

NEERAJ

New Delhi, Delhi, India

+91-9871520187

neerajhhhh225@gmail.com

Linkedin

Github

Portfolio

SUMMARY

Machine Learning & MLOps Practitioner with experience in building scalable AI solutions. Skilled in deploying models using AWS, Docker, and CI/CD pipelines. Passionate about automating ML workflows and optimizing model performance for real-world applications. Experienced in ML model orchestration, experiment tracking, and cloud computing.

SKILLS

Languages: Python, SQL (PostgreSQL)

Data Science: A/B Testing, Hypothesis Testing, Data Science Pipeline (Cleaning, Wrangling, Visualization, Modeling, Interpretation), EDA, Statistics

Machine Learning: Python (e.g. scikit-learn, numpy, pandas, matplotlib), Machine learning Pipeline (Data Ingestion, Transformation, Validation, Model Training and Evaluation), ETL, Feature Engineering

MLOps: Apache Airflow, MLflow, CI/CD (GitHub Actions), Docker, FastAPI, Flask, Grafana

Cloud Technologies: AWS (EC2, S3, ECR, Beanstalk, SageMaker), Azure (Web App, Container Registry)

Tools/Technologies: Git, GitHub, Jupyter Notebook, Zenml, VS Code, APIs, Streamlit, DVC

PROJECTS

Malicious URL Detection System

OCT 2024 - Dec 2024

- Developed an end-to-end MLOps solution for malicious URL detection, achieving 97.6% F1-score.
- Tech Stack: Python, MongoDB, MLflow, Airflow, FastAPI, Docker, GitHub Actions, AWS (S3, ECR, EC2)
- Built a FastAPI backend for model training and batch predictions, reducing inference time by 35%.
- Implemented an automated model training pipeline and prediction pipeline using Apache Airflow, scheduled weekly to ensure regular model retraining.
- Tracked experiments using MLflow, hosted on AWS EC2, with artifacts stored on AWS S3.
- Performed Data validation by ensuring schema correctness, detecting data drift, & generating drift reports.
- Deployed the application on AWS EC2 using Docker containerization and the Docker image stored in ECR.
- Implemented CI/CD pipeline using GitHub Actions for automated deployment of application to AWS EC2, reducing deployment time by more than 30%.

Machine Predictive Maintenance Classification

JAN 2025 - FEB 2025

- Developed an end-to-end MLOps solution to predict machine failures, achieving an F1-score of 96.4%.
- Tech Stack: Python, Flask, Docker, GitHub Actions, AWS (Beanstalk, EC2, CodePipeline), Azure, Xgboost
- Evaluated ML models to predict machine failure types, selecting the best model based on F1-score.
- Optimized model performance through hyperparameter tuning using GridSearchCV.
- Deployed the application on AWS EC2, AWS Elastic Beanstalk, and Azure web app, ensuring scalability, reliability, and 99.9% uptime in production.

Customer Satisfaction Prediction

SEP 2024 - SEP 2024

- Built an end-to-end ml pipeline for continuously deploying the machine learning model and prediction, alongside a data application that utilizes the latest deployed model for the business to consume.
- Tech Stack: Python, ML Pipeline Orchestration (ZenML), Model Deployment (MLflow), Streamlit
- Built a training pipeline using ZenML and MLflow for experiment tracking and model deployment.

EDUCATION

Netaji Subhash University of Technology, New Delhi, DELHI

2020 – 2022

Completed Coursework: Calculus, Intro to Computer Programming, Algebra, Probability Theory.

ACHIEVEMENTS

- Created and shared technical content on AI, ML, and MLOps, building an engaged community of 7K+ tech professionals and practitioners in 8 months, leading to Recognition from established professionals in the field.