

Mounusha Ram Metti

+1 (623) 272-8817 · mounushametti12@gmail.com · linkedin.com/in/mounusha-ram-metti · github.com/Mounusha25 · mounushametti.com

TECHNICAL SKILLS

Languages: Python, SQL, R, TypeScript, JavaScript

Machine Learning & Modeling: scikit-learn, XGBoost, LSTM, Isolation Forest, Random Forest, Regression, Time-Series Forecasting

Statistics & Experimentation: Statistical Modeling, ANOVA, A/B Testing, Correlation Analysis, Feature Engineering, Model Evaluation

Generative AI & Tools: RAG, LangChain, FAISS, MLflow, Tableau, Excel, Pandas, NumPy, Git

Databases & Cloud: PostgreSQL, BigQuery, TimescaleDB, AWS RDS, AWS S3

PROFESSIONAL EXPERIENCE

Arizona State University (SSEBE) · Data Science Researcher

Oct 2025 - Present | Tempe, AZ

- Analyzed **1.58M+ sensor records** with correlation, multi-feature regression, ANOVA, and seasonal decomposition, finding intake humidity drives harvest rate (r of about 0.65), explains about **75%** of variance, and station gaps reach **15%**, with monsoon yield up **33%**.
- Benchmarked **5 model families**, including LSTM, Isolation Forest, and Random Forest, for anomaly attribution across a **14-round** process on 645 labeled fault events, reaching **0.80 F1**, 0.85 recall, 0.75 precision, and under 0.1% FPR on 7,352 held-out windows.
- Gathered requirements directly from **SRP**, a real external client, and delivered a Next.js and FastAPI monitoring dashboard presented to SRP stakeholders, now in active production use by **10+** users and supporting **4,830+ liters** of harvested water across 9 stations.

Arizona State University (SSEBE) · Data Scientist

Feb 2025 - Oct 2025 | Tempe, AZ

- Modeled river flow with **sklearn GradientBoosting** regression at **0.87 R2**, extending analysis from gauged rivers to 337,000 ungauged reaches using elevation-gradient, network-centrality, and flow-accumulation features engineered from federal datasets.
- Automated an **ELT pipeline** integrating CWNS 2022 and NHDPlus federal datasets with schema validation, null checks, and bad-record rejection, cutting data-processing time **60%** and delivering model-ready tables for wastewater risk analysis to **3** research groups.
- Independently optimized **memoized BFS** over a 1.07M-node river DAG for 1,100 intake points, caching overlapping downstream subpaths and cutting traversal time from **95 seconds to 12 seconds**, an **8x** speedup, so repeated national routing queries became practical for research use.

AWL Metaverse · Machine Learning and Analytics Software Engineer Intern

May 2024 - Dec 2024 | India

- Benchmarked logistic regression as an interpretable baseline, then trained **XGBoost** on **500K+** user activity records and served the better classifier through a FastAPI REST endpoint, improving retention prediction performance **18%** over rule-based scoring.
- Deployed a **LangChain** and OpenAI embeddings **RAG** assistant with Chroma and FAISS retrieval inside the course interface, letting students self-serve answers in free text and reducing dashboard-tracked support ticket volume about **30%** across the learning platform.

Dharte Inc. · Marketing and Data Analytics Intern

May 2023 - Aug 2023 | India

- Designed 2 website variants and ran an **A/B test** across **3K+ visitors**, structuring the comparison to isolate UX effects from audience targeting and measuring a **15% relative** CTR lift that informed the UX decision for marketing and product stakeholders.
- Applied **behavioral segmentation** across **25K+** customer and campaign interactions in SQL and Python to isolate 4 high-value customer groups, then measured a **7%** increase in targeted cohort interaction rate following the segmentation-driven initiative.

PROJECTS

Retail Demand Forecasting Engine | Python · BigQuery ML ARIMA_PLUS · Apache Beam · PySpark · FastAPI · Tableau

- Benchmarked a **BigQuery ML ARIMA_PLUS** time-series forecasting model against 3 baselines on a real 30-day holdout across 1,115 Rossmann stores, reaching **16.92 MAPE** versus 21.91-24.74 MAPE for seasonal-naive and moving-average forecast alternatives.
- Added **80% confidence intervals** to 33,450 30-day store forecasts, making demand uncertainty visible for staffing and replenishment decisions, and structured **4 Tableau-ready reporting views** so analysts could explore forecast and baseline results.

Decision-Optimized Churn Prediction | Python · Decision Engine · Budget-Constrained Optimization · Profit-Curve Analysis

- Modeled campaign economics by ranking customers on expected net value under a hard **\$50,000 monthly budget**, encoding a \$300 false-positive penalty and \$50 campaign cost to set a 0.71 breakeven probability, with **0.9068 ROC-AUC** on point-in-time-correct features.
- Optimized a budget-constrained plan targeting the top 659 customers, about 0.3% of the base, for about **105% ROI** on \$52,000 expected monthly net value, capturing **96.4%** of optimal profit, and used profit-curve analysis to locate maximum profit near \$1.24M at 25,100 customers.

EDUCATION

Arizona State University — M.S. Data Science | **GPA 4.0/4.0**

Dec 2026 expected

IIT Kharagpur — B.Tech Civil Eng. (integrated M.Tech Environmental Eng.) | **GPA 3.8/4.0**

May 2024

ACHIEVEMENTS

- 3x AI Hackathon Winner** - placed top 1 of 10,000+ participants across carbon-aware AI governance and distributed agent network events.
- IBM Data Science Professional Certificate** - coursework covering statistical modeling, regression, and applied machine learning methods.
- Passed Claude 101 (Anthropic) certification** - trained in prompt engineering, agents, and responsible AI use in production workflows.