

Kartik Joshi

408-531-7646 | kxj230026@utdallas.edu | [linkedin.com/in/kartikjos](https://www.linkedin.com/in/kartikjos) | github.com/KJ14GOD | x.com/kjcodes | Dallas, TX

EDUCATION

University of Texas at Dallas - Academic Excellence Scholarship

GPA - 3.8

Bachelor of Science in Computer Science

May 2028

Coursework: Data Structures and Algorithms, Computer Architecture, Machine Learning, Database Systems

TECHNICAL SKILLS

Languages: Java, Python, C, C++, TypeScript/JavaScript, SQL

Frameworks: Spring Boot, React, FastAPI, Node.js, Electron

Backend & Data: REST APIs, distributed systems, PostgreSQL, OpenSearch, AWS cloud services, DynamoDB, Docker

Engineering: C++/Python systems, concurrency, caching, testing, CI/CD, AWS CDK, Linux, Git, metrics/logs/traces

Agent Systems: Claude Code, MCP, LLM APIs, agent evaluation

Honors: Goldman Sachs Intern Hackathon Winner; HackRice 2025 – 1st Place; AIME Qualifier; USACO Gold

PROFESSIONAL EXPERIENCE

Engineering Summer Analyst

May 2026 – Aug 2026

Goldman Sachs

Dallas, TX

- Built an agentic change-intelligence control plane that decomposed Anaplan model diffs into parallel schema, dependency, risk, and downstream-impact investigations, reducing reviewer effort by **67%**.
- Designed a durable agent runtime with shared state, resumable checkpoints, idempotent tasks, tool permissions, and queue backpressure so hundreds of section-level analyses could run concurrently without losing work.
- Built tool-using agents that traversed report history, searched model context, compared prior changes, and assembled a provenance graph linking every recommendation to its source evidence.
- Added verifier and adjudicator agents that challenged unsupported findings, reconciled conflicting analyses, and escalated ambiguous changes to human reviewers, cutting end-to-end enrichment time by **10×**.

Software Engineer Intern

July 2025 – September 2025

Scend.ai

Remote

- Owned the agentic market-research workflow behind Scend's M&A platform, turning natural-language market theses into structured research on growth drivers, key players, consolidation trends, and acquisition targets.
- Built the production data-acquisition layer with Python, Node.js, and Firecrawl, orchestrating web discovery across **5 private-market entity types** and normalizing results covering **189K+ companies** and **14K+ investors**.
- Hardened the research pipeline across **125K+ enrichments** with structured-output validation, result filtering, deduplication, response caching, and PostgreSQL-backed persistence for repeatable market reports.

Agentic AI Researcher @ Professor Xinda Wang Lab

May 2025 - Aug 2025

NSF REU

Dallas, TX

- Selected for a competitive **NSF REU** focused on agentic AI and software security research.
- Developed **MAVUL**, a multi-agent vulnerability detection framework using **LangGraph** and **ReAct** loops in which agents collaborate, critique, and refine security findings.
- Evaluated MAVUL on **879 real-world C/C++ vulnerable code pairs** from PrimeVul against GPTLens, VulTrial, and GPT-4o.
- Published "MAVUL: Multi-Agent Vulnerability Detection via Contextual Reasoning and Interactive Refinement" at the **IEEE TPS 2025 Conference**.

PROJECTS

Novus – Multiplayer Coding-Agent Control Plane | TypeScript, Electron, Fastify, PostgreSQL, Claude Code July 2026

- Built Novus as a multiplayer mission room for teams to collaborate around coding-agent work, sharing context, live directions, approvals, handoffs, checkpoints, and isolated Git worktrees.
- Designed concurrent agent coordination across multiple workstreams and sessions with scoped parallel writers, queued directions, durable control leases, and session continuity.
- Implemented supervised autonomy with **5 permission profiles**, role-based capabilities, approval routing, force-interrupt recovery, secret isolation, and governed project skills and MCP servers.
- Built the delivery lifecycle from agent execution to Git checkpoints, verification evidence, branch comparison, pull-request creation, review readiness, merge, and a snapshotted completion receipt.

Robinhood Agentic | Python, FastAPI, asyncio, Anthropic Claude API, SQLAlchemy, PostgreSQL

May 2026

- Built an autonomous portfolio research engine that grounds recommendations in read-only Robinhood holdings, yfinance market data, RSS news, SEC EDGAR filings, and persistent research memory while keeping trade execution human-controlled.
- Implemented a custom **Anthropic Claude tool-use loop** with **16 typed tools** and a 12-step execution budget for screening stocks, inspecting portfolios, retrieving signals/news/charts, managing watchlists and missions, and logging recommendations.
- Orchestrated researcher, analyst, monitor, and judge workflows with **FastAPI + asyncio**, routing scheduled and event-triggered work through cooldowns, failure isolation, persistent state, and streaming audit traces.
- Built an **agent evaluation and self-improvement harness** with LLM-as-judge scoring, human failure labels, one-pass repair gates, and shadow-portfolio benchmarking against SPY and sector indexes.