Manishwaran Murugan

Santa Clara, CA 95050
408-375-7136 muruganmanishwaran@gmail.com linkedin.com/in/Manishmrgn github.com/Manishmrgn

Education

Santa Clara University **Sep 2025 - Jun 2027 (Expected)**
Master of Science in Computer Science and Engineering; GPA: 3.79/4.0 *Santa Clara, CA*

- Coursework: Machine Learning, Deep Learning, Reinforcement Learning, High Performance Computer Architecture, Artificial Intelligence, Data Structures & Algorithms

Vellore Institute of Technology, B.Tech. in Computer Science and Engineering, Vellore, India **Sep 2020 - May 2024**

Experience

Dimaag-AI **Jun 2026 - Aug 2026**
AI Engineer Intern *Fremont, CA*

- Built a synthetic SFT data-generation pipeline for a domain-specific LLM tool-calling assistant, producing multi-turn dialogues grounded in real backend tool outputs.
- Designed multi-tool routing with versioned YAML constraint prompts; replaced hard-coded output patches with regex validators and corrective retry loops to keep synthetic artifacts out of training data.
- Developed a Qwen3 32B agent harness with network-isolated code execution and LLM-based context compaction; deduplicating tool schemas cut fixed per-call cost 45% (9.2k to 5.0k tokens), leaving 85% of a 32k context window free.
- Benchmarked vision-language models (Qwen3-VL 32B, Qwen2.5-VL 72B); built document VLM and RAG ingest tools.

ITC Infotech Ltd. **Sep 2024 - Aug 2025**
Business Analyst *Bengaluru, India*

- Built a time series price forecasting model (Type-2 fuzzy logic regression) for the ITC Ltd. agri business, reaching 92.7% forecast accuracy (7.3% MAPE).
- Automated data pipelines, ran controlled experiments, and evaluated models with MAE, RMSE, and ablations.

Indian Institute of Science **Jun 2023 - Jul 2024**
Research Intern (Capstone) *Bengaluru, India*

- Built end-to-end 3D reconstruction with CUDA-optimized NeRF and Structure from Motion (Meshroom) under Dr. Suresh Sundaram, adding ArUco-based metric calibration to recover true scale for as-built modeling.
- Automated large-scale synthetic data generation in Unreal Engine 5 with C++, Blueprints and Python by randomizing marker IDs and sizes, scene and camera poses, lighting, motion blur, occlusions and lens distortion.
- Trained and evaluated 10 YOLOv8 object detection + cv2.aruco pipelines (k-fold CV, adaptive thresholding, corner refinement) to 98.6% mAP@50, recovering absolute scale within 0.2 cm on an 18 cm marker.
- Engineered parallelized KD-Tree/FLANN and Octree point cloud background subtraction, cutting runtime on 10M-point clouds from 27 h to 40 s (2,400x faster).

Projects

AI-Town Narrative Simulation | *Python, LLaMA-3, Docker, React, Node.js* **2025**

- Extended the AI-Town framework into an interactive fantasy world, running an offline LLaMA-3 model in Docker for local NPC dialogue and world updates.
- Built a real-time narration system that converts NPC conversations into story events and added player-NPC interactions that can change the narrative.

UAV-based Coastal Debris Detection | *YOLOv8, Deep Learning, UAV Imaging* **2024**

- Benchmarked 5 YOLOv8 variants with hyperparameter tuning, Wise-IoU and augmentation on UAV debris imagery, raising YOLOv8m from 81.7% to 86.1% mAP@50; found Wise-IoU degraded the largest models.

Technical Skills

Languages: Python, C++, C#, JavaScript
ML / LLM: PyTorch, TensorFlow, Hugging Face, CUDA, vLLM, llama.cpp, Ollama, RAG, SFT, LLM Agents
Computer Vision: OpenCV, YOLOv8, NeRF, Structure from Motion, COLMAP, Meshroom
Tools: Docker, Linux, Git/GitHub, Unreal Engine 5, Unity, React