Vignesh Balamurugan M.B

+91 6381538287 | vigneshbalamurugan.m23@iiits.in | linkedin.com/in/vignesh-balamurugan | github.com/mbvb1312

Education

Indian Institute of Information Technology, Sri City 2023 – 2027
*B.Tech in Artificial Intelligence & Data Science | CGPA: **8.8*** Sri City, India

Edison G Agoram Memorial School 2021 – 2023
*CBSE XII | **94%*** Chennai, India

Professional Experience

Summer Intern 05/2025 – 07/2025
National Institute of Technology (NIT), Trichy Trichy, Tamil Nadu

- Researched Human Activity Detection and Recognition using **Vision Transformers**, temporal sequence modeling with **LSTMs**, and video preprocessing pipelines; delivering **18% accuracy improvement**.
- Contributed to a research paper under review and engineered a **video captioning pipeline** from video input.
- Minimized **inference latency by 22%** through optimized preprocessing and lightweight fine-tuning strategies.

Projects

RAG Fact Verification System | *DeBERTa-v3, DPR, FAISS, RRF, MLflow* [GitHub]

- Constructed a multi-hop fact-verification pipeline combining Wikipedia retrieval, **DeBERTa-v3 NLI**, DPR dense retrieval, and Reciprocal Rank Fusion (RRF).
- Integrated **3-hop evidence chaining** with hybrid TF-IDF + dense retrieval alongside experiment tracking via MLflow, DVC, and Prometheus/Grafana.

Crowd Stampede Risk Predictor | *CANNet, ResNet-34 FPN, DroneCrowd, Temporal Risk Modeling* [GitHub]

- Architected a dual-head **CANNet (ResNet-34 FPN)** framework for simultaneous crowd-density estimation and person-localization from aerial surveillance footage.
- Formulated a temporal crowd-risk propagation engine utilizing grid-level density evolution for proactive identification of **stampede-risk regions**.

Meta-Agent — Multi-Agent LLM Orchestration Framework | *LangGraph, Redis, ChromaDB, Kubernetes HPA* [GitHub]

- Devised a graph-driven multi-agent orchestration framework supporting sequential, parallel, DAG-based, and hierarchical LLM execution pipelines.
- Incorporated Redis/ChromaDB-backed persistent memory and an LLM-as-judge evaluation loop with Kubernetes HPA autoscaling support.

STRIDE — Intelligent Transportation System | *Optical Flow, ROI Tiling, Dynamic Lattice Modeling, OpenCV* [GitHub]

- Proposed a **13-stage STRIDE architecture** leveraging a 2-tier parallel virtual lattice (8x8 ROI grid) for localized density, velocity, and directional traffic-flow estimation.
- Engineered adaptive region-based traffic decomposition through dynamic ROI tiling and motion-vector aggregation for congestion hotspot identification.
- Attained **0.70 IoU accuracy** and **0.69 directional-flow accuracy** on region-level traffic flow estimation benchmarks.

Chennai Night Safety Intelligence | *QGIS, OSRM, rasterio, geopandas, Spatial MLP* [GitHub]

- Integrated **12+ geospatial datasets** encompassing CCTV density, police coverage, streetlight intensity, and crime heatmaps for urban risk modeling.
- Orchestrated spatial-query pipelines generating **61 CCTV nodes, 60 streetlight clusters, and 43 patrol regions** using raster-vector geospatial analysis.
- Trained a spatial-context MLP utilizing 9-channel neighborhood features and constructed a route-risk comparison engine yielding **10.5% peak risk reduction** with only **+0.25 km** average detour overhead.

RL for CRISPR Gene Editing Optimization | *SARSA, Policy Iteration, MDPs, Drift-Aware RL* [GitHub]

- Modeled CRISPR gene-editing optimization as a 2D grid-world MDP and employed **Value Iteration, Policy Iteration, and SARSA** for adaptive policy learning.
- Augmented the environment with non-stationary dynamics using a drift-aware hybrid RL agent for accelerated post-drift recovery.

Technical Skills

AI & Machine Learning: PyTorch, TensorFlow, Scikit-Learn, OpenCV, HuggingFace Transformers, FAISS, LangChain, LangGraph, RAG, Reinforcement Learning, Vision Transformers
Cloud & MLOps: Docker, Kubernetes, GitHub Actions, MLflow, DVC, AWS, Prometheus, Grafana
Backend Development: FastAPI, Flask, Node.js, REST APIs, SQLite
Data & GIS: Pandas, NumPy, QGIS, geopandas, rasterio, SQL
Tools & Infrastructure: Docker, Git, Linux, Postman, Vercel, Jupyter Notebook
Concepts: Data Structures and Algorithms, Object-Oriented Programming, System Design, Distributed Systems, Agile
Languages: Python, Java, C, SQL, JavaScript

Certifications & Achievements

Supervised ML — Coursera (99.83%) [Verify] | **Advanced Learning Algorithms — Stanford/DeepLearning.AI (100%)** [Verify]
Data Analytics Job Simulation — Deloitte via Forage | **Global Game Jam 2024 — Unity game shipped**