

AKSHAY ANTONY

(470) 754-8727 | akshaymuth@gmail.com | LI: akshay-antony | Web: akshayantony.com | GS: scholar.akshay

EDUCATION

Carnegie Mellon University

Master of Science in Mechanical Engineering - Applied Advanced Study (GPA: 4/4)

Pittsburgh, PA

December 2022

- Courses: Machine Learning, Deep Learning, Computer Vision, Visual Learning and Recognition, Optimization

Indian Institute of Technology, Roorkee

Bachelor of Technology in Mechanical Engineering (GPA: 9.077/10)

Roorkee, India

May 2021

SKILLS

Robotics Skills: Deep/Machine Learning, Computer Vision, Path Planning, Reinforcement Learning, Optimization

Programming Languages: Python (Pytorch, NumPy, OpenCV, TensorRT), C++ (ros, Eigen, OpenCV), CUDA

EXPERIENCE

Wayve Technologies

Sunnyvale, US

Machine Learning Engineer

July 2025 - Present

- Worked on curating and mining scenario-specific datasets by analyzing model failures and identifying gaps in coverage, ensuring stronger training signals and more representative evaluation across all scenarios of driving.
- Currently fine-tuning Wayve's end to end world foundation model with curated data and additional input signals to improve performance and reliability in both highway and urban driving situations.

Autowx Technologies

San Jose, US

Software Engineer: Perception and Prediction

April 2025 - July 2025

- Enhanced MapTR for HD map prediction by speeding up inference and adding new lane/road feature classes, then extending it to output specific map change classes; trained on multiple RTX 3090 GPUs, achieving an average IoU of 0.91.
- Curated and improved limited multi-sensor datasets by filtering failed or mislabeled samples and creating new augmentation methods, enabling the model to learn effectively from scarce data and perform reliably in diverse urban driving scenes.

Arcrest Technologies

Fort Smith, US

Robotics Engineer, Perception

April 2023 - April 2025

- Developed 3D LiDAR calibration using keypoint matching and the Kabsch algorithm, achieving $\sim 0.40^\circ$ rotation and ~ 4 cm translation error, enhancing spatial accuracy. Engineered a novel LiDAR-to-robot base calibration via ground alignment and feature tracking, achieving $\sim 0.55^\circ$ rotation and ~ 2 cm translation error, improving perception across all company robots.
- Led the migration of the perception team's point cloud processing pipeline from CPU to GPU using CUDA programming on Nvidia Orin, boosting output frequency of processed clouds by 54% to 20 Hz and reducing CPU utilization by 15%.
- Contributed to the design of an end-to-end Autonomous Barcode Detection pipeline with oriented bounding boxes using YOLOv8, achieving a MAP of 0.89, reducing freight misplacement and enhancing inventory tracking.
- Developed a fused 3D LiDAR and camera-based ground segmentation pipeline with SegFormer, achieving 0.93 IoU, improving obstacle detection and safer navigation by identifying low-lying ground hazards missed by LiDAR alone.
- Designed a multimodal 3D Object Detection pipeline integrating 2D detections and LiDAR data to generate accurate 3D bounding boxes, enhancing object localization for autonomous systems by leveraging sensor fusion techniques.

Carnegie Mellon University

Pittsburgh, US

Research Assistant

February 2023 - April 2023

- Developed MOFGPT, a decoder only transformer model generating Metal-Organic Frameworks (MOFs) from SMILES strings with a perplexity of 1.2. Fine-tuned for adsorption energy prediction and implemented the reinforce algorithm ensuring generated MOFs meet energy, validity, and novelty criteria. ([publication link](#))
- Designed a novel multimodal pretraining approach using BERT for text info, and graph convolution for structural data of adsorbate-catalyst systems. Fine-tuned for energy predictions, reducing MAE loss from 0.71 eV to 0.35 eV. ([publication link](#))
- Fine-tuned Prot-Bert to predict variations in GPCR motifs to obtain a test accuracy of 97.6% and analyzed 3D structures to understand higher-order interactions within receptor conformations. ([publication link](#))

Thordrive

Cincinnati, US

Perception Engineer

January 2023 - February 2023

- Built a semantic segmentation pipeline for LiDAR data using a modified Cylinder3D model; trained on 4 NVIDIA 3090 Ti GPUs using PyTorch DDP, achieving 93.4% IoU for the aircraft class, and integrated results into existing perception modules.
- Made the model faster by 66.6% to run at 30 fps utilizing TensorRT and MinkowskiEngine library for fast sparse convolution and integrated it with the robotics stack.