Mohammed Sanan Moinuddin

(781) 323-9754 | sanan.mohd18@gmail.com | LinkedIn | GitHub | Portfolio

Summary

Computer Science graduate student with strong foundation in machine learning and software engineering. Experience developing ML systems using PyTorch, TensorFlow, and scikit-learn for forecasting and recommendation systems. Proficient in Python, FastAPI, and deploying models on cloud platforms. Seeking Summer 2026 SWE/ML internship.

Education

Northeastern University September 2025 – May 2027
Master of Science in Computer Science | GPA: **3.8/4.0** Boston, MA

Osmania University August 2019 – May 2023
Bachelor of Engineering in Computer Science & Engineering | GPA: **3.5/4.0** Hyderabad, India

Technical Skills

Languages: Python, Java, SQL, C, JavaScript, TypeScript
Development: PyTorch, TensorFlow, Scikit-Learn, Pandas, NumPy, React, Next.js, Node.js, FastAPI
Tools: Git, Docker, Kubernetes, PostgreSQL, MySQL, MongoDB, Redis, Apache Kafka, AWS, Azure DevOps, Jenkins
Concepts: Machine Learning, Deep Learning, Automation Testing, Regression Testing

Experience

Software Engineer August 2023 – August 2025
PepsiCo

- Developed **Python automation framework** using **Selenium WebDriver** with Page Object Model design pattern, **reducing manual testing effort by 90%** and saving 40+ hours weekly across the QA team.
- Designed **Azure DevOps CI/CD pipeline optimization** by parallelizing test execution, **reducing test execution time from 4 hours to 1.5 hours** and enabling faster release cycles.
- Built **RESTful microservices** using **FastAPI** with **PostgreSQL** and **Redis caching** for test data management, **improving test throughput by 65%** across distributed environments.
- Engineered **automated test reporting dashboard** using **React** that aggregated 500+ test scenarios in real-time, **reducing defect triage time by 40%**.
- Led comprehensive regression testing across **500+ scenarios** using data-driven testing methodology, **achieving 99.5% automation coverage with zero critical production defects**.

Projects

Predictive Infrastructure Scaling System Github
Tech Stack: Python, PyTorch, FastAPI, Kafka, PostgreSQL/TimescaleDB, XGBoost, LightGBM, Prophet

- Built **multi-model ML ensemble** combining **PyTorch Transformer**, **XGBoost/LightGBM/CatBoost**, and **Prophet** with **quantile regression** for probabilistic load forecasting 15 minutes to 7 days ahead.
- Engineered **100+ feature pipeline** with cyclical time encodings, rolling statistics, **Fourier components** for seasonality detection, and business event modeling with **prediction calibration** for uncertainty quantification.
- Implemented **real-time streaming** using **Kafka** for metrics ingestion from **Prometheus/Kubernetes**, with async processing and **dead letter queue** for fault tolerance.
- Designed **multi-objective decision engine** with **risk assessment** evaluating cost, stability, and spot interruption probability across **AWS/GCP/Azure** to optimize instance allocation.

Multimodal Short Video Recommender System Github
Tech Stack: Python, PyTorch, HuggingFace Transformers, FastAPI, PostgreSQL, Docker, GCP

- Designed a **two-tower neural network** that combines ResNet-50, Wav2Vec2 and BERT encoders with **gated multimodal fusion** for content-aware modality weighting across video, audio, and text.
- Achieved efficient **transfer learning** by freezing pretrained backbones (ResNet-50, Wav2Vec2, BERT) and training only fusion and projection layers, reducing memory usage and training time.
- Implemented **vector similarity search** using PostgreSQL with **pgvector** cosine distance operators and async database operations for scalable video retrieval.
- Built production-ready **FastAPI** service with batch indexing, **top-K retrieval**, horizontal scaling support, and cross-platform GPU acceleration (**CUDA/MPS/CPU**), containerized with Docker.