

Akshant Jain

(312) 561-0879 | Princeton, NJ | akshant.jain2@gmail.com | linkedin.com/in/akshantjain | akshantjain.com

PROFESSIONAL SUMMARY

Software Engineer with 5 years of experience and Employee of the Year 2024, owning a production BI analytics module and its AI agent. Builds agent harnesses, tool retrieval, evals and sandboxed execution, from open-source packages on PyPI to a multi-tenant agent platform.

TECHNICAL SKILLS

AI & LLM: LLM agents, multi-agent systems, agent evaluation, LangGraph, LangChain, MCP, RAG, hybrid search (BM25 + embeddings), prompt engineering, context management, Claude, OpenAI and Gemini APIs

AI Dev Tools: Claude Code (hooks, skills, AGENTS.md), Cursor, Google Antigravity, MCP, GitHub Actions/git hooks

Languages & Frameworks: Python, TypeScript, JavaScript, Java, SQL, FastAPI, Next.js, React, Node.js, Spring Boot

Data & Infrastructure: PostgreSQL, MySQL, Redis, Weaviate, DuckDB, Firebase, Spark SQL, Docker, GCP, CI/CD

Testing & Security: Playwright, OAuth 2.0, JWT, RBAC, AES-256 encryption, multi-tenant isolation

PROFESSIONAL EXPERIENCE

InnCreTech LLC

Oct 2021 – Present

Software Engineer (Associate), BI Analytics Lead (Oct 2023 – Present) | Princeton, NJ

Previously Programmer Analyst (Oct 2021 – Sep 2023)

*Project: **DataGOL** (Oct 2023 – Present)*

- Owner and lead of DataGOL's BI analytics module: lead 2 backend engineers, 1 AI engineer and 1–3 QA testers, own the system design and AI agent, and turn product and stakeholder needs into specs.
- Built the first AI features on OpenAI (natural-language charts and insights, ~90% less manual effort), then rebuilt them as a multi-layer BI agent whose design the AI team adopted for other agents.
- Own DAVE, DataGOL's AI analytics agent: it writes queries on selected tables, finds patterns and anomalies, builds charts, widgets and reports, then reads the results back to check its insights.
- Brought back twice to take over the agent after AI-team versions fell behind the BI module; on DAVE, raised eval scores from 2–3 to 4–5 (out of 5), cut latency ~50%, and cut errors and bad configs ~2/3.
- Restructured DAVE's skills, prompts and workflows with a config validation loop, tests and a Playwright + content-check eval framework, keeping token use in check and matching BI widget features.
- Took over a client's AI agent (grocery app) on a 2-week deadline: fixed its pipelines and logic for 2–3x lower latency, cut prompt size and cost, and advised on next features like memory and alerts.
- Designed a transformation layer on top of the BI module so manual and AI-generated analytics share one format, letting new structures be added later without changing the BI internals.
- Architected the dashboard engine (drill-down, conditional formatting, optimized SQL, caching, parallel queries) for sub-second responses, plus white-label embedding and PDF/Excel/CSV reports.

*Project: **TermGrid** (Oct 2021 – Present)*

- Helped build and scale TermGrid's SaaS deal platform, which processes \$50B+ in transactions for 600+ institutional clients at 99.9% uptime, through its venture funding with Asymmetric Capital Partners.
- Built an AI document feature that turns uploaded financial documents into structured financial data and lets users ask plain-English questions about the extracted results.
- Migrated the platform from Angular.js to React.js, cutting API calls by 65% and page load time by 55%.
- Built an NDA security layer so sponsors can flag confidential deal information to lenders, and a coverage module that keeps the full history of lender–sponsor interactions.
- Built SOC 2-compliant multi-tenant collaboration tools for sponsors, lenders and advisors (encryption, RBAC, audit trails) and REST microservices handling 1M+ daily API calls at 99.95% reliability.

- Served on the production monitoring rotation, fixing ~10 production bugs and shipping ~5 patches per month.
- Moved into a lead and mentor role from 2024 as focus shifted to DataGOL: guided newly hired engineers, reviewed code, and ran technical interviews for new candidates.

Project: CoAM (Additive Manufacturing Platform) (Jan 2023 – Jan 2024)

- Automated material planning and calculations across sites, cutting errors 50% and manual effort 95%.
- Built a library of reusable React components for the platform's new React front-end framework.

AI ENGINEERING PROJECTS

VisvoAI, AI Agent Platform

Sep 2025 – Present

Open-core AI agent platform: two open-source Python packages on PyPI and a multi-tenant product built on them (FastAPI, LangGraph, Next.js, PostgreSQL, Redis, Weaviate).

Open-source packages (visvoai-ai, visvoai-core)

- Built visvoai-ai: a LangChain provider layer with a live model registry (pricing, context, capabilities), cost metering and normalized reasoning levels across Anthropic, Gemini and OpenAI-compatible.
- Built visvoai-core: a LangGraph agent runtime with override hooks for nodes, routing, state, checkpointing and persistence so products extend it without forking and a step cap forcing a final answer.
- Wrote hand-built BM25 + cosine tool search, generalized from the platform into the core package.

Platform

- Engineered the agent harness and loop on LangGraph: a custom tool-execution node with parallel batching, repeat-call blocking, schema-guided recovery from bad arguments, step caps and context compaction.
- Built dynamic tool binding to stop loading every tool (~57% of each prompt): per-step retrieval over Gemini-embedded MCP tools; rewrote tool descriptions to lift recall@8 from 91% to 97% (211 cases).
- Built a tool system of ~40 tools with roles and per-tool permission checks before execution, plus human-in-the-loop approvals (approve, edit, approve-all, reject) that pause and resume the agent.
- Built custom agents, 17 on-demand skills and subagent calling with isolated context, depth caps and parallel runs that stream live to the parent's UI and run as the parent's user.
- Built versioned artifacts (docs, code, slides, sheets, charts) that agents and users co-edit with conflict checks, SQL over sheets in DuckDB, and a network-isolated sandbox where agents run code.
- Built RAG over a per-user document Library: PDF/DOCX/PPTX parsing, heading-aware chunking, Gemini embeddings and hybrid BM25 + vector search in Weaviate with citations, plus long-term chat memory.
- Wrote an MCP client and OAuth connectors from scratch (Google, Microsoft, Slack, Notion, HubSpot); it syncs server tools by content hash, repairs malformed LLM arguments, and moves large results to files.

Repolane, Open-Source AI Developer Tool

repolane.dev | Sep 2026 – Present

Open-source (Apache 2.0) tool for running AI coding agents like Claude Code safely across repos.

- Designed a workspace model on git worktrees (one clone per repo, one workspace per task) so several AI coding agents can work in parallel across repos without conflicts or mixed-up context.
- Built a Python rules engine on Claude Code hooks that checks every agent action, including shell commands and file paths, to keep agents in scope, block access to secrets and protect main branches.
- Guarded the rules with a 184-case automated test suite and live self-checks.

ACHIEVEMENTS

Employee of the Year 2024, InnCreTech LLC

- For leading the BI analytics module and its contribution to winning new customers.

EDUCATION

University of Illinois at Chicago, Chicago, IL

May 2021

- Bachelor of Science in Computer Science
- Minor in Mathematics