VENKATESH R

AI Researcher

hessiancurve@gmail.com | +91 8668135371 | LinkedIn | GitHub | Portfolio | Dharmapuri, India

Applied AI researcher building multimodal ML systems — face recognition, NLP, speaker identification, and hybrid quantum-classical models. Designed cross-modal embedding architectures (384/512/256-dim) with clustering and retrieval pipelines deployed to production. Research work spans climate modeling (LSTM + SAR), quantum ML (PennyLane VQC), and zero-knowledge cryptography. 5x national hackathon winner, ₹12L+ in government prizes.

TECHNICAL SKILLS

ML / Deep Learning	PyTorch, TensorFlow, InsightFace (ArcFace), Florence-2, PaddleOCR, YOLO, Seamless-M4T, DBSCAN, XGBoost, LightGBM, LSTM, Isolation Forest, Scikit-learn, model evaluation, ablation studies
LLM / NLP	Gemma 3, vLLM, Ollama, DSPy (prompt optimization), LangChain, LangGraph, DistilBERT, DeBERTa, NER, sentiment analysis, RAG, structured generation, fine-tuning (QLoRA/LoRA), RLHF concepts
Embeddings & Retrieval	Qdrant, FAISS, sentence-transformers (384-dim), VGGFace2 (512-dim), Resemblyzer (256-dim), cosine similarity, HNSW indexing, cross-modal retrieval, embedding evaluation
Computer Vision	Object detection (Florence-2, YOLO), face recognition (ArcFace, DBSCAN clustering), OCR (PaddleOCR, Tesseract), video analysis, crowd density, on-device inference (TFLite, CoreML)
Quantum ML	PennyLane, Variational Quantum Circuits (4-qubit VQC), hybrid quantum-classical pipelines, XGBoost/LightGBM baselines
Languages & Infra	Python, C++, Rust, SQL

RESEARCH & PRODUCTION EXPERIENCE

AI Engineer (Consultant) — Saptang Labs, Chennai *(Jun 2024 – Present)*

Multimodal Media Analysis Platform — designed and built the complete ML research and inference layer for a distributed platform (5 services, 7 databases) processing audio, video, and image data across multiple modalities.

- Designed **cross-modal embedding architecture** — text (384-dim, sentence-transformers), face (512-dim, ArcFace), voice (256-dim, Resemblyzer) — 3 Qdrant collections with HNSW indexing. Evaluated retrieval precision and clustering coherence across modalities.
- Built **face recognition pipeline**: RetinaFace → landmark alignment → 512-dim ArcFace → DBSCAN clustering (cosine, eps=0.65) for unsupervised person grouping. **Reduced identification from days to seconds.**
- Replaced CLIP zero-shot with **Florence-2 generative detection** — 16 fixed labels → open-vocabulary bounding-box detection. Measured recall improvement on unseen categories.
- Engineered **speaker diarization + cross-file matching**: Seamless-M4T ASR → speaker separation → 256-dim voice embeddings → Qdrant retrieval across 500+ media files.
- Built **number plate deduplication** — OCR + Levenshtein distance clustering (threshold 0.85) to consolidate partial reads across sampled video frames.

Nucleus — 11 NLP Research Packages shipped as production Python wheels with typed Pydantic interfaces and dual Ollama/vLLM backends:

- **NER + Relationship Extraction** — two-step Gemma 3 12B pipeline (~3.3 sec/post). Entity detection with alias resolution → relationship extraction → ArangoDB graph edges for multi-hop queries.
- **DSPy NER** — BootstrapFewShot metric-driven prompt optimization. Compared structured output consistency vs. raw prompting — measurable format adherence improvement.
- **Risk Classifier** (DeBERTa zero-shot, IT Rules 2021) | **Psychological Classifier** (6 behavioral metrics) | **Sentiment** (80% acc, 0.22s/post, 10K+/day) | **Translation** (14 languages, code-mixed) | **Summarizer** | **Event Extractor**.

LangGraph Automated Report Agents — automated multi-day analyst workflow: 94+ article ingestion → Qdrant vector deduplication → LLM relevance scoring → formatted .docx report generation in **25 minutes (99% time reduction)**.

RESEARCH PROJECTS & PUBLICATIONS

Glofsense — Multimodal Climate Risk Prediction *(C-DAC HimaShield 2025 Winner, ₹5L)* — PIB GLOF early warning: Sentinel-1 SAR imagery + LSTM temporal modeling + XGBoost risk classification + ESP32 sensor fusion + water flow simulation for evacuation planning.

Eigencurve — Hybrid Quantum-Classical Threat Detection *(Terrier Cyber Fest 2025, 3rd Prize)* PennyLane VQC (4-qubit) vs. XGBoost/LightGBM on NetFlow data. Evaluated quantum advantage on non-linear feature spaces. VirusTotal IOC enrichment.

Nirnay — Real-Time Misinformation Detection *(NFSU 2025 Winner, ₹1L)* FFMPEG live capture → Whisper ASR → DistilBERT claim extraction → Fact Check API → Streamlit dashboard with confidence scoring.

Xylem — Zero-Knowledge Identity — ZK-SNARK (Circom, Groth16) for privacy-preserving KYC. 500ms proof, <1KB size. DPDP Act compliant.

Blackfence — Isolation Forest anomaly detection on CICIDS2017 → auto-generated ModSecurity WAF rules with z-score explanations.

Meristem — Geospatial analysis of 125M Aadhaar enrollment records — surge detection, coverage gaps, demographic insights.

EDUCATION

Degree	Institution	GPA	Duration
BS Data Science & Applications	Indian Institute of Technology Madras	7.00	2023 – 2026
B.Tech Biotechnology	Bannari Amman Institute of Technology	7.93	2022 – 2026

Certifications: CS50 AI (Cambridge) · Banking & Finance (IIT Kanpur)

Awards: 5x National Hackathon Winner (₹12L+) · Hedera DLT Climate Winner ($9K) · BRICS 2nd Global (Cybersecurity) · NFSU 2023 Winner (₹5L)