Rohin Siddhartha

Website | GitHub | pvrohin@gmail.com | LinkedIn

Education

Worcester Polytechnic Institute — Worcester, MA
Master of Science – Mechatronics, Robotics, and Automation Engineering — *Aug 2022 – May 2024*

PSG Institute of Technology and Applied Research — Coimbatore, India
Bachelor of Engineering – Electronics and Communications Engineering — *2017 – 2021*

Technical Skills

Tools/Frameworks: PyTorch, TensorFlow, Keras, Scikit-Learn, OpenCV, ROS, Gazebo, RViz
Programming Languages: Python, C/C++, MATLAB
Tools & Platforms: Linux, Git, Docker, CUDA, cuDNN

Experience

Atlas Copco — Auburn Hills, MI
Computer Vision Engineer — *May 2024 – Present*

AI-Driven Quality Inspection

- Implemented and validated SGZ (Zero-Shot Low-Light Enhancement Network) using depthwise separable convolutions, achieving 93% detection accuracy on factory imagery and eliminating strobe lighting dependency — reducing power consumption, footprint, and cycle time across production lines.
- Led Joint Analytics, an AI vision system for SPR quality assessment using SAM-based segmentation for rivet geometry extraction, crack detection, and inclusion identification; built synthetic data pipelines (Gmsh, Plotly) with tolerance-aware scoring for BMW and JLR.
- Integrated Joint Analytics into Databricks (Jointstream), automating metric logging and anomaly detection and reducing manual annotation time from 5+ minutes to under 1 minute per joint.

Industrial Robotics & Vision Integration

- Spearheaded the commissioning of multi-robot cells, integrating 3D vision systems with FANUC Robots for autonomous guidance and inspection.
- Optimized robot-vision communication by modifying FANUC Karel logic, increasing program efficiency by 20% through improved bit reception.
- Engineered an automated perpendicularity correction system for Robotic Flow-Drill processes using real-time 3D vision feedback.

Pace Robotics — Bengaluru, India
Robotics Intern — *Jan 2022 – Jun 2022*

- Built a complete AMR simulation stack using URDF, Gazebo, and RViz in ROS2, integrating real-time LiDAR, Stereo Camera, and IMU sensor inputs as ROS topics for perception and autonomy experiments.
- Developed data collection and validation pipelines for AMR performance testing in lab and field environments, leveraging the simulation to benchmark robot autonomy algorithms across diverse deployment scenarios.

Indian Institute of Technology — Project Page Palakkad, India
Research Intern — *Dec 2020 – Jun 2021*

- Implemented Mask R-CNN-based fruit segmentation system achieving 79.23% precision, 81.2% recall, and 80.61% IoU, coupled with a RANSAC sphere fitting pipeline on extracted point clouds for accurate 3D fruit localization.
- Validated the full perception stack in Gazebo simulation with ROS integration, enabling closed-loop robotic control for autonomous fruit picking applications.

Projects

SfM and NeRF – Classroom project on CV aimed at building a traditional Structure from Motion pipeline from scratch and comparing with deep learning variants such as Neural Radiance Fields. Github

Panorama Stitching – Classroom project on CV aimed at stitching images using traditional and DL-based approaches including supervised and unsupervised homography via Spatial Transformer and DLT methods. Github

Robust Trajectory Tracking for Quadrotor UAV using Sliding Mode Control – Developed a robust control scheme to track desired trajectories in the presence of external disturbances. Interfaced the controller with Gazebo simulation of Crazyflie quadrotor using CrazyS and mav_comm ROS packages. Designed boundary layer-based sliding mode control laws for the z, ϕ, θ, ψ coordinates and achieved RMSE under 0.1 m. Github

Visual-Inertial Odometry (VIO) Using Extended Kalman Filter – Implemented EKF-based VIO system processing 100,000+ IMU and camera frames at 100Hz. Developed 15-state model with covariance estimation reducing position error by 30% through optimized sensor fusion. Validated system through extensive testing in simulation and real-world scenarios. Github

Mars Rover Prototype – Indian Rover Challenge – Led outdoor autonomous navigation development using stereo cameras and SLAM for obstacle avoidance. Developed a custom ROS package for skid-steering teleoperation and low-level motor control. Fused IMU and visual odometry using RTAB-Map SLAM for robust position estimation in challenging terrain. Github Project Page

INS-GNSS Integration with Unscented Kalman Filter for Autonomous Navigation – Developed INS-GNSS sensor fusion system achieving 0.08m position accuracy using a 15-state UKF with accelerometer/gyroscope bias estimation, a 60% improvement over feedforward. Implemented nonlinear state estimation for 9 navigation variables using sigma point filtering with WGS84 Earth and Somigliana gravity models. Validated navigation algorithms across 23,000+ timesteps at 1Hz, maintaining filter stability via haversine analysis. Github

Autonomous Multi-Trailer Parking Path Planning Optimization – Developed a path planning algorithm for multi-trailer parking, cutting solution time to 800s from 1500s. Designed a Hybrid RRT – Local A* algorithm using Reeds-Shepp metrics for improved trajectory generation and obstacle avoidance. Optimized simulations with dynamic graphing and Gaussian sampling for high-dimensional maneuvering. Github

Self-Driving Car – Mahindra & Mahindra Driverless Car Challenge – Aimed at building a self-driving car using a full scale Mahindra Reva electric vehicle. Executed Lane Detection and Imitation Learning with data augmentation. Implemented YOLO-V4 for Traffic Sign Detection using the GTSDB dataset, achieving 90.3% test accuracy. GithubProject Page

PD Control with Gazebo and ROS – Implemented a position controller for a 3-DOF robot in Gazebo, using ROS services to pass joint reference values and publishing high-frequency joint effort commands to drive joints to target positions via closed-loop PD control. Github

Goal-Driven Autonomous Navigation using Deep Reinforcement Learning – Implemented TD3 (Twin Delayed Deep Deterministic Policy Gradient) for goal-directed mobile robot navigation in Gazebo, using a multi-modal state representation combining 24-dim 3D LiDAR scans and 50x50 RGB camera inputs for simultaneous obstacle avoidance and goal reaching. Integrated TensorBoard for training visualization and monitored policy convergence across randomized environments in ROS Noetic. Github

End-to-End Transformer-Based Image Captioning – Built an end-to-end image captioning model using a SWIN Transformer encoder with a Transformer decoder in TensorFlow (62M parameters), achieving 70% accuracy through hyperparameter optimization with limited local compute. Github

Biofilm Cell Characterization from 3D Microscopy Images using Deep Learning –

- Engineered a multi-stage feature extraction pipeline for 3D CLSM microscopy volumes, reducing dimensionality by 99.84% (562,500 $\rightarrow$ 905 features) via axis-wise 2D projections and max-pooling compression, enabling efficient learning on high-dimensional biomedical data.
- Designed and trained a Fully Connected Neural Network (FCN) with hierarchical 5-fold cross-validation to predict 5 biophysical cluster properties (cell count, Euclidean diameter, aspect ratio, shape anisotropy, asphericity) from 3D image sub-volumes, achieving $\sim$99% accuracy and outperforming Random Forest and LASSO baselines.
- Benchmarked CNN, FCN, Random Forest, and PCA-LASSO regression models on a dataset of 1,550 labeled 3D sub-volumes across 4 matrix-disruption treatment conditions, demonstrating superior generalization of deep learning approaches for multi-output regression on volumetric scientific data.