Kartik Soni

Bhilwara, Rajasthan, India | kartikksoni0@gmail.com
LinkedIn | GitHub

Summary

A 3rd-year B.Tech AI&ML student at NMIMS MPSTME with hands-on experience building physics-informed deep learning models, multimodal RAG systems, and spatio-temporal forecasting pipelines. Demonstrated ability to work at the intersection of applied ML and real-world problem domains - from government smart city initiatives to oceanographic data interfaces. Actively seeking research internships where I can contribute to and learn from production-grade AI infrastructure.

Education

SVKM NMIMS Mukesh Patel School of Technology Management and Engineering 2023 – 2027
B.Tech in Artificial Intelligence and Machine Learning

Central Academy Sr. Sec. School CBSE
Class XII (PCM)

Research & Projects

Physics-Informed ML for PM2.5 Forecasting — AISEHack (ANRF) 2026
Python, PyTorch, ConvLSTM, Wind Advection, Spatial Attention, SMAPE, Episode Detection

- Secured **1st place** across both phases of ANRF AISEHack 2026 (IIT Delhi + IBM Research India), forecasting PM2.5 concentration fields over India 16 hours ahead on a 140×124 spatial grid using stacked ConvLSTM cells with GroupNorm; Phase 1 private score 10.481, Phase 2 score 0.8925.
- Implemented a physics-informed Wind Advection Warp module using u10/v10 wind fields and differentiable grid sampling to simulate pollutant transport; in Phase 2 (no future met inputs), extended this with a learned WindPredictor head that forecasts future wind implicitly from 10-hour historical momentum.
- Diagnosed and fixed silent normalization failure on emission features (NOx, SO2, NH3) at 1e-11 scale using log1p scaling; added a multi-scale dilated SpatialEncoder, EpisodeDetector head, and episode-weighted SMAPE loss to explicitly target extreme pollution spike events evaluated under Phase 2 composite metric.

Voice Classification Deep Learning (VoxCeleb) Aug – Oct 2025
Python, PyTorch, CNN, Mel Spectrogram, Audio Augmentation

- Built an end-to-end speaker identification system on the VoxCeleb dataset, classifying 1000 speakers from raw audio using a 2D CNN trained on mel spectrograms.
- Designed a lightweight VGG16-inspired CNN with batch normalization and dropout, balancing model capacity with regularization for multi-class speaker classification.
- Applied on-the-fly mel spectrogram augmentation (time masking, frequency masking) to improve robustness and reduce overfitting on variable-length audio clips.
- Achieved **86.6% validation accuracy** using CrossEntropy loss, Adam optimizer, and early stopping.

Offline Multimodal RAG Platform 2025
Python, Flask, FAISS, CLIP, Whisper, Tesseract, OpenCV, Ollama

- Built a fully offline retrieval-augmented generation system performing semantic search across text, images, and audio — eliminating cloud API dependencies for privacy-focused deployment.
- Integrated FAISS vector retrieval with Whisper transcription and CLIP-based image embeddings for cross-modal querying.

FloatChat — Natural Language Climate Data Query System 2025
Python, FAISS, PostgreSQL, NetCDF/Xarray, LLM Integration

- Developed an intelligent NL query interface for Argo oceanographic datasets, enabling non-technical researchers to explore complex climate data without manual scripting.
- Reduced manual report generation time by **4**×, significantly improving researcher workflow efficiency.

Underground Pipe Leak Detection — Dhule 2.0 Smart City Initiative 2025
Python, Signal Processing, ML Classification, Hydrophone Sensors

- Designed a machine learning pipeline for detecting water leakage in underground pipes using acoustic signals from hydrophones and flow sensor data, developed for the Dhule 2.0 government smart city initiative.
- Presented at **Innovent Sobus** and selected as finalist, demonstrating real-world deployment potential for municipal infrastructure.

Achievements & Competitions

- **1st Place** - ANRF AISEHack Theme 2: Pollution Forecasting (Physics-Informed ML for PM2.5)
- **2nd Place** - Yugantar 2025 (Smart India Hackathon Internal Round)
- **Round 5** - HackRx 6.0 (Bajaj Finserv National Hackathon 2025)
- **Finalist** - Innovent Sobus (Dhule 2.0 Smart City - Underground Leak Detection)

Extracurricular

- **Member & Volunteer**, ACM Student Chapter — NMIMS MPSTME Shirpur Oct 2025 – Present
  Volunteered at *Initial Commit: AI Edition*, an AI/ML workshop, assisting attendees with technical doubts and concepts during session.

Certifications

- Machine Learning — Banaras Hindu University / SWAYAM (May 2025, **Score: 77%**)
- Deep Learning - IIT Ropar / NPTEL (Oct 2025, **Score: 82%. Silver + Elite**)

Technical Skills

Languages: Python, C++
ML / DL: PyTorch, TensorFlow, scikit-learn, AutoKeras, AutoML
NLP / LLM: BERT, LangChain, OpenAI API, Claude API, RAG pipelines, Agentic AI, MCP
Data & Scientific: NumPy, Pandas, SciPy, FAISS, NetCDF/Xarray, Spatio-Temporal Modelling
Vision & Audio: OpenCV, CLIP, Whisper, Matplotlib, Seaborn, Plotly
Web & Deployment: Flask, FastAPI, Streamlit
Databases & BI: SQL, MySQL, Power BI, SSIS
Tools: Git, Google Colab, Jupyter, LaTeX