Bavin Saravanan

bsaravan@andrew.cmu.edu
bavin-hub.github.io · github.com/bavin-hub
+1 412-616-5834

Education

Carnegie Mellon University — Pittsburgh, PA
M.S. in Mechanical Engineering; GPA: 3.92/4.0 — *Aug 2025 – Expected Dec 2026*
Relevant coursework: Dynamics and Control, Reinforcement Learning, Deep Learning.

SSN College of Engineering — Chennai, India
B.E. in Mechanical Engineering; GPA: 8.74/10.0 — *Jul 2019 – May 2023*
Relevant coursework: Mechanics, Introduction to Robotics, Mechatronics.

Research Interests

Data-efficient dynamics learning; planning under uncertainty; sim-to-real transfer; contact-rich robotics.

Publications and Manuscripts

[1] **LatNav: Monocular Obstacle Avoidance for Quadrotors via Collision Latents and Learned Odometry**.
Bavin Saravanan, Muqing Cao, Sebastian Scherer.
Under review, IEEE Robotics and Automation Letters (RA-L), 2026.

[2] **AM-Bench: A Modular Simulation Suite and Benchmark for Aerial Manipulation Policy Learning**.
Yutong Wang, Dongjae Lee, Xiaofeng Guo, Yuanzhu Zhan, Yufei Jiang, **Bavin Saravanan**, Muqing Cao, Jia Xie, Chenyang Mao, Sebastian Scherer, Junyi Geng, Guanya Shi.
Conference on Robot Learning (CoRL), 2026. [paper] [project]

[3] **Drone Soccer: Learning to Manipulate with Multicopter Downwash**.
Neelay Joglekar, **Bavin Saravanan**, Yutong Wang, Varun Kandiyappan, Junyi Geng, Sebastian Scherer.
IROS Workshop on Sim2Real and Classical Control, 2026 (workshop paper). [paper]

[4] **Fly-by-Code: Embodied Coding Agents for Aerial Manipulation with Active Visual and Physical Feedback**.
Jeongyeon Seo, Sihyun Cho, Jaewoo Lee, Gyeongrak Choe, Yutong Wang, **Bavin Saravanan**, Jia-Bin Huang, Furong Huang, Sebastian Scherer, Guanya Shi, H. Jin Kim, Seungjae Lee, Dongjae Lee.
Under review, International Conference on Learning Representations (ICLR), 2027. [project]

Research Experience

AirLab, CMU Robotics Institute — Pittsburgh, PA
Research Assistant; Advisor: Prof. Sebastian Scherer — *Nov 2025 – Present*

Monocular Quadrotor Navigation – LatNav

- Developed a navigation policy pairing RGB collision latents with learned odometry; deployed onboard a Starling 2 Max quadrotor using only an RGB camera.
- Modeled vehicle dynamics and delays for sim-to-real transfer, enabling flights of up to **12 minutes** across diverse environments.

Aerial Manipulation Benchmark – AM-Bench

- Built task environments and ran experiments across 12 tasks on a fully actuated hexarotor (FA-Hexa); the best fine-tuned policy averaged 52% success, with contact-rich tasks remaining challenging.
- Added a ground-effect model, reducing near-ground simulation-to-real end-effector vertical RMSE from 1.91 to 0.99 cm (**48%**) relative to simulation without ground effect.

Downwash-Based Manipulation – Drone Soccer

- Trained a ball-manipulation policy using a simplified downwash model; demonstrated **10 consecutive passes** with a human on hardware.
- Directed the ball to its target in three of four real-world setups varying drone, ball, and goal positions. Remaining redirection failures motivated further modeling of the interaction dynamics.

Legged Systems Group, CMU Robotics Institute — Pittsburgh, PA
Research Assistant; Advisor: Prof. Hartmut Geyer — *May 2026 – Present*

Learning-Based Prosthesis Control

- Proposed and deployed a reinforcement learning joint-torque policy for the CMU prosthetic leg, combining human motion priors with the user's intact-leg femur pose.
- Learned residual dynamics from **6 minutes of walking data** to correct simulated next-state predictions for policy optimization, enabling **20 continuous minutes** of level-ground walking.
- Currently extending the policy to anticipate gait changes under uncertainty in user intent and robot dynamics.

Industry Experience

Skytex Unmanned Aerial Solutions Bengaluru, India
AI and Robotics Engineer *Jul 2023 – Jul 2025*

Cactus Swarm – Decentralized VTOL Autonomy

- Developed vector-field guidance around aircraft turning limits after simulation mismatch limited RL control; supported **15 autonomous long-range flights** for a decentralized swarm of 12 VTOL UAVs.
- Built inter-UAV communication and onboard autonomy pipelines using ROS2, PX4, and Micro XRCE-DDS for flocking, foraging, and cooperative payload delivery.

Bumblebee – GPS-Denied Quadrotor Navigation

- Developed onboard navigation for GPS-denied forests using LiDAR SLAM, local A* planning, and a Lissajous search strategy on a Jetson Orin NX quadrotor.
- Implemented detection, tracking, and Siamese-network target verification in cluttered environments.

Projects

CrazyMusician – Drone Keyboard Playing Sep – Dec 2026
ROS2, PX4, ADRC *[code]*

- Developed the flight-control and deployment pipeline for a quadrotor keyboard-playing system; deployed active disturbance rejection control (ADRC) on a Starling 2 Max through a ROS2–PX4 offboard interface.
- Established the VOXL 2/PX4 deployment and communication