

Vansh Patel

Ahmedabad, Gujarat, India

+91 6358766799 | patelvanshj.17@gmail.com | linkedin.com/in/vansh-jpatel | github.com/vanshpatel-017

Summary

Third-year Computer Science undergraduate specializing in **Machine Learning, Deep Learning, and Computer Vision Systems**. Experienced in building **production-grade AI pipelines** involving real-time inference, multi-object tracking, transformer architectures, and scalable deployment systems. Strong focus on **computer vision, industrial AI, MLOps, and end-to-end ML system engineering**.

Education

Institute of Infrastructure Technology Research and Management

Sep 2023 – May 2027

B.Tech in Computer Science (CPI: 7.6)

Ahmedabad, India

Technical Skills

Programming: Python, C++, SQL, Java, Bash

Techniques: Feature Engineering, PCA, SMOTE

ML/DL: PyTorch, TensorFlow, Scikit-Learn, XGBoost

MLOps: Docker, AWS, FastAPI, CI/CD

Computer Vision: YOLO11, OpenCV, ViTs, ByteTrack

Tools: Git, Linux, REST APIs

Experience

Machine Learning Engineer Intern

Sep 2025 – Oct 2025

Codveda Technologies (Remote)

- Led development of a **production-ready churn prediction pipeline**, improving model accuracy from 80% to **95%** using XGBoost and hyperparameter optimization.
- Designed robust data preprocessing workflows using **SMOTE, feature engineering, and scaling**, achieving **ROC-AUC of 0.92**.
- Implemented **cross-validation and performance evaluation pipelines** to ensure model generalization.
- Applied **K-Means clustering with PCA** for large-scale user segmentation, retaining 95.8% variance.
- Collaborated in a distributed team to deliver scalable, production-grade ML solutions.

Projects

LeafCure – Multi-Crop Disease Detection System

GitHub

PyTorch, Vision Transformer, FastAPI, React, Firebase

- Built a **production-grade end-to-end computer vision system** for real-world plant disease detection across **38+ classes** using Vision Transformers.
- Designed and deployed scalable **REST API-based inference services** using FastAPI for real-time predictions.
- Optimized **model loading and inference latency** for reliable deployment in cloud-constrained environments.
- Engineered full-stack platform with analytics, user management, and monitoring for continuous usage.

MindScan – AI Mental Health Detection | *Flask, CNN-LSTM, WebSockets*

Infotsav '24 Finalist

- Developed a **CNN-LSTM based pipeline** for ECG signal classification, achieving **97.2% precision (0.96 recall)**.
- Built **real-time monitoring infrastructure** with WebSockets for continuous health tracking and alerting.
- Optimized model architecture for **edge deployment (Raspberry Pi)** enabling low-latency inference.

ForgeGuard – Industrial Safety Intelligence System | *PyTorch, YOLO11, ByteTrack, OpenCV* GitHub

- Developed a **real-time industrial safety monitoring system** using **YOLO11 and ByteTrack** for PPE detection, worker tracking, and hazard-zone analysis in construction environments.
- Engineered a **multi-stage computer vision pipeline** integrating compliance intelligence, temporal safety monitoring, and persistent multi-object tracking for scene-aware surveillance.
- Built a **modular production-style architecture** with real-time inference pipelines, analytics visualization, and centralized configuration management.

Achievements & Certifications

- Top 5% on Kaggle (Top 150/3000+ teams)
- National Finalist – Infotsav '24 (Rank 4)
- Research paper on Hierarchical Attention for deep learning submitted to IEEE Access
- IBM, DeepLearning.AI, Deloitte certifications