

SIDDARTH DAYASAGAR

617-560-0540 | dayasagar.s@northeastern.edu | [LinkedIn](#) | [GitHub](#) | [Portfolio](#)

EDUCATION

Northeastern University, Boston, MA

Expected Dec 2026

Master's in Robotics

CGPA 3.76

Relevant Coursework:

Mobile Robotics, Robotics Mechanics and Control, Control Systems Engineering, Formal methods for dynamic systems, Legged Robotics, Pattern Recognition and Computer Vision

Jain University, Bangalore, India

June 2023

Bachelor of Technology in Robotics and Automation

CGPA 3.8

SKILLS

- **Programming:** Python, C++, MATLAB
- **Libraries:** OpenCV, PCL, ROS2 (humble, Jazzy), CasADi, Pinocchio
- **Simulation:** Gazebo Classic, Gazebo harmonic, Mujoco, IsaacSim, IsaacLab

WORK EXPERIENCE

Robotics Research Intern

Aug 2025 – Dec 2025

SpaceData Inc., Tokyo, Japan

- **Engineered** a PPO-based quadruped locomotion policy in Isaac Sim, improving performance by approximately **70% over the company's pretrained baseline** by refining rewards, optimising hyperparameters, and implementing domain randomisation.
- **Integrated** ROS 2 Nav2 with an **MPPI local planner** for indoor autonomous navigation, combining global path planning with multi-sensor fusion for **GPS-denied disaster-response** scenarios.
- Built an **outdoor navigation pipeline** using elevation-based grid mapping and an MPC local controller, extending quadruped deployment from controlled indoor environments to **unstructured multi-terrain** settings.
- **Architected Space Station OS**, a modular ROS 2 autonomy framework that combined subsystems into a unified control architecture, using behaviour trees, system diagnostics, and multi-node coordination to create a scalable, fault-tolerant robotics simulation platform.

Robotics System Integration Intern

Sept 2022 – Apr 2023

Flomobility, Bangalore, India

- **Integrated** a full-stack vision-based localisation pipeline using **AprilTag boards**, merging real-time pose estimation with the robot's navigation stack to **eliminate localisation drift**.
- **Implemented** the design and deployment of **Bang-Bang and PID control laws** for physical actuators, stabilising hardware response and **maximising system operational efficiency**.
- **Orchestrated** the enterprise-level **robotics software stack**, serving as the technical owner for **version control stability** and system-wide integrity to ensure a **robust and scalable production framework**.
- **Spearheaded** a "Product First" methodology by troubleshooting **sensor-to-actuator latency** bottlenecks, delivering **high-precision pose estimation** for commercial-grade autonomous robots.

PROJECTS

BHEEMA: Bipedal Humanoid Engineered for Extensible Manipulation and Autonomy

Jan 2026 - April 2026

- Implemented a **centroidal MPC locomotion controller** for the **Unitree G1 humanoid** using single rigid-body dynamics, OSQP as the QP solver, and friction cone constraints, running at **50 Hz**.
- **Developed** a 200 Hz **whole-body controller** using Pinocchio for analytical Jacobians, gravity compensation, and joint-space PD tracking, with operational-space impedance control for swing leg trajectories following the MIT Cheetah 3 framework.
- **Achieved** stable bipedal walking with alternating gait for 20+ seconds by resolving a coordinate frame mismatch between Pinocchio and MuJoCo, and trained a flat and martain terrain locomotion policy in MjLab, deploying it via ONNX Runtime for real-time inference

PRANA: Perception-condition Robotic network with attention

Feb 2026- Apr 2026

- **Designed** a flow-matching **vision-action policy** that iteratively denoises correlated noise into 50-step action trajectories, replacing deterministic L1 regression with a velocity-field formulation inspired by the 2025 BEHAVIOR Challenge first-place solution.
- **Iterated** across three architecture versions (**self-attention baseline, flow matching, DINOv2 backbone**), conducting ablation studies on vision encoders (ViT-Tiny, ConvNeXt, DINOv2)
- **Deployed** the trained policy on a **physical 7-DoF arm at 50 Hz**, with checkpoint-bound normalisation statistics and a deque-based action buffer that amortises inference cost across the full action chunk.

Learning-Based Extension of Classical Pick-and-Place using PPO

Nov 2025 - Dec 2025

- **Extended** a classical pick-and-place manipulation pipeline by training a **PPO** policy for a **7-DoF Franka Panda arm** in MuJoCo, using Cartesian end-effector action increments instead of explicit inverse kinematics.
- **Implemented** training pipeline in **PyTorch** with Gaussian policies, GAE, clipped objectives, and entropy regularisation for stable learning in continuous action spaces.
- **Designed** task-specific reward shaping and safety penalties (reach-to-object, joint limits, smooth motion, collision avoidance) and validated pick-and-place performance using MuJoCo.