

Asmita Gupta

Irvine, CA | (248) 595-9462 | asmitagpt1109@outlook.com | [LinkedIn](#) | [GitHub](#)

Education

MS BUSINESS ANALYTICS | GPA: 3.74 | MAY 2026 | UNIVERSITY OF MASSACHUSETTS AMHERST

Coursework: Data Science, Machine Learning, Database Systems, AI for Business

BE COMPUTER SCIENCE | GPA: 3.2 | MAY 2022 | PUNE UNIVERSITY

Coursework: Data Structures & Algorithms, Database Management

Skills & Abilities

- Python, Java
- SQL, PostgreSQL, ETL Pipelines, Snowflake, MLflow
- PyTorch, TensorFlow, HuggingFace Transformers, XGBoost, Regression, Classification, Recommendation Systems, Forecasting, Reinforcement Learning
- Tableau, Metabase, Streamlit
- Amazon Web Services: S3, SageMaker, Lambda, EC2
- Git, GitHub, Jira, Docker, Claude Code
- LLMs, LangGraph, Prompt Engineering, RAG
- Anthropic Claude API, FastAPI, REST APIs

Experience

SOFTWARE ENGINEER (AI/ML) INTERN | AMOTIONS INC. | SEP 2026 | SAN FRANCISCO, CA

- Architected production **SaaS** applications with **NLP/LLM integrations (OpenAI, Gemini/Vertex AI, Claude)**, including **provider-abstraction** and **fail-safe response validation**.
- Built real-time **inference pipelines** combining **speech-to-text (Deepgram)** with LLM-based analysis for live **emotional-cue detection** and in-call coaching guidance.
- Owned end-to-end delivery of AI surfaces: real-time guidance, AI roleplay, seller assessments, multimodal post-call analytics - from model inference to customer-facing delivery.

SOFTWARE ENGINEER | APPDIRECT PVT LTD | JUL 2022 – DEC 2023 | PUNE, INDIA

- Resolved 8+ monthly transactions and invoice discrepancies threatening financial compliance by tracing root causes through data flows across a large-scale **Java** monolith and **Spring Boot** services.
- Engineered **ETL pipelines** across 2+ TB of e-commerce data, cutting query execution time by 20% through index and join refactoring; automated **data quality validation**, catching 12 critical issues monthly before reaching production.
- Shipped **REST API** integrations automating data ingestion from external service feeds into the core billing platform using **Java** and **Spring Boot**, eliminating manual reconciliation steps and improving data freshness for downstream analytical workflows.

Projects

CINEIQ – BEYOND STATIC RECOMMENDATION | PYTHON, LIGHTGBM, PYTORCH, ANTHROPIC CLAUDE API, AWS SAGEMAKER, MLFLOW, STREAMLIT | 2026

- Trained an **NCF recommender** with **BPR** ranking loss on **33.7M** interactions using **AWS SageMaker**, achieving **NDCG@10 = 0.80**, with automated **S3** artifact sync.
- Built a **LightGBM demand forecaster** (**82% RMSE** improvement over a naive baseline) and a **LinUCB** contextual bandit for exploration–exploitation; integrated the **Claude API** for real-time personalized explanations.
- Orchestrated the **MLOps pipeline**: **S3** artifact versioning, **CloudWatch** monitoring and one-command reproducible **SageMaker** training across 4 components.

DISASTERSENSE – MULTIMODAL DISASTER SEVERITY DETECTION | PYTHON, PYTORCH, HUGGINGFACE, FASTAPI, POSTGRESQL, DOCKER | 2026

- Fine-tuned **EfficientNet-B0** (**64% accuracy, up from a 58% baseline**) and **Twitter-RoBERTa** (75% accuracy) on CrisisMMD, handling 4.65x class imbalance with **weighted loss** and differential learning rates; validated per class with precision/recall.
- Deployed the multimodal severity-scoring system end to end: **FastAPI** at **~200 ms** CPU latency, **Docker**, a public **Hugging Face Spaces** demo, and live **PostgreSQL**-backed monitoring in **Metabase**.

SPEECH EMOTION RECOGNITION | PYTHON, PYTORCH, HUBERT, BILSTM, TORCHAUDIO, LIBROSA | 2025

- Extracted **HuBERT** embeddings from **7,462 CREMA-D clips** and trained a **BiLSTM-attention classifier**, reaching **71.17% accuracy** and **0.72 macro F1** across 6 emotions, 4x the random baseline.
- Built an audio pipeline from raw waveforms through **augmentation (pitch shift, time stretch, Gaussian noise)**, **feature extraction** and per-class evaluation.