Abhinav Reddy Kowkuntla

Boston, MA — kowkuntla.a@northeastern.edu — 617-943-9956
LinkedIn — Portfolio

Education

Northeastern University, Boston, MA Dec 2025
Master of Science in Robotics, Concentration: Electrical and Computer Engineering
Relevant Coursework: Robotics Sensing and Navigation, Mobile Robotics, Autonomous Field Robotics, Computer Vision

Amrita Vishwa Vidyapeetham, India Jun 2020
Bachelor of Technology in Electronics and Communications Engineering

Technical Skills

Computer Vision: Homography, Segmentation, Feature Matching, Factor Graphs, Bundle Adjustment, CLAHE.
Deep Learning: CNNs, Transformers (ViT, Attention), RNNs (LSTM, GRU), Generative Models (GANs, Diffusion), TensorRT Optimization.
Robotics: SLAM, Sensor Fusion, Localization, Navigation, MPPI, Motion Planning, Kalman/EKF/Particle Filters, Camera Calibration (Kalibr), ROS Bag Processing.
Programming: Python, C++, Linux, Data Structures and Algorithms, Machine Learning.
Libraries/Tools: PyTorch, NumPy, SciPy, Scikit-learn, MATLAB Simulink, ROS1, ROS2, Gazebo, GitHub, Jenkins, CI/CD, Docker. Git. Bash. version control, CANoe, CANape, DDS, Software Architecture.
Hardware: NVIDIA Jetson (Orin Nano), CUDA, LiDAR, Stereo Cameras, Embedded platforms, Edge Devices, SSH, GPU.

Work Experience

Northeastern Field Robotics Lab, Boston, MA Jan 2024 – Present
Graduate Research Assistant

- Designing a custom multimodal stereo setup integrating RGB and infrared/thermal cameras with unified calibration for robust perception across diverse weather conditions.
- Building end-to-end deep learning pipelines for dense correspondences enabling depth estimation, optical flow, and holistic scene understanding in challenging environments.
- Developed a lightweight alternative to RAFT-Stereo using cross-attention Transformers, achieving **14× faster inference (100 ms vs. 1.4 s)** and **50% lower memory usage (2.7 GB vs. 5.6 GB)** for real-time disparity estimation on Jetson Orin Nano.
- Supported calibration and ROS bag conversion pipelines for multi-sensor data collection and analysis.

Robert Bosch, India Jan 2021 – Dec 2023
Function Owner / Lead System Developer

- Led development of model-based automotive-grade software in Simulink and C++ for ADAS applications (Automatic Parking Systems), delivering over $2M in production features compliant with AUTOSAR, with cross-functional teams.
- Improved development efficiency by introducing code reusability and automating repetitive workflows, reducing release timelines by 1–2 weeks.

Projects

Multi-Robot SLAM and Multi-Modal Dataset Collection Spring 2025

- Deployed a fleet of five custom mobile robots equipped with RGB, LiDAR, and IMU payloads for large-scale multi-modal data collection.
- Designed and executed coordinated exploration missions ensuring spatial overlap and precise temporal synchronization across all platforms.
- Implemented hardware-level time synchronization for all sensors to ensure accurate cross-modality alignment and timestamp consistency.
- Collected over 1 TB of synchronized data enabling SLAM benchmarking and evaluation using frameworks such as RTAB-MAP, LeGO-LOAM, and LIO-SAM.

Multi-view 3D Reconstruction using Bundle Adjustment Fall 2024

- Reconstructed 3D point clouds from 24 monocular images using bundle adjustment and factor graph optimization in GTSAM.
- Utilized the Levenberg–Marquardt optimizer in GTSAM for nonlinear least-squares pose refinement and error minimization.
- Implemented SIFT feature extraction with non-maximum suppression and loop closure for improved geometric consistency.
- Reduced average pixel reprojection error from 1400 to 600 pixels through iterative optimization and pose adjustment.

GPS–IMU Fusion using Extended Kalman Filter Fall 2024

- Implemented a 2D pose estimation pipeline in Python by fusing GPS and IMU measurements using a custom Extended Kalman Filter.
- Defined process and measurement models, tuned covariance matrices, and validated consistency through trajectory evaluation.
- Achieved improved navigation accuracy, reducing position RMSE compared to individual sensor estimates.