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, TypeScript, JavaScript

Machine Learning & Modeling: PyTorch, scikit-learn, XGBoost, LSTM, Anomaly Detection, Feature Engineering, Model Evaluation

Generative AI & NLP: LangChain, RAG, FAISS, ChromaDB, Embeddings, LLMs, Prompt Engineering, Function Calling

MLOps & Backend: MLflow, FastAPI, Docker, GitHub Actions, Prometheus, Grafana, Git, AWS EC2

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

PROFESSIONAL EXPERIENCE

Arizona State University (SSEBE) · *Machine Learning Engineer* Oct 2025 - Present | Tempe, AZ

- 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.
- Tracked **132** training runs with **MLflow** and narrowed modeling scope from 10 sensor features to **4** core parameters after exposing frozen-sensor contamination and an 8x station-coverage imbalance, roughly doubling model scores and surfacing the winning 0.80 F1 configuration.
- 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.

AWL Metaverse · *Machine Learning and Analytics Software Engineer Intern* May 2024 - Dec 2024 | India

- 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.
- 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.
- Logged chatbot conversations in **PostgreSQL** and analyzed customer question patterns with SQL and Python queries, surfacing **50+ recurring content gaps** in support content and frequent student-question themes for a 5,000+ user 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.
- Automated weekly stakeholder reporting with **Python, SQL, and Excel Power Query**, consolidating fragmented sources and cutting turnaround from **3-4 hours to about 1 hour** while improving output consistency across marketing campaign performance metrics.

PROJECTS

TrustMed-AI | *Python · FastAPI · ChromaDB · FAISS · Cross-Encoder Reranking · Next.js*

- Indexed **9K+ medical chunks** into **3 ChromaDB collections** with domain-routed retrieval, added FAISS candidate refinement and cross-encoder reranking, and exposed a **3-step ReAct** trace with source cards, confidence, and response time so every answer is auditable.
- Built a **100-prompt** evaluation set across **5 medical domains** and **3 benchmark scripts** measuring **6** retrieval and grounding signals, including Precision@K, MRR, and grounded score, served through 5 FastAPI endpoints and 2 ElevenLabs voice endpoints.

StyleNova-AI | *Python · PyTorch · CLIP · scikit-learn · FastAPI · Next.js*

- Built a hybrid recommender combining **CLIP** ViT-B-32 512-dimensional visual embeddings, TF-IDF content similarity, and collaborative kNN with **3** adaptive cold-start weighting stages, reaching **0.73 Precision at 10**, 0.68 diversity, and **89%** catalog coverage.
- Precomputed **512-dimensional** product embeddings at ingest, cutting recommendation latency from about 2 seconds to under **200 ms**, and added GPT-4o-mini function calling for preference edits that cut swipes needed to escape a mismatched cluster from **12 to 4**.

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.