

ADITYA CHOUDHARY

Chicago, IL, USA | 312-608-6144 | adityacaug15@gmail.com | [LinkedIn](#)

Profile Summary

AI engineer with 2 years of experience building production LLM agents and RAG systems on Azure and AWS, from scoping and prompt engineering through evaluation and post-launch monitoring. Strong Python, SQL, and cloud foundation, with recent focus on RAG architecture, agent drift detection, and LLM-driven workflow automation. Comfortable owning the full technical lift and tying builds to measurable outcomes. AWS-certified, H1B-transferable.

Technical Skills

- **AI & Agents:** RAG architecture, LLM prompt engineering, ground-truth evaluation, agent drift monitoring, retrieval tuning, confidence thresholds, OpenAI / Anthropic APIs, LangChain, agentic coding workflows (Claude Code)
- **ML & Data:** Feature engineering, model evaluation, ETL pipeline design, advanced SQL, Pandas, embeddings, vector retrieval
- **Programming Languages:** Python, Java, C, C++, SQL
- **Cloud & Infrastructure:** Azure (AKS, Functions, OpenAI Service, Blob Storage, SQL), AWS (Lambda, S3, Redshift, Bedrock), Docker, Hybrid Cloud
- **Data & Analytics:** Tableau, PowerBI, Looker, Informatica, Excel (advanced modeling)
- **Systems & Integration:** Distributed systems, API connectivity, IAM, RBAC, cloud security, data governance

Experience

RAMS Contracting, Cloud Automation Engineer

January 2025 - Present

- Built an end-to-end internal RAG chatbot on Azure, owning scoping, prompt engineering, retrieval tuning, and deployment, serving ~50 internal users across field and office operations with ~150 queries handled per week
- Established a ground-truth evaluation set and feedback logging to detect agent drift, tightening grounding prompts and adding confidence thresholds to prevent out-of-corpus hallucinations in production
- Designed retrieval pipelines over ERP, project management, and field data sources, integrating Vista ERP, Procore, and PlanGrid as grounding corpora for the agent
- Partnered with leadership to define the AI and automation roadmap, prioritizing agent and LLM workflow opportunities across operations, finance, and field teams
- Translated technical proposals into ROI-driven business cases, securing approval for a Python ETL pipeline rebuild by framing capacity reclaim and error-rate reduction
- Authored advanced SQL queries for ETL validation, cross-system reconciliation, and performance auditing across Vista ERP, Procore, and PlanGrid

Noble House Solutions, Software Developer Intern

June 2024 - August 2024

- Improved ML model accuracy by ~25% through feature engineering, optimized preprocessing, and iterative model refinement for conversion prediction
- Developed AWS ETL pipelines using Python, Docker, Lambda, and S3, increasing processing throughput by ~35% to support scalable ML training and inference workloads
- Integrated and modeled data from 12 disparate sources into AWS Redshift, improving query efficiency and supporting downstream ML workflows

Amazon, Business Operations Specialist (QA)

January 2022 - July 2022

- Coordinated requirements, testing, and validation across engineering, product, and operations teams in support of high-impact system migrations
- Drove root-cause analysis and defect resolution through structured triage workflows and cross-functional collaboration, accelerating product improvements

Projects

OptionSpace Beta, UI/UX Strategy Audit (Product review, stakeholder analysis)

- Partnered with the product team to audit novice-user workflows across Vol Smile, Term Structure, and Vol Screener screens, producing a prioritized friction-reduction report grounded in stakeholder interviews and user journey mapping

Stock Market Performance Dashboard (Power BI, Python, SQL, Azure)

- Designed Azure-based ETL architecture using Functions, Blob Storage, and SQL to automate ingestion, reducing manual reporting by ~50% via a scalable refresh framework
- Built an analytics dashboard with Sharpe Ratio, CAGR, and risk indicators enabling AI/ML-assisted insights for portfolio decisions

Education

Master's in Computer Science, DePaul University, Chicago, IL

September 2022 - August 2024

Certifications

AWS Certified Cloud Practitioner (November 2024) | Oracle Certified Professional, Java SE 8 Programmer (March 2021)