

Abdullah Rahman

Senior AI Engineer • LLMs • RAG • Agentic Systems • MLOps
Queens, NY • abdullah2891@gmail.com • 929-231-6062 • LinkedIn • GitHub

PROFESSIONAL SUMMARY

Senior AI/Software Engineer with 8+ years building production-grade AI and distributed systems. Specialized in Large Language Models (LLMs), Retrieval-Augmented Generation (RAG), agentic AI pipelines, and natural-language interfaces. Hands-on with GPT-4, GPT-4o, LangChain, Azure AI/Vector Store, prompt engineering, and end-to-end MLOps. Track record of securing stakeholder buy-in, leading cross-functional teams, and delivering measurable business outcomes at enterprise scale.

EXPERIENCE

Evernorth (Cigna)

New York, NY

Senior AI / Software Engineer

Oct 2019 – Present

- **GenAI & LLM Leadership:** Technical lead for all GenAI/LLM initiatives — architected an agentic Scenario Agent (LangChain + GPT-4o) that orchestrates existing APIs for in-product scenario creation (secured full stakeholder buy-in), and shipped a GPT-4-powered natural-language query interface giving senior leadership self-serve access to complex analytics data.
- **RAG Pipeline & AI Chatbot Productionization:** Led end-to-end productionization of an enterprise AI chatbot (GPT-4o) through a rigorous internal Center of Excellence (CoE) production readiness review. Built a full RAG pipeline with LangChain, agentic retrieval, and Azure Vector Store; engineered a custom web crawler with intelligent chunking and data-cleaning logic for high-quality embeddings and reliable retrieval.
- **Scalable Backend & Distributed Systems:** Designed high-throughput distributed backend (Python, PostgreSQL, AWS, MongoDB) supporting 1,000+ facilities each processing ~10,000 claims. Built responsive React.js frontends with Cypress end-to-end testing for rigorous quality assurance.
- **ETL & Data Engineering:** Built AWS Batch, AWS Glue, and Docker ETL pipelines ingesting and transforming 150k+ hospital contracts, validated against 10M+ historical claims with 90% test coverage from scratch — earning the Dragon Award for significantly enhancing user experience and trust.
- **Feature Expansion & Database Migration:** Introduced pinpoint contract analysis and comparison features, migrated from a non-relational to a relational (PostgreSQL) database, and integrated Ambulatory Payment Classification (APC) payment systems — growing the user base to 1,000+ contractors org-wide.
- **Technical Leadership & Mentorship:** Spearheaded multiple projects from inception to production, owning requirements gathering, architecture, and delivery. Mentored junior engineers and introduced AI-assisted developer workflows that measurably improved team velocity.

AT&T

Durham, NC

Software Engineer — Web Applications

July 2016 – May 2019

- **Full-Stack Development:** Built and maintained a large-scale internal web application for AT&T, integrating multiple internal tools into a unified platform and delivering 400+ story points across features and bug fixes, earning commendations for exceptional performance.
- **Technology Stack:** Developed using Ember.js, React.js, and Perl, ensuring scalability, reliability, and consistent cross-browser performance.
- **Database & Systems Integration:** Integrated 25 SQL Server databases (200+ tables) into a complex ticketing system; managed security permissions and access control for thousands of users, user groups, and global clients.
- **Code Quality:** Performed daily code reviews for efficiency, correctness, and maintainability using CodeCloud, consistently maintaining high engineering standards across the team.

EDUCATION

CUNY City College — B.S., Electrical Engineering

New York, NY

Relevant Coursework: Algorithms, Numerical Analysis, C++ for Engineers

TECHNICAL SKILLS

AI / ML: LLMs (GPT-4, GPT-4o), RAG, LangChain, Agentic AI, Prompt Engineering, Embeddings, Vector Databases, NLP, Azure AI / Azure Vector Store, Generative AI, MLOps

Languages: Python, JavaScript, Java, Kotlin, Swift, Rust

Backend: Django, Flask, REST APIs, Microservices

Frontend: React.js, Angular, Ember.js

Databases: PostgreSQL, MySQL, MongoDB, Azure Vector Store

Cloud/DevOps: AWS (Batch, Glue, S3, Lambda), Azure, Docker, Terraform, CI/CD