

EDUCATION

GITAM University, Visakhapatnam
B.Tech, second year

2025 – 2029

PROJECTS

VIYGAN-2: Multi-temporal GAN super-resolution of Sentinel-2 Python · PyTorch · SIH26142 (NTRO)

- Fine-tuned a multi-date ESRGAN (RRDBNet, 16.7 M params) on 103,936 paired patches from 55 Indian AOIs, mapping an 8-epoch Sentinel-2 L2A stack from 10 m GSD to a 2.39 m grid.
- On 7 held-out AOIs (2,800 tiles): edge-F1 0.642 and LPIPS 0.182 vs 0.524 / 0.382 for the Satlas baseline, ahead of ESA SEN2SR, LDSR-S2 and Real-ESRGAN; 0.010 px geometric drift.
- Built s2colour, a variant with an in-loop spectral colour lock per 10 m cell, plus a per-pixel confidence map; served via FastAPI GPU inference, XYZ tile pyramid and GeoTIFF export (viygan-2.pages.dev).

Rail Section Controller: CP-SAT decision support for Indian Railways Python · OR-Tools · FastAPI

- Builds any corridor on demand from nationwide open data (8,966 stations, 5,208 trains, 417,080 schedule rows) with a precomputed national OpenStreetMap rail graph.
- CP-SAT model with NoOverlap line occupancy and Cumulative platform capacity, priority-weighted tardiness objective, FCFS warm start and staged relaxation; cut average delay 47.90 → 35.14 min on an 80-station corridor in a 10 s budget.
- Wall-clock-driven simulation with disruption injection, what-if scenario forking and live WebSocket state.

KalJahaz: Dry-bulk freight forecasting for India's East Coast ports Smart India Hackathon · kaljahaz.pages.dev

- Online adaptive recursive least squares per vessel class on Baltic sub-indices, China PMI, India IIP and iron-ore prices; voyage-cost forecasts at T+0/7/30/60 days with the optimal fixing window.
- Leg-by-leg voyage model: great-circle routing via lane waypoints, Kwon speed loss, currents, tides and port LOA / beam / draft constraints; live data ingestion every 6 h with per-series provenance.

Alfred: Offline agentic voice assistant llama.cpp · Ollama · sherpa-onnx

- Local quantised 9B LLM on an 8 GB laptop GPU (32k context, 8-bit KV cache, vision) with tool calling for apps, web, screen, files and reminders.
- Persona via activation steering (control vector, no prompt or fine-tuning); at coefficient 0.1 GSM8K rose 90 → 93 and BFCL 87 → 90.
- Streaming speech stack: Silero VAD, Parakeet-TDT ASR (~80 ms), Kokoro TTS with barge-in.

ML data cleaning & normalisation pipeline (demo: sign language to speech) Python · NumPy · pandas · scikit-learn

- Model-agnostic preprocessing shared by training and inference: One Euro denoising of every record, NaN-safe handling of degenerate input, translation / scale / in-plane rotation normalisation, and 63 raw values engineered into 36 invariant features; every invariance unit-tested, stratified evaluation split.
- Demo on edge hardware: MediaPipe landmarks from a live camera into a 100-tree random forest over 44 hand signs (2,526 samples) at 2.15 ms per prediction, spoken via TTS on a Raspberry Pi.

NIP: Low-latency PC → Raspberry Pi screen streaming C · ffmpeg · GStreamer

- DXGI capture with hardware-encoder auto-selection (NVENC / QuickSync / AMF / x264); Pi-side hardware decode straight to KMS with no compositor or jitter buffer; LAN, Wi-Fi AP or USB-gadget transport.

CRT Terminal Portfolio · Embedded & Robotics three.js · WebGL2 · ESP32 · TETRIS PRIZM

- This portfolio: a working shell rendered through custom GLSL CRT shaders, with a five-flag CTF and client-side crypto tools.
- ESP32-CAM streaming with YOLO detection, GNSS tracking, Bluetooth rover, SPI TFT display, TETRIS PRIZM robotics.

RESEARCH

- Persona control via activation steering:** contrastive activation addition from 300 Wodehouse lines, injected into layers 8–24; dose-response measured on XSTest, GSM8K and BFCL.
- Autonomous continuous cognition** (ongoing): a never-prompted local LLM with graph memory instead of a context window.

SKILLS

Languages	Python, C, C++, JavaScript, GLSL, SQL, Bash
Deep learning	PyTorch, PyTorch Lightning, torchvision, timm, Hugging Face (transformers, Hub, safetensors), GANs / ESRGAN / BasicSR, LPIPS, OpenCLIP, ONNX Runtime
LLMs & speech	llama.cpp, Ollama, GGUF quantisation, activation steering, tool calling, sherpa-onnx, Silero VAD, Kokoro TTS, pyttsx3
Data & vision	NumPy, pandas, SciPy, scikit-learn, Matplotlib, OpenCV, scikit-image, Pillow, Kornia, MediaPipe, Ultralytics YOLO, joblib
Geospatial	rasterio, pyproj, Shapely, GeoTIFF, Sentinel-2 / Copernicus Data Space, OpenStreetMap / Overpass, Leaflet
Backend & MLOps	FastAPI, Pydantic, Uvicorn, OR-Tools CP-SAT, SQLite, PostgreSQL, Redis, Docker, Playwright, pytest, Weights & Biases, TensorBoard, Git / Git LFS, Cloudflare
Embedded	Raspberry Pi, ESP32-CAM, ESP8266, Arduino, TETRIS PRIZM, Linux
Security	OWASP Top 10, Nmap, Wireshark, Linux hardening, hashing and JWTs, CTFs