long-term memory benchmark), evaluated with an LLM-as-judge Gemini reader
- Designed a retrieval pipeline combining semantic, temporal, BM25 keyword, and reranking channels, merged via Reciprocal Rank Fusion and re-weighted by a rule layer that learns from prior query outcomes
- Engineered a context-budget-aware pipeline (CABAR) combining greedy token allocation, lexical MMR, recency weighting, and extractive summarization, cutting context usage by **72–76%** at equal benchmark performance

Sealevel: Solana Language Model

Apr. – May 2026

Python, PyTorch, SGLang, Qdrant, MCP, Next.js  

- Coding LLM fine-tuned specifically for Solana/Anchor development: QLoRA fine-tune of Qwen2.5-Coder-7B on **270K Solana-specific records**, scoring **85%** on an 80-task Solana/Anchor benchmark and **90%** on general HumanEval
- Built a RAG pipeline over current Solana docs using Qdrant to keep responses accurate as Anchor syntax evolves
- Served via an OpenAI-compatible API across web, CLI, VS Code, and MCP, with direct integration into Claude Code, Cursor, and Windsurf; **1,500+ installs** on PyPI

Orky

Dec. 2025 – Feb. 2026

Next.js, FastAPI, PostgreSQL (pgvector), Google Gemini  

- Built an enterprise AI assistant answering employee questions across internal tools, with role-based access control enforced at query time so users only see sources they're authorized for
- Integrated 17 live app clients (ServiceNow, AWS, Jira, Salesforce, SAP, Snowflake, SharePoint, and others) and designed a multi-step agentic workflow engine that sequences tool calls with variable passing between steps and mid-workflow file-upload pausing

ACHIEVEMENTS

- Solana Foundation × CoinDCX India Grant:** \$4,000 for tradeX (Feb 2025) + \$10,000 for Lavinth (May 2025)
- Dialect Blinks Instagrants:** \$1,000 for X-Mailer (Jan 2025)
- Agents in Action Hackathon:** 3rd, Akash Track (Best Overall Project), \$2,000 Akash Credits, for PayPerUse
- Zo House Hackathon:** Data Management AI Agent Track, 420 USDT, for DockMind

TECHNICAL SKILLS

Languages: JavaScript/TypeScript, Python, Rust, Solidity, C++

AI/ML: PyTorch, Fine-tuning (QLoRA), RAG Pipelines, Vector DBs (Pinecone, Qdrant, pgvector), Gradient Boosting

Frameworks: React, Node.js, Next.js, FastAPI, Hono

Developer Tools: Git, Docker, Prisma, AWS, Akash, Redis

Databases: PostgreSQL, MySQL, MongoDB, Neo4j