

AMOGH RAMAGIRI

Tel: +1 571-478-2290 | Email: amoghr@gwmail.gwu.edu | [LinkedIn](#) | [Github](#) | [Publications](#) | [Portfolio](#) | Washington, DC

EDUCATION

George Washington University

Master of Science in Data Science (CCAS Dean's Award)

Aug 2024 – May 2026

Washington, DC

Presidency University

Bachelors of Technology in Computer Science and Engineering

Sept 2020 - May 2024

Bengaluru, KA

PROFESSIONAL EXPERIENCE

George Washington - IT Department

AI Engineer

Jan 2026 - May 2026

Washington, D.C.

- Built an agentic ticket triage system (**LangChain** orchestration, **Pinecone** vector store) serving **44,000+** users, autonomously classifying and clustering **100,000+** tickets to surface **120+** recurring patterns.
- Engineered a **RAG pipeline** with **Azure OpenAI** summarization that auto generated **knowledge base articles** for top recurring issues, cutting **analyst triage** effort by **43%**.
- Shipped the **system end to end**, **React**, **TypeScript**, and **Tailwind CSS** frontend over a **PostgreSQL** database, deployed on **AWS EC2**, **RDS Postgres** in production.

Fulton Bank

Data Scientist

May 2025 – Dec 2025

Lancaster, PA

- Built mortgage default prediction models using Random Forest processing **1M+** data points, identifying high-risk applications with **91% recall** and reducing potential loan losses by **18%**
- Streamlined **20+ SQL queries** for credit risk reporting, reducing processing time by **15%** and improving data accuracy for the Lending team across business lines
- Migrated on-premises data infrastructure to **Azure** (Data Factory, Synapse), achieving **20% improvement** in data quality and **43% faster** system performance
- Rebuilt and optimized **37+ Power BI dashboards** and published to Microsoft Fabric on Azure, enabling real-time analytics and supporting faster decision-making across major business lines

Data Science for Sustainable Development, GWU

Data Scientist, AI Systems

Sep 2024 – Oct 2025

Washington, DC

- Developed conversational AI assistant using **LangChain** and fine-tuned **LLaMA-3**, processing **55,000+** queries and cutting research time by **85%**, reducing inference costs by **60%** compared to GPT-4 baseline through local model deployment
- Fine-tuned **LLaMA-3-8B** using **LoRA** on **10,000+** conflict resolution documents, evaluated with RAGAS framework (faithfulness, context precision, answer relevance), improving response quality by **38%** and reducing hallucination rate by **45%**
- Implemented vector search using **PostgreSQL with pgvector** across **107,000+** humanitarian records with retrieval evaluation pipeline tracking precision@k and MRR, improving search accuracy by **55%** over BM25 baseline

Factocart, Velabh Technologies Pvt. Ltd.

Data Scientist

Mar 2024 – Aug 2024

Bengaluru, KA

- Processed **2M+** transaction records from AWS RDS using **PySpark** and SQL, built demand forecasting models with Scikit-learn achieving **92% accuracy**, reducing stockouts by **13%** across **10+** categories
- Conducted hypothesis testing and A/B experiments to evaluate marketing campaigns, identifying high-impact strategies that increased conversion rates by **15%** and drove **11% sales growth**
- Developed LLM-powered insight pipeline using **LangChain**, OpenAI APIs, and **Pinecone**, reducing manual analysis time by **60%** and cutting API costs by **35%** through prompt chaining and few-shot optimization

TECHNICAL SKILLS

Programming: Python, SQL, JavaScript/TypeScript, R, C

AI/ML Frameworks: PyTorch, TensorFlow, scikit-learn, XGBoost, Hugging Face Transformers, SHAP

LLMs & GenAI: OpenAI API, LangChain, CrewAI, RAG, FAISS, ChromaDB, Prompt Engineering, Fine-Tuning (LoRA), MCP

NLP: BERT, Sentence Transformers, spaCy, TF-IDF, NER, Text Classification, Embeddings

Full-Stack & APIs: React, Next.js, FastAPI, Flask, RESTful APIs, WebSockets, SSE

Cloud & MLOps: AWS (EC2, RDS, S3, BedRock, EMR, DynamoDB), Azure, Hadoop, Docker, Git, CI/CD, MLflow

Data & Visualization: PostgreSQL, DuckDB, MongoDB, PySpark, Pandas, NumPy, Tableau, Power BI

PROJECTS

Cloudant.ai, Agentic Infrastructure Design Tool (AWS)

2026

- Built an **agentic AI** tool that crawls a GitHub repo and generates a **production ready AWS architecture** diagram with cost estimates in **under 60 seconds**, replacing hours of manual infrastructure planning.
- Designs complete AWS architectures from a plain language **product description**, mapping **VPC networking** and security across **200+** AWS services for **production** deployment.

Conversational Data Analyst Agent

2025

- Deployed LLM-powered text-to-SQL agent using **LangChain** and **GPT-4** with RAG-based schema retrieval, achieving **90%+ query accuracy** and reducing manual SQL requests by **50%**, containerized with **Docker**, served via **FastAPI** (sub-300ms latency), with MLflow tracking and GitHub Actions CI/CD