

ABHISHEK SHANKAR

• 774-519-0926 • Mountain View, CA • abhishekshankar88@gmail.com • www.linkedin.com/in/abhishekshankar97/

Work Experience

RhythmX AI

AI Engineer Consultant

December 2024 – Present

Palo Alto, CA

- Led the end-to-end scaling of a LLM optimised clinical reasoning engines, expanding coverage from 4 to 200+ medical conditions by optimizing prompt orchestration, PyTorch-based model pipelines, and GPU-accelerated containers on Azure
- Designed and executed LLM evaluation benchmarks across 5+ GPT model variants to assess accuracy, latency, and cost trade-offs, driving model selection decisions that reduced API overhead by \$4,000/week while sustaining 98%+ accuracy
- Engineered GenAI orchestration and observability dashboards, accelerating cross-functional debugging by 50% and providing senior leadership with real-time analysis to optimize system performance.
- Deployed production-ready AI features on Azure Container and App Services, using Dockerized workloads, CI/CD pipelines, and Kafka-based orchestration, enabling scalable and reliable model deployment in a cloud-native environment
- Engineered high-concurrency inference workflows in Python, slashing system latency by 50% through identifying architectural bottlenecks and eliminating redundant upstream data processing

Boehringer Ingelheim

Data Scientist Co-op

July 2023 – December 2023

Ridgefield, CT

- Built a GraphRAG pipeline using Stardog and LLMs to automate knowledge extraction from 1,000+ unstructured supply chain PDFs, drastically increasing reporting transparency
- Automated GTM (Go-To-Market) dashboards for 10+ drug launches using AWS CDK and Databricks, providing leadership with real-time AI-driven market strategy recommendations
- Engineered custom ETL scripts in Python to ingest and validate high-volume parquet data into AWS Athena, ensuring data integrity for downstream Business Intelligence (BI) applications
- Spearheaded the technical transition to Databricks by developing a suite of 25+ Python and PySpark tutorial notebooks, standardizing AI development practices across the enterprise team

Adqvest Capital Pvt. Ltd.

Data & Software Engineer

April 2020 – June 2022

Chennai, India

- Deployed a fleet of 100+ Python-based scrapers and managed AWS RDS (SQL) instances to process gigabytes of alternative financial data for real-time market analysis
- Implemented custom NLP pipelines using SpaCy and PageRank algorithms to generate automated sentiment reports, directly contributing to the retention of 5+ high-value institutional clients
- Architected and deployed 10+ statistical analysis projects and interactive dashboards, enabling data-driven capital allocation and risk assessment for complex financial portfolios

Projects

SupplyChainInsights - Advanced GraphRAG based Chatbot ([GitHub](#) & [Medium](#))

May 2024 – August 2024

Python, Neo4j, OpenAI and Streamlit

- Built a GraphRAG-based chatbot using Neo4j and OpenAI to analyze relationships across 1000+ HIV drug supply chain entities
- Implemented a multi-agent workflow that delivers graphs, summarized answers, Cypher queries, and key metrics for analyzing AI-driven responses based on user queries
- Leveraged Graph Data Science similarity algorithms and OpenAI embeddings along with LangChain to utilize vector searches & create recommendations

Education

Northeastern University

Master of Science in Information Systems (Major in Data Science)

August 2024

Boston, MA

SRM Institute of Science and Technology

Bachelor of Technology in Mechanical Engineering

May 2019

Tamil Nadu, India

Skills

LLM & GenAI: Prompt Engineering, RAG (Vector Graph), LLM Evaluation, MCP (Anthropic Certified), Model Benchmarking, Temporal Agentic Frameworks, LangChain, Claude Code

Programming & ML: Python (Expert), JavaScript, PySpark, SQL, Scikit-Learn, TensorFlow, PyTorch

Cloud & Data: Redis, Google Cloud Platform (GCP), AWS (CDK, Athena, RDS), Databricks, Snowflake

Databases: Neo4j, Stardog (Knowledge Graphs), PostgreSQL, MongoDB, Clickhouse

DevOps & Tools: CI/CD, Streamlit, Git, Docker, Jira/Confluence, Tableau, PowerBI