

SRECHARAN SELVAM

+1 917-495-3410 srecharan@gmail.com [linkedin.com/in/srecharan](https://www.linkedin.com/in/srecharan) Srecharan.com github.com/Srecharan

Education

Carnegie Mellon University

Master of Science, Research: Machine Learning

Relevant Coursework: Machine Learning, Deep Learning, Advanced Computer Vision, Visual Recognition

May 2025

Pittsburgh, PA

GPA - 4.0/4.0

Experience

Kantor Lab, Carnegie Mellon University [GitHub ↗]

Aug 2023 – Present

Graduate Research Assistant

- Developed a Vision Language Action (VLA) system through LoRA fine-tuning of LLaVA-1.5-7B (CLIP + Vicuna) foundation model, adapting multimodal reasoning from text generation to action prediction for robotic leaf manipulation.
- Automated data creation via self-supervised learning pipeline, eliminating **100%** manual annotation for GraspPointCNN.
- Trained attention-based GraspPointCNN using MLflow to track **60+** model experiments for grasp point optimization.
- Boosted inference (**20→27 FPS**) by parallelizing 3D projection with CUDA kernels & compiling models with TensorRT.
- Deployed Dockerized VLA-enhanced grasping stack to a 6-DOF robot, achieving **82%** leaf grasp success rate in field tests.

Hanon Systems [GitHub ↗]

Jan 2023 – Jun 2023

Machine Learning Engineer Intern

- Built GRU-based 3D hand gesture (dynamic/static) recognition model trained on custom dataset with EKF smoothing for real-time AR interface (Unity), allowing **100+** automotive technicians to safely simulate HVAC assembly procedures.
- Optimized inference using CUDA kernels and ONNX quantization, cutting latency by **33%** & memory footprint by **50%**.
- Exposed gesture and depth modules via Flask REST APIs for modular model serving to downstream applications.
- Containerized complete pipeline by utilizing Docker for consistent, scalable deployment across **3** training centers.

Vee Ess Engineering [GitHub ↗]

Jul 2022 – Dec 2022

Computer Vision Engineer Intern

- Engineered distributed Apache Spark pipeline to boost recyclable material recovery on high-speed conveyors by handling terabytes of multi-camera footage and auto-annotating frames with Mask R-CNN to create **43,000+** segmented images.
- Streamlined dataset storage & versioning via AWS S3 to reduce I/O overhead by **24%** and accelerate training iterations.
- Integrated custom-trained YOLOv5 model across real-time camera streams yielding **96% mAP@[0.5:0.95]** with **<15ms** latency, enabling conveyor speed modulation based on detection density to cut manual sorting labor by **6+ hrs/week**.

Projects

VLM-Based Tool Recognition System for Industrial Safety Applications [GitHub ↗]

Jan 2025 – May 2025

- Fine-tuned Qwen-2.5-VL-7B & LLaMA-3.2-11B-V using LoRA on custom dataset via LLM-guided prompt engineering (**8K** images, **29K** annotations) for real-time multi-modal industrial tool recognition and safety guidance generation.
- Built a RAG pipeline using LangChain and Pinecone to ground tool-specific information, reducing hallucinations by **55%**.
- Implemented RLHF (GRPO) on AWS SageMaker to optimize preference learning on paired responses for VLM alignment.
- Orchestrated LLM-based evaluation pipeline (OpenAI API) with Kubernetes, scoring **4K+** outputs for **8** model variants.

Multi-Model Stock Prediction with NLP and Automated Trading [GitHub ↗]

Oct 2024 – Feb 2025

- Spearheaded distributed training infrastructure using data parallelism across 4x V100 GPUs to train ensemble ML models (bidirectional LSTM + XGBoost with **35+** features) for algorithmic trading system across multiple timeframes.
- Created Kafka streaming pipeline to ingest **9K** financial events/day from multiple APIs for low-latency trading decisions.
- Automated Apache Airflow workflows managing FinBERT sentiment analysis, boosting prediction accuracy by **~5%**.
- Designed automated trading system with CI/CD model retraining and Tradier API execution, delivering **58.5%** win rate.

GenAI for Synthetic Data Augmentation: GANs, VAEs & Diffusion [GitHub ↗]

Apr 2024 – Jul 2024

- Trained WGAN-GP, β -VAE, and DDPM on CUB-200-2011 to generate synthetic bird images for classifier augmentation.
- Produced **18K** synthetic images across **200** bird classes with optimized mixing ratios to address training data limitations.
- Validated synthetic data utility through ResNet-50, achieving **5.1%** accuracy gains and **15%** boost in low-data scenarios.

Skills

Languages & Frameworks: Python, C++, SQL, PyTorch, TensorFlow, OpenCV, scikit-learn, Transformers, ONNX, Git
ML Training: RAG, RLHF (GRPO, DPO), SFT, PEFT, LoRA, VLM/LLM Fine-tuning, Distributed Training, LangChain
Infrastructure: AWS (SageMaker, EC2, S3), GCP, Docker, Kubernetes, TensorRT, CUDA, MLflow, Kafka, Spark, Flask

Publications

Srecharan Selvam, Abhisesh Silwal, George Kantor "Self-Supervised Learning for Robotic Leaf Manipulation: A Hybrid Geometric-Neural Approach", [arXiv:2505.03702](https://arxiv.org/abs/2505.03702), Under review at ICCV 2025. [Project Page ↗] [PDF ↗]