

NIHAR GUDADHE

Data Scientist | AI ML Engineer

Nagpur, India | +91- 9168844536 | gudadhenihar@gmail.com | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

SUMMARY

Data Scientist & Analyst with 2+ years of experience across the full machine learning lifecycle - EDA, feature engineering, model development, and deployment - for regression, classification, and time-series problems. Proficient in Python and SQL with a strong foundation in statistics. Productionized a Random Forest pricing model that cut quotation turnaround from days to seconds, and built automated CRM/ERP data pipelines to streamline the Quote-to-Cash process. Actively building skills in Generative AI (LLMs, prompt engineering, RAG) alongside core ML and Power BI analytics work.

TECHNICAL SKILLS

Programming & Databases: Python, SQL, PostgreSQL, MySQL, Excel

Machine Learning & Statistics: Regression, Classification, Clustering, Time Series Forecasting, Ensemble Methods (Random Forest, XGBoost, LightGBM), PCA, Anomaly Detection, Feature Engineering & Selection, Cross-Validation, Hyperparameter Tuning, SMOTE, Hypothesis Testing, A/B Testing

Deep Learning & NLP: ANN, CNN, RNN, LSTM, GRU, Encoder-Decoder, Attention Mechanism, Tokenization, Word Embeddings, Named Entity Recognition, Sentiment Analysis

Generative AI: Large Language Models (LLMs), Prompt Engineering, RAG, Hugging Face Transformers, LangChain

Computer Vision: Image Classification, Object Detection (YOLO), Image Segmentation, OpenCV

Libraries & Frameworks: Pandas, NumPy, Scikit-learn, TensorFlow, PyTorch, Matplotlib, Seaborn, NLTK, spaCy, PySpark, Databricks

Data Visualization: Power BI, DAX, Power BI Service, Tableau, Excel

Engineering & Deployment: FastAPI, Docker, Git, GitHub, CI/CD, AWS, Render, Jupyter Notebook, VS Code

Core Concepts: Statistical Inference, Confidence Intervals, EDA, Data Cleaning, Data Modeling, Data Warehousing, Data Labeling

PROFESSIONAL EXPERIENCE

Data Scientist & Analyst | Influx Control Panel, Nagpur Feb 2024 - Present

- Enterprise Data Integration & Warehousing:** Architected and maintained an automated ELT pipeline extracting raw CRM leads and billing records into a MySQL staging lake, transforming data via dbt, and housing clean analytical tables in a PostgreSQL data warehouse.
- Predictive Pricing Engine:** Developed and productionized a custom Random Forest Regression model using Scikit-Learn to automate cost estimations for complex industrial items, containerized behind a real-time FastAPI endpoint, cutting quotation turnaround times from days to seconds.
- Operational Efficiency:** Developed a Production Line Efficiency Dashboard that tracked cycle times and machine idle time, identifying bottlenecks that led to a **12%** increase in shop-floor throughput.
- Sales Performance Analytics:** Designed a 360-degree customer dashboard in Power BI that empowered the sales team to prioritize high-value segments, contributing to a **10%** improvement in sales conversion rates.
- Quality Assurance:** Applied anomaly detection techniques in an automated defect-analysis report that correlated raw material batches with support tickets, helping the engineering team reduce product return rates by **14%**.

PROJECTS

Scipe - Enterprise Analytics & ML Pricing Platform (Internal Project)

Python | Scikit-learn | PostgreSQL | MySQL | dbt | FastAPI | Docker | Power BI

- Engineered automated ETL scripts to sync CRM webhooks and Tally transaction logs into a structured relational warehouse schema.
- Built feature-engineered models evaluating non-linear component costs, switchgear configurations, and historical win-loss ratios to standardize pricing.
- Deployed model inference services via FastAPI inside multi-container Docker environments, orchestrated with GitHub Actions CI/CD.
- Developed SQL-driven Power BI dashboards to monitor industrial pipeline risk and throughput metrics.

Shoppers Intent - MLOps Classification Pipeline (Independent Project)

Python | FastAPI | MySQL | PostgreSQL | Scikit-learn | XGBoost | LightGBM | MLflow | Evidently AI | Docker | Render

- Built an end-to-end MLOps system predicting online shopper purchase intent on a **12,330-row**, 17-feature dataset, auto-selecting the best of Random Forest, XGBoost, and LightGBM (tuned via GridSearchCV, SMOTE-balanced) by ROC AUC, achieving **91.8%** accuracy and **97.3%** ROC AUC.
- Developed an interactive web frontend and FastAPI backend (with Swagger docs) for live session risk scoring, achieving a **92.1%** F1 score in production.
- Implemented MLflow experiment tracking and model registry, automated daily model retraining via APScheduler, and containerized deployment on Render with Evidently AI drift monitoring - validated end-to-end with a Pytest suite and GitHub Actions CI/CD.

OpinionMetrix - AI Sentiment Transformer (Independent Project)

Python | NLP | Scikit-learn | Streamlit | MySQL | Git | GitHub

- Developed an end-to-end NLP sentiment-analysis web app that classifies text as Positive, Negative, or Neutral with confidence scoring.
- Built a preprocessing pipeline (HTML/URL/emoji stripping, tokenization, stopword filtering, lemmatization) paired with TF-IDF and an ML classifier, with a MySQL backend logging inputs, predictions, and confidence metrics for auditing.
- Deployed the app on Streamlit Community Cloud with secure secrets management for database credentials.

EDUCATION

B.E. in Computer Science Engineering

Sipna College of Engineering, Amravati

CERTIFICATIONS

- Diploma in Machine Learning using Python - Alison
- Data Science Essentials with Python - Cisco Networking Academy