

Varun Gupta

9313228834 | varunygupta123@gmail.com | linkedin | github

EDUCATION

DJ Sanghvi College of Engineering

B.E. in Artificial Intelligence & Data Science; GPA: **8.80/10.0**

Mumbai, India

Aug 2023 – June 2027

Vedant International School

CBSE Class XII: **86.8%** — Class X: **84.67%**

Ahmedabad, India

2021 – 2023

TECHNICAL SKILLS

Languages: Python, SQL (PostgreSQL, MySQL, DuckDB)

AI/ML & GenAI: PyTorch, TensorFlow Lite, scikit-learn, Transformers, Hugging Face, FAISS, LangChain, LangGraph, LLM Orchestration

Backend & Infra: FastAPI, Docker, Redis, PostgreSQL, Git, Streamlit, React

Specializations: Agentic AI, Multi-Agent Systems, RAG Systems, Responsible AI, Anomaly Detection, NLP, Financial Data Engineering, Edge ML

PROFESSIONAL EXPERIENCE

Python Developer Intern

March 2026 – Present

GJ Map Solutions

Remote

- Working on geospatial ML pipelines – fine-tuned a U-Net model on satellite imagery for land-use segmentation and connected the outputs to ArcGIS and QGIS for visualization and downstream analysis

Python Developer Intern

June 2024 – August 2024

VaultofCodes

Remote

- Fixed bugs in production Python codebases and built automation scripts and utility tools during team sprints, helping improve code quality and reliability across the project

TECHNICAL PROJECTS

Epistemic Debate Engine – Multi-Agent Financial Analysis System

Agentic AI | Financial Markets

Python, FastAPI, LangGraph, Groq, Gemini, Ollama, PostgreSQL

- Built a financial analysis system with **6 AI agents**, each representing a different investment philosophy (EMH, Behavioral Economics, Reflexivity, Fundamental Analysis, Adaptive Markets, Global Macro); agents debate a given stock or asset by generating investment theses and challenging each other's reasoning in real time over WebSockets
- Set up an LLM routing layer across Groq (Llama 3.3 70B), Gemini (Gemma 2 9B acting as a debate judge), and a local Ollama fallback; the judge rates each argument on logical strength from 0 to 1 and updates each agent's credibility score live in PostgreSQL as the debate runs
- A Synthesis Agent reads the full debate weighted by each agent's final credibility and outputs a **BUY/SELL/HOLD** call with confidence level, price targets, key risks, and conditions that would reverse the thesis

InvestGuard – Agentic Investment Data Quality Intelligence System

Agentic AI | Financial Domain

Python, DuckDB, PyTorch, FAISS, scikit-learn, Streamlit, LangChain

- Developed an anomaly detection system for investment data across **6 financial domains** – using a layered approach of rule-based checks, statistical outlier detection (Z-score/IQR), and ML models (Isolation Forest + LSTM Autoencoder) – hitting **F1 > 0.78** on injected test anomalies across securities, trades, positions, prices, performance, and client records
- Built an agentic triage layer with **4 tools** – anomaly classification, RAG-based SOP lookup via FAISS, team routing using RACI mappings, and LLM-generated summaries – that auto-creates **exception reports** and sends them to the right domain owners
- Created a Streamlit dashboard showing per-domain data health KPIs and anomaly trends; added TF-IDF based reconciliation to match and resolve record conflicts across noisy, multi-source security identifiers

On-Device KYC Pipeline – Privacy-Preserving Identity Verification

Edge ML | Banking Compliance

Python, TFLite, EfficientNetB0, MobileNetV3, MobileFaceNet, Kotlin

- Combined three deep learning models into a single on-device KYC pipeline – EfficientNetB0 for ID document classification, MobileNetV3-Small for **liveness detection** (trained on CelebA-Spoof using Focal Loss to handle class imbalance), and MobileFaceNet for face matching using ArcFace loss with 512-dimensional embeddings
- Quantized the full pipeline to INT8, cutting model size by **75%** with under 1% accuracy drop; deployed as a TFLite model on Android (Kotlin) where all inference runs on the device – **no biometric data is ever sent to a server**, making it suitable for bank-grade KYC verification

POSITIONS OF RESPONSIBILITY

Training & Placement Coordinator | DJ Sanghvi College of Engineering

August 2025 – Present

- Handle the full placement cycle for 800+ students – coordinating with companies on schedules, shortlisting criteria, and slot allocation, while managing student communications across both technical and non-technical drives

Vice Chairperson | DJS NSDC

May 2024 – May 2026

- Organized nationwide hackathons and skill-development events; managed chapter operations including student programs, workshops, and partnerships with external organizations under NSDC's mandate

Data Science Mentor | DJS Infomatrix

August 2025 – May 2026

- Taught second-year students data science fundamentals including geospatial analysis, time series forecasting, Power BI dashboard building, and dimensionality reduction; guided their project work and helped them prepare for competitions

ACHIEVEMENTS & CERTIFICATIONS

4x Hackathon Winner – HackBuild, Tech Expo 2025, 2026

Advanced Machine Learning – Kaggle | **AWS Academy Cloud Foundations** – cloud computing and scalable infrastructure

500+ Rating on CodeForces – solved 50+ algorithmic problems, demonstrating strong problem-solving skills