

# ANIRUDH SARVE

Mumbai, Maharashtra, India | +91 93225 81894 | anisarve2006@gmail.com | anirudhsarve.netlify.app  
github.com/anisarve2006 | linkedin.com/in/anirudhsarve

## PROFESSIONAL SUMMARY

Computer Engineering undergraduate specializing in Machine Learning, Computer Vision, and Deep Sensor-Fusion systems. Experienced in designing physics-constrained Kalman filters, cross-modal remote sensing inpainting networks, and graph-masked beam-search vision models. Adept at compiling high-efficiency neural models to resource-constrained edge runtimes (ONNX, WebAssembly, ARM Cortex-A) with mathematically guaranteed bounds and real-world benchmarked accuracy.

## EDUCATION

**Vidyalankar Institute of Technology (VIT), Mumbai — B.E. in Computer Engineering**

Expected May 2028

**Cumulative GPA: 8.90 / 10.0** | Affiliated with University of Mumbai

**Relevant Coursework:** Machine Learning, Deep Learning, Computer Vision, Sensor Fusion, Data Structures & Algorithms, Linear Algebra, Probability & Statistics, Operating Systems.

**Sonopant Dandekar Shikshan Mandali — HSC (Class XII, PCMB)**

May 2024

**Score: 92.60%** | Distinction in Mathematics, Physics & Computer Science

## TECHNICAL SKILLS

**AI & Deep Learning:** PyTorch, TensorFlow, Hugging Face, Torchvision, ONNX, Scikit-learn, OpenCV, Faster-Whisper, Albumentations, NumPy, Pandas.

**Architectures & Methods:** 15-State Error-State Kalman Filter (ESKF), Sensor Fusion, Temporal Attention, ResU-Net, ResNet-101, EfficientNet-B0, Mel-Spectrogram Forensics, Cross-Modal Attention, Constrained Beam Search, ZUPT.

**Deployment & Systems:** ONNX Runtime, WebAssembly (WASM), WebGL, FastAPI, Flask, Docker, C++, Python, Git, Linux.

**Databases & Cloud:** SQLite (Vector BLOB scanning), Qdrant, MongoDB, PostgreSQL, Google Cloud Run, AWS.

## PROFESSIONAL EXPERIENCE

**NxtChaptr — Software Developer (Edge AI & Full-Stack)**

Nov 2024 – Aug 2025

- Architected and deployed on-device edge computer vision for automated photo matting, integrating a U<sup>2</sup>-Net segmentation model via **WebAssembly/WebGL** to process product images client-side in under **3.5 seconds**.
- Eliminated cloud-GPU inference costs (~\$0.02 per upload) and mobile network upload overhead by executing background removal directly inside seller browsers prior to Cloudinary cloud sync.
- Engineered automated catalog metadata extraction microservices utilizing **Google Gemini Vision**, reducing average listing creation latency from 8 minutes to under 30 seconds.
- Integrated Upstash Redis atomic distributed locks (SET NX EX) and Ably WebSockets to eliminate race-condition double purchases on 1-of-1 vintage catalog items during live pop-up drops.

*Exposure: Edge AI (WASM, WebGL, U<sup>2</sup>-Net), Python, Node.js, Next.js, Redis, MongoDB, Cloudinary.*

**Wroomie — Software Developer (Backend & Applied ML)**

Nov 2023 – Jan 2024

- Developed multi-attribute compatibility scoring algorithms clustering user lifestyle preferences across 4 behavioral vectors (sleep schedule, cleanliness, dietary rules, study habits), reaching a **94% match satisfaction score**.
- Built automated media ingestion and validation pipelines via Cloudinary and OpenCV, standardizing room verification images and KYC profile assets.
- Optimized MongoDB geospatial room discovery queries, decreasing query latency by **35%** across high-density university rental zones.

*Exposure: Applied Clustering, Python, Node.js, Express, MongoDB, RESTful API Architecture.*

## FEATURED TECHNICAL PROJECTS

**Neural-Inertial Dead-Reckoning Engine with 15-State Error-State Kalman Filter (ISRO SIH)**

*Technologies: PyTorch, 15-State ESKF, ONNX, C++*

- Developed a hybrid dead-reckoning engine for ISRO (PS 26168) holding autonomous vehicle position drift to  $\leq 1.85$  m across 120-second complete GPS blackouts (a **98.4% drift reduction** versus standard EKF baselines).
- Fused 10 Hz raw IMU strapdown mechanization with a 15-state Error-State Kalman Filter ( $\delta p, \delta v, \delta \theta, \delta b_a, \delta b_g \in \mathbb{R}^{15}$ ), enforcing physical Non-Holonomic Constraints (NHC) and ZUPT stationary zero-velocity updates.
- Designed a 5-head temporal self-attention GRU neural network predicting forward velocity and dynamic measurement covariance ( $R_v$ ), allowing the filter to adaptively weight network confidence against road irregularities.
- Benchmarked across the IO-VNBD dataset (73 journeys, 953k samples, 26.5 driving hours); compiled model to a compact **24.6 KB ONNX graph** running at sub-3 ms CPU latency on ARM Cortex-A ECUs with zero cloud dependencies.

