

AMOGH RAMAGIRI

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

EDUCATION

George Washington University

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

Aug 2024 – May 2026

Washington, DC

Presidency University

Bachelor's of Technology in Computer Science and Engineering (Spl. in AI and ML)

Sept 2020 – May 2024

Bengaluru, KA

PROFESSIONAL EXPERIENCE

Fulton Bank

Data Scientist

May 2025 – Dec 2025

Lancaster, PA

- Built **mortgage default prediction** model (Random Forest) on **140,000+ loan records**, achieving **91% recall** on high-risk applications and reducing potential portfolio losses by **18%** (~\$2M+ annually).
- Optimized **20+ SQL queries** for credit risk reporting, cutting **execution time 15%** and eliminating 4+ hours of weekly manual data reconciliation across the Lending team.
- Led **Azure cloud migration** (Data Factory, Synapse Analytics), improving data quality **20%** and accelerating pipeline throughput **43%**, enabling real-time reporting.
- Redesigned **37+ Power BI dashboards** on Microsoft Fabric, consolidating 6+ data sources and cutting **executive reporting cycle from weekly to daily**.

Factocart, Velabh Technologies Pvt. Ltd.

Data Scientist

Mar 2024 – Aug 2024

Bengaluru, India

- Built SKU-level **demand forecasting** models (Random Forest, XGBoost) achieving 92% accuracy after **hyperparameter tuning** via cross-validation, **reducing stockouts 13%** across 10+ product categories.
- Engineered **30+ behavioral and transactional features** from **2M+ customer** records (AWS RDS, Python), surfacing acquisition patterns that **drove 11% sales lift** across targeted segments.
- Designed and executed **A/B experiments** with statistical **hypothesis testing** across marketing campaigns, identifying high-performing strategies that **lifted conversion rates 15%**.

Embedded System Lab, NCUE

Research Scientist

Sep 2023 – Mar 2024

Changhua, Taiwan

- Engineered **multi-modal biometric authentication** system (DeepFace CNN + MFCC audio features, PyTorch) for offline edge deployment on Raspberry Pi, improving authentication accuracy **30%** over single-modal baseline.
- Reduced **false rejection rate 25%** via statistical root-cause analysis on **10,000+ authentication failures**, identifying and resolving 4 critical failure modes across lighting and audio variance conditions.
- Architected **edge-to-cloud data pipeline** (Flask, REST APIs) streaming biometric records from Raspberry Pi cluster to centralized training infrastructure with **99% data integrity** and sub-200ms latency.

Wodo

Data Analyst

Jun 2022 – Sep 2023

Bengaluru, India

- Automated **marketing performance reporting pipeline** in Python and SQL across **13+ client accounts**, consolidating data from 5+ ad platforms into a unified dashboard and reducing weekly reporting time by **~75%**.
- Built **real-time engagement and purchase dashboards** enabling clients to reallocate marketing budgets dynamically, reducing **wasted ad spend by 20%** and improving ROAS across paid channels.
- Conducted **customer segmentation** using K-Means and hierarchical clustering on behavioral data, identifying 5 high-value audience cohorts that improved campaign targeting precision and **lifted conversion rates 12%**.

TECHNICAL SKILLS

Programming and Tools: Python, R, SQL, Flask, FastAPI, REST API, PowerBI, SSMS, Linux

Data Analysis and ML: Pandas, NumPy, Matplotlib, Seaborn, ggplot2, Scikit-learn, TensorFlow, PyTorch, PySpark, Airflow

GenAI: LangChain, RAG, OpenAI API, Anthropic API, pgvector, Pinecone

Cloud and MLOps: AWS (EC2, S3, RDS, Redshift, SageMaker), Azure (Data Factory, Synapse), Docker, MLflow, CI/CD, Git

Databases and Data Platforms : PostgreSQL, MongoDB, Snowflake, Databricks, Supabase

PROJECTS

Ticket Triage Knowledge Base Gap Analysis and Generation

2026

- Deployed production **agentic NLP system** serving **44,000+ users**, autonomously processing **100,000+ IT support tickets** via classification, semantic clustering, and vector search, identifying **120+ recurring issue patterns**.
- Integrated **Azure OpenAI GenAI** summarization layer with **RAG pipeline**, reducing IT triage analyst effort by **43%** and enabling automatic knowledge base article generation for top recurring issues.

Dynamic Pricing Optimization using Machine Learning

2025

- Engineered **price elasticity and demand features** from **500,000+ transaction records**; deployed **XGBoost** price sensitivity model achieving **10.7% revenue lift** through data driven pricing recommendations.
- Improved **demand forecasting accuracy by 18%** over baseline using gradient boosting with custom behavioral features, enabling real time **price optimization** across product tiers.