

MAYANK DESHPANDE

Austin, TX, 78758

📞 240-481-8418 ✉ msdeshp4@umd.edu 🌐 [Website](#) [LinkedIn](#) [Github](#) [Medium](#)

Summary

Perception engineer who owns production 3D perception end-to-end: multi-camera RGB-D and LiDAR sensing, deep-learning segmentation, 6-DoF pose estimation, and the data pipelining, model training and evaluation loop that keeps them working in the field. Sole owner of the perception stack on a deployed autonomous-robot fleet, with a track record of developing highly reliable production software with high ownership and accountability.

Experience

Contoro Robotics

Oct. 2025 – Present

Robotics Perception Engineer

Austin, TX

- Own the perception stack for entire fleet of autonomous box unloading robots. RGB-D cameras and 2D/3D LiDAR feeding segmentation, point-cloud reconstruction, 6-DoF grasp poses, and obstacle detection.
- Raised segmentation recall 15 points on unseen sites and cut merged/split boxes 40%. Rebuilt the dataset pipeline around visual clustering and negative mining based on field heuristics.
- Train and ship the segmentation and detection models the robots run: PointRend, DETR, SAM and ViT backbones in PyTorch and Detectron2. Redesigned the complete ML development architecture to be scalable and maintainable allowing developers to systemically evaluate and improve model performance.
- Cut inference time on the robot by 35% and got two models sharing one GPU, by fixing how frames are preprocessed and batched and when weights are held in memory. Exported through ONNX/TensorRT and quantized where accuracy allowed.
- Developed a per-site Adaptive learning loop that learns from operator corrections while the robot keeps unloading, then swaps the new weights asynchronously
- Took box dimension error from about 4 cm down to under 5 mm by measuring faces from mask pixels projected onto the fitted plane, improving the major pipeline algorithms and point cloud manipulation.
- Own calibration for the sensor rig: intrinsics, extrinsics, hand-eye. Solved as rigid-body transforms in SE(3), checked against fiducial targets, and developed an auto calibration pipeline.
- Wrote the point-cloud and 3D LiDAR layer: plane fitting, clustering, outlier rejection, occupancy mapping to construct a high accuracy World Model of the scene at a tight latency budget
- Developed strategies to tackle the common depth issues like dark surfaces drop out, depth edges throw flying pixels, and multipath interference. That work drives sensor choice and mounting geometry for the future hardware decisions.

Intuitive Surgical Inc.

May 2024 – Dec 2024

System Software Engineer Co-op

Sunnyvale, CA

- Improved optical flow accuracy by over 40% in low-texture endoscopic scenes by building a custom dataset from the robot's high-precision kinematic motion and fine-tuning a RAFT model for visual-motion verification.
- Cut manual validation effort by 15 hours per week and enabled proactive failure detection by automating the robot life-test pipeline and its data analysis in Python.

GAMMA AI Lab, University of Maryland

Jan 2024 – May 2024

Research Assistant

College Park, MD

- Improved pedestrian trajectory prediction accuracy ~8% over baseline at 2 ms inference latency by architecting a spatio-temporal graph convolutional network (ST-GCNN) conditioned on the robot's planned future path to model human-robot interaction.
- Validated forecasting performance in simulation and real-world robot experiments in crowded environments, demonstrating robustness and real-time viability on-robot.

CodelatticeLabs Pvt. Ltd.

May 2022 – July 2023

Robotics Software Engineer

Bengaluru, IN

- Reduced system latency ~5% and improved multi-robot localization accuracy in production by developing real-time C++ firmware for ESP32 microcontrollers and a submap-fusion algorithm under strict embedded constraints.
- Reduced simulated traffic collisions by 30% at intersections by implementing multi-agent coordination and trajectory tracking algorithms using reinforcement learning.

Projects

Multi-Sensor Fusion for Robust 3D Perception | Python, PyTorch, Sensor Fusion, LiDAR, CARLA

October 2024

- Engineered a deep sensor-fusion framework that beat a camera-only baseline by over 30 mAP in dense fog, using a Transformer-based BEV architecture (BEVFusion) to fuse camera and LiDAR features, validated on a custom adverse-weather CARLA dataset.

Self-Supervised Pre-Training for Multi-Modal 3D Perception | PyTorch, LiDAR, Self-Supervised

Sept. 2024

- Improved downstream 3D object detection by 10% while using 95% less labeled data, developing a self-supervised framework that pre-trains a perception backbone on cross-modal consistencies between unlabeled LiDAR and camera streams.

Vision-Based Localization with Particle Filtering | C++, Python, ROS2, Deep Learning

May 2024

- Developed a real-time visual localization and odometry system for aerial robots using a particle filter over learned CNN, VecKM, and Histogram-of-Features encoders, validated in a Gazebo PX4 software-in-the-loop environment.

3D Reconstruction using Structure from Motion | C++, OpenCV, Eigen, Open3D

March 2024

- Built a C++ Structure-from-Motion pipeline for 3D reconstruction, implementing feature matching, epipolar geometry and pose recovery, triangulation, and bundle adjustment for global optimization.

Technical Skills

Languages: C++ (11/14/17), Python, CUDA, MATLAB, SQL, Bash
3D Perception & Geometry: Point clouds (Open3D, PCL), RANSAC & plane fitting, registration/ICP, camera calibration (intrinsic, extrinsic, hand-eye), coordinate transforms (TF2), epipolar geometry, occupancy mapping, 6-DoF pose estimation, sensor fusion
Deep Learning & ML: PyTorch, Detectron2, Mask R-CNN, PointRend, DETR, SAM, DINOv2, ViT/Transformers, instance & semantic segmentation, detection & tracking, model training & evaluation (AP/recall), dataset curation, active learning, domain adaptation, ONNX, TensorRT, quantization, GPU inference
Robotics & Sensing: ROS1/ROS2, RGB-D cameras (Orbbec, RealSense, IFM3D), 2D/3D LiDAR, IMU, AprilTag, sensor synchronization, MoveIt, Gazebo, CARLA
Production & Tooling: Docker, GitHub Actions CI/CD, pytest & GoogleTest, rosbag replay harnesses, telemetry & observability (TimescaleDB, Grafana, CloudWatch), AWS/S3, Linux/ARM, Git

Publications

Behavioral Analysis of ROS Motion Planners Integrated with Robotics Middleware Framework (RMF) | IEEE, 2022
Evaluated the integration of the Robotics Middleware Framework with Free Fleet, benchmarking path-planning algorithms in multi-robot scenarios to improve autonomous mobile robot fleet management.

Education

University of Maryland, College Park <i>M.Eng. Robotics, GPA: 3.98/4</i>	Aug. 2023 – May 2025 <i>College Park, MD</i>
Ramdeobaba College of Engineering and Management <i>B.E. Mechanical Engineering, GPA: 9/10</i>	Aug. 2019 – May 2023 <i>Nagpur, IN</i>