AKSHAY ANTONY

(470) 754-8727 | akshaymuth@gmail.com | LinkedIn: https://www.linkedin.com/akshayantony

EDUCATION

Carnegie Mellon University   **Pittsburgh, PA**

Master of Science in Mechanical Engineering - Applied Advanced Study (GPA: 4/4)   *December 2022*

- Courses: Machine Learning, Deep Learning, Computer Vision, Visual Learning and Recognition, Optimization, Geometry   Based methods in vision, Advanced Engineering Computation (C++)

Indian Institute of Technology, Roorkee   **Roorkee, India**

Bachelor of Technology in Mechanical Engineering (GPA: 9.077/10)   *May 2021*

SKILLS

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

Application Software: ROS, AWS, Docker, Linux (Ubuntu), CARLA, Git

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

EXPERIENCE

Arcbest Technologies   **Fort Smith, AR**

Robotics Engineer, Perception I   *April 2023 - Present*

- Engineered a calibration technique for aligning two 3D LiDAR systems, utilizing keypoint matching on a diagonal board using the Kabsch algorithm, resulting in a rotation error of ~0.40 degrees and translation error ~2 cm enhancing sensor fusion accuracy.
- Designed a LiDAR to robot base link calibration method using ground alignment and feature tracking, achieving a translation error of ~3 cm and a rotation error of ~0.55 degrees, improving the robot's perception accuracy.
- 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%.
- Developed a freight dimensioning system using LiDAR data, employing motion-based filtering to track fork movement and isolate pallet points. Applied advanced computer vision algorithms to process the point cloud, achieving 2-5 inch dimensional accuracy in cluttered environments.
- 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.

Carnegie Mellon University   **Pittsburgh, PA**

Research Assistant   *Feb 2023 - April 2023*

- Developed MOFGPT, a decoder-only transformer model for synthesizing novel Metal-Organic Frameworks (MOFs) from SMILES strings with a perplexity of 1.2. Created a reward model by fine-tuning for adsorption energy prediction and implemented the REINFORCE algorithm to ensure generated MOFs meet energy, validity, and novelty criteria.
- 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)

Thordrive   **Cincinnati, OH**

Deep learning perception intern, Perception Engineer   *May 2022 - Aug 2022, Jan 2023 - Feb 2023*

- Built a deep learning-based semantic segmentation pipeline for lidar data by a modified Cylinder 3D model to obtain an IoU of 93.4% for aircraft class and integrated results into existing perception modules. Made the model faster by 66.6% to run at 30 fps utilizing the MinkowskiEngine library for fast sparse convolution and integrated it with the robotics stack.
- Created an image processing algorithm for lane detection after transforming the front-center camera image to a bird's eye view by employing concepts of geometric transformation and homography.

RESEARCH PROJECTS

Camera-Lidar Autocalibration   *Oct 2022 - Dec 2022*

- Performed targetless camera-lidar extrinsic calibration by optimizing an objective that maximizes the agreement between semantics of image and lidar, using Powell's Conjugate method to obtain an MSE of 0.07085 degrees and 0.032 cms.

Weakly Supervised Lidar Object Detection   *Mar 2022 - May 2022*

- Constructed a Deep learning model for 3d object detection capable of being trained with only class labels and 3d bounding box proposals produced by density-based clustering, without expensive bounding box ground truth labels
- Fused 3D proposals with camera image and trained WSDNN to acquire an AP (Average Precision) of 29 for car class on KITTI.

Deep Learning-based Point Cloud Registration   *Mar 2022 - May 2022*

- Developed a graph attention-based point cloud registration algorithm, reducing the MSE loss of 3D rotation matrix by 14% on the ModelNet40 dataset, outperforming previous methods.