ANIRUD MOHAN

College Park, MD | P: +1 7037058375 | amohan26@umd.edu | LinkedIn | GitHub

EDUCATION

University of Maryland | College Park, MD **Expected Graduation Date - 05/26**
Master of Science in Applied Machine Learning **GPA: 3.76 / 4.00**
Coursework: Principles of Machine Learning, Principles of Data Science, Probability and Statistics, Optimization, Algorithms and Data Structures for ML

Misrimal Navajee Munoth Jain Engineering College | Chennai, IN **Jul 2019 – May 2023**
Bachelor of Engineering in Computer Science & Engineering **GPA: 8.82 / 10.00**
Coursework: Database Management System, Artificial Intelligence, Software Engineering, Operating Systems

SKILLS

Languages: Python, R, SQL (MySQL,PostgresSQL), Java
ML/AI: Scikit-learn, TensorFlow, Pytorch ,HuggingFace , LangChain ,LlamaIndex , Guardrails
Data Analytics: Pandas, NumPy, Seaborn, Matplotlib, Plotly, ggplot2, PowerBI, Microsoft Excel, PySpark,Jupyter ,SciPy
Cloud: AWS, GCP, IBM Cloud, DataBricks
Tools: Docker, Git, LangGraph, ClearML, Flask, Gradio

WORK EXPERIENCE

Thapovan Info Systems **Chennai, IND**
Junior Machine Learning Engineer Oct 2023 – Jun 2024

- Developed and Deployed Chatbots: Designed end-to-end LLM-powered chatbot systems using Hugging Face Transformers.
- Integrated Guardrails for LLMs: Improved interaction accuracy by 30% by creating structured input/output validation pipelines and implementing Retrieval-Augmented Generation (RAG).
- Increased accuracy upto 96%: By effectively implementing advanced Retrieval Augmented Generation techniques.
- Collaborated Cross-Functionally: Used Jira and Agile to manage 3+ project tasks, ensuring timely delivery and team collaboration

Azentio Software **Chennai, IN**
Software Developer Intern Feb 2023 – Sept 2023

- Architecture Optimization: Contributed to refactoring the monolithic architecture into a multi-server structure, improving system scalability and performance.
- API Development: Separated UI and API into two projects, improving security and scalability by exposing only the logistics API.
- Innovated methods using Python libraries like Scrapy and BeautifulSoup4 for high-volume dataset preparation, resulting in findings that rectified three critical bottlenecks during system performance evaluations.

COURSEWORK & CAPSTONE PROJECTS

Document Classification using SVM (Numpy, Pandas, Python, Sklearn) Aug 2024 – Dec 2024

- Developed an end-to-end document classification pipeline, implementing SVM from scratch and optimizing TF-IDF feature extraction for high-dimensional text data.
- Engineered an optimized sparse matrix representation, reducing memory usage significantly and accelerating computational speed by 3x in large-scale text processing tasks.
- Implemented 5-fold cross-validation and hyperparameter tuning, achieving 90% accuracy on the 20 Newsgroups dataset through feature scaling and normalization.

Early Diagnosis of Alzheimer's Disease (Python, TensorFlow, PyTorch, Streamlit, Kaggle, ROBEX) Jan 2023 – May 2023

- Designed and implemented deep learning architectures, including CNNs and transformers, to analyze MRI scans for early Alzheimer's detection, optimizing model performance and scalability.
- Engineered an end-to-end pipeline for data preprocessing, model training, and real-time inference using TensorFlow and PyTorch, improving computational efficiency.
- Leveraged ensemble learning techniques and hyperparameter tuning to enhance model robustness, achieving over 97% accuracy on the ADNI dataset.

University Admit Eligibility Predictor (Python, Flask, Scikit-Learn, HTML, CSS, Docker) Aug 2022 – Dec 2022

- Designed and deployed a machine learning-powered web application that processed 1,000+ student profiles, achieving 95%+ prediction accuracy on university admission chances, with a Flask-integrated backend ensuring seamless real-time inference.
- Constructed an efficient prediction pipeline utilizing linear regression algorithms with optimized hyperparameters, leading to improved accuracy and achieving a consistent inference time of less than two seconds across varying datasets.
- Containerized the application using Docker, ensuring efficient deployment and cross-platform accessibility, while optimizing model inference time for real-time predictions.