Harshith Mohan Kumar

San Francisco, California | hiharshith18@gmail.com | +1(951)-347-6462 | harsh188.github.io

Education

University of California, Riverside | M.S in Computer Science **Mar. 2025**

GPA: 3.92/4.0 | Coursework: Deep Learning, Advanced Computer Vision, Reinforcement Learning

PES University, Bengaluru India | Bachelor's in Computer Science **2023**

Coursework: Data Analysis, Natural Language Processing, Image Processing, Operating Systems, Computer Networks

Technical Skills

Programming Languages:	Python, R, C++17, C, JAVA, JavaScript
Machine Learning:	TensorRT, PyTorch, TensorFlow, OpenCV, Wandb, Scikit-learn
Software:	Docker, Kubernetes, Profiling, AWS, GCP, Carla, GitHub Actions

Experience

UC Riverside Riverside, CA

Graduate Student Researcher 06/2024 – Present

- **Synchronized and calibrated** a RGB-Thermal stereo camera to develop the first thermal optical flow network using self-supervision with low-rank domain adaptation leading to 0.98 SSIM on reconstructed 1.57 mega pixel images using PyTorch.
- Optimized model inference and automated monitoring using **CUDA** for multi-scale parallel photometric evaluation and memory profiling to enable a memory efficient **real-time motion tracking** using less than eight gigabyte of GPU memory.

Intel Bengaluru, KA

Machine Learning - Autonomous Driving 01/2023 – 08/2023

- Developed severity-ranking metrics for ADAS collision alerts using **multimodal 3D scene understanding** (LiDAR, dashcam, telemetry) and **domain adaptation** to improve out-of-distribution robustness, reducing false positives by 30%.
- Coordinated with a cloud architecture team to design a periodic job pulling data from AWS S3 & MongoDB, enabling continuous model updates with minimal downtime across clusters.

Google Summer of Code Los Angeles, CA (Remote)

Open-Source Contributor 06/2022 – 08/2022

- Efficiently **time segregated terrabytes of digitized VCR recordings** using early stage multimodal, audio, text and image, feature extraction and DBSCAN clustering with CUDA optimized parallel execution.
- Utilized thread-safe consumer-producer pattern, achieving 8x concurrent file processing, while maintaining stable memory consumption through profiling; thereby localizing and minimizing compute overhead.
- Designed and orchestrated a multi-stage **largescale distributed ML pipeline** using SLURM. **[Website]**

Enrole Inc. Boulder, CO (Remote)

Data Science - Recommender Systems 06/2020-08/2020

- Architected and implemented RESTful APIs using Python Flask, following test-driven development with PyTest achieving 95% test coverage and reducing integration bugs by 60% through systematic code reviews.
- **Containerized and deployed** a clustering-based recommender system using scikit-learn, implementing proper version control for both code and model artifacts using Google Cloud Run service for mobile application.

Academic Projects

Layer-wise Safety Alignment of Vision Language Models using RLHF

- Utilized reinforcement learning for safety alignment against adversarial attacks within the encoder layers of LLaVA model.

Unsupervised Surveillance Tracking

- Developed an unsupervised multi-object tracking system using pre-trained deep networks (ResNet-50, YOLOv8) for feature extraction, achieving consistent tracking across 10,000+ frames.
- Containerized the system using Docker with GPU-optimized CUDA support, enabling seamless deployment on edge devices (NVIDIA Jetson) and cloud instances (AWS EC2).

3D Perception and Localization in CARLA

- Built a sensor data pipeline in CARLA for real-time merging of LiDAR point clouds and camera frames, achieving consistent 3D bounding box detection within 5ms latency.

Publications

OCTraN: 3D Occupancy Convolutional Transformer Network, **CVPR Transformers for Vision Workshop** - 2023

Fusing Pseudo Labels with Weak Supervision for Dynamic Traffic Scenarios, **ICCV Workshop** - 2023

GraphCoReg: Co-Training for Regression on Temporal Graphs, **ECML-PKDD Workshop** - 2022 (Best student paper award)