Twisha Patel

Vadodara, Gujarat twishap534@gmail.com +91 8160893358 twishapatel12.github.io
Twisha Patel twishapatel12

Summary

Software Engineer (AI/ML) with hands-on experience in computer vision, NLP/LLMs, and speech-to-text. Builds end-to-end systems from data ingestion and training to deployment and monitoring. Strong in Python and open-source ML stacks (PyTorch, TensorFlow, OpenCV, Hugging Face), with experience contributing production-grade RAG pipelines to large open-source projects.

Education

Indian Institute of Technology, Madras *Dec 2025 – Jun 2026*
Advanced PG in Agentic AI Workflows and Agentic Systems Development

ITM SLS Baroda University, Vadodara, Gujarat *Aug 2021 – May 2025*
Bachelor's Degree in Computer Science Engineering — CGPA: 8.6/10

Shri J.R Shah Bright English Medium School, Vadodara, Gujarat *June 2020 – July 2021*
Higher Secondary School — Grade: B2

Experience

Associate AI/ML Engineer *July 2025 – Present*
Sundaram Technologies, Vadodara, India

- Built and deployed computer vision models for industrial automation, improving detection accuracy **F1** and reducing preprocessing latency by 35%.
- Productionized models with TensorFlow, PyTorch, and OpenCV, and containerized services with Docker/CI/CD to shorten release cycles from 2 weeks to 3 days.

AI/ML Intern *Jan 2025 – July 2025*
Artem Health Tech, Ahmedabad, India

- Developed a machine-integration service to sync device data to the company portal, cutting manual uploads by **40%** and improving data freshness from **daily to hourly**
- Implemented real-time object detection with MMDetection (RTMDet-tiny) on webcam streams at **25–30 FPS** with sub-**50 ms** per-frame preprocessing.
- Built a speech-to-text pipeline leveraging open-source models with agentic **behavior**; achieved **12–15%** WER on internal evaluation data.
- Fine-tuned a finance-focused LLM (Flan-T5) on curated datasets; improved exact-match accuracy for finance Q&A by **18%** on held-out tests.

Skills

Programming Languages: Python, JavaScript, PHP, SQL, Bash

Web Technologies: HTML5, CSS3, Bootstrap, jQuery, REST APIs

Databases: MySQL, SQLite

Frameworks/Libraries: TensorFlow, PyTorch, Keras, Scikit-learn, OpenCV, Hugging Face, Transformers, Pandas, NumPy, Matplotlib, Seaborn, Streamlit, FastAPI

Machine Learning/AI: Deep Learning, NLP, Computer Vision, Transfer Learning, Model Explainability (GradCAM, SHAP), Hyperparameter Tuning, Human-in-the-Loop AI

Large Language Models: Fine-tuning, Prompt Engineering, Model Deployment, Custom Q&A Solutions

Tools/Platforms: Git, Docker, GitHub Actions, CI/CD Pipeline, VS Code, Jupyter Notebooks, Google Colab

Other: Data Visualization, Exploratory Data Analysis, Report Generation (CSV, PDF, JSON), Batch Processing, API Integration, Research Documentation

Projects

Pneumonia Detector AI

- Built an interactive, explainable AI system for pneumonia detection in chest radiographs using EfficientNet, PyTorch, and Streamlit, with GradCAM visualizations for model transparency.
- Engineered a robust human-in-the-loop workflow enabling real-time user feedback, **orchestrated** retraining, image-quality checks, and audit/session logging to drive continuous improvement.
- Delivered a privacy-first solution with batch analysis and downloadable reports (CSV, PDF, JSON, FHIR), packaged with Docker for seamless deployment.

AutoML Pipeline Service

- **Architected** and deployed an end-to-end AutoML service (Streamlit + FastAPI) for dataset profiling, model selection, training/inference, versioning, and performance logging.
- **Systematized** dataset profiling and experiment tracking with configurable retention and reproducible benchmarks.

Object Detection with MMDetection

- Implemented a real-time object detection system using MMDetection's RTMDet-tiny and OpenCV on webcam video.
- Rendered live bounding boxes and class labels using **pretrained** COCO weights.
- Gained hands-on experience with MMDetection's API, model inference workflow, and visual result rendering.

Traffic Sign Recognition

- Developed and trained a deep learning model in Python to classify German traffic signs using the GTSRB dataset.
- Used OpenCV for advanced image preprocessing and data augmentation to enhance robustness.
- **Built** and fine-tuned a convolutional neural network (CNN) with TensorFlow/Keras, achieving high test accuracy.

Voice Assistant

- Developed a desktop voice assistant in Python that executes voice commands (reminders, music, web search).
- Integrated speech recognition and TTS libraries for seamless two-way interaction.
- Enhanced natural language understanding for more accurate and flexible command processing.

Finance Based Small LLM

- Fine-tuned Flan-T5 on curated financial datasets for finance-specific Q&A.
- Applied prompt engineering and evaluation to improve contextual relevance and precision.

Open Source Contributions

LlamaIndex — Open Source Contributor *2026*

- Contributed a production-ready **FastAPI-based Retrieval-Augmented Generation (RAG)** example using LlamaIndex with **fully local Ollama models**.
- Designed a document ingestion and query pipeline leveraging vector indexing and local embeddings, enabling **offline, no-API-key** RAG workflows.
- Integrated FastAPI endpoints for querying indexed documents and validated the example via automated CI, linting, and maintainer review.
- Collaborated with maintainers to resolve architecture, embedding-model selection, and formatting feedback; contribution merged into the main repository.