## Multimodal Optical-SAR Fusion ResU-Net with Cross-Modal Attention (ISRO LISS-IV)

*Technologies: PyTorch, ResU-Net, SAR, GeoTIFF, GDAL*

- Developed an operational generative AI framework for automated cloud removal and surface reconstruction in ISRO Resourcesat-2/2A LISS-IV satellite imagery (5.0 m resolution, Green/Red/NIR bands).
- Engineered a multi-encoder ResU-Net with cross-modal attention bridging cloudy optical LISS-IV observations, cloud-penetrating Sentinel-1 SAR radar (VV/VH backscatter), and multi-temporal clear-sky reflectance priors.
- Enforced hard clear-pixel preservation ( $M \cdot G + (1 - M) \cdot \text{Orig}$ ) to maintain 100% bit-exact original observations for cloud-free pixels, restricting generative synthesis strictly to occluded regions.
- Formulated a differentiable Spectral Angle Mapper (SAM) loss and aleatoric Laplacian uncertainty head, achieving **32.09 dB cloud-region PSNR**, **0.7167 SSIM**, and a tight **6.17° SAM angle** on real ISRO validation scenes over the Assam corridor.

## Graph-Constrained Hierarchical Spatial Attention Model for Visual Taxonomy Classification

*Technologies: PyTorch, ResNet-101, LSTM, Scrapy*

- Engineered a sequence-based vision-to-taxonomy classifier mapping fashion imagery to hierarchical category paths across 52 categories, eliminating invalid hierarchy transitions through graph-masked beam search.
- Coupled a ResNet-101 spatial feature extractor (14×14 feature grids) with an LSTM spatial-attention decoder, dynamically attending to discriminative garment sub-regions (collars, cuffs, hemlines).
- Implemented decode-time taxonomy adjacency masking directly on transition logits, mathematically guaranteeing **100% structurally valid prediction paths** under production catalog ingestion SLAs.
- Assembled and curated a domain dataset of 184,000 multi-category clothing images from 13 e-commerce platforms using custom distributed Scrapy spiders.

## Cross-Modal Spectral-Visual Biometric Deepfake Forensic Classifier & Temporal Tamper Detector

*Technologies: PyTorch, EfficientNet-B0, Mel-Spectrogram, OpenCV*

- Built a cross-modal forensics verification pipeline fusing visual facial manipulation clues with acoustic speech synthesis artifacts to expose deepfake videos in real time.
- Implemented an EfficientNet-B0 visual classifier detecting boundary-blending seams and frequency-domain latent anomalies; coupled with a mel-spectrogram analyzer identifying synthetic vocoder phase mismatches and robotic micro-jitter.
- Designed a late-fusion probabilistic scoring engine outputting executive authenticity confidence alongside localized spatial-temporal tamper heatmaps for forensic explainability.

## Air-Gapped Multimodal Media Understanding Pipeline with Prompt-Conditioned Beam Decoding

*Technologies: Faster-Whisper, pHash, SQLite Vector, NumPy*

- Architected an entirely offline, air-gapped multimodal video intelligence pipeline operating on 16 GB RAM commodity hardware with zero external API calls or third-party cloud subscriptions.
- Implemented prompt-conditioned beam search decoding in Faster-Whisper using contextual domain priming packs, accurately capturing conversational Hindi-English (Hinglish) code-switched technical speech without model fine-tuning.
- Developed an in-memory SQLite float32 BLOB dot-product vector search achieving **sub-12 ms latency** across 50,000 video segments; integrated perceptual hashing (pHash) to filter **85% of redundant visual frames** in <1 ms on CPU.

## HONORS & LEADERSHIP

- **Lead Architect, ISRO Smart India Hackathon (SIH 2026):** Authored core AI-ESKF sensor-fusion architecture for problem statement PS 26168 (Smart Vehicles).
- **Autonomous Navigation Member, Indian Rover Challenge (IRC):** Programmed onboard kinematic telemetry, IMU state estimators, and computer vision guidance pipelines.
- **Scholastic Distinction:** Ranked in top percentiles with 8.90/10.0 CGPA at VIT Mumbai and 92.60% in HSC Maharashtra State Board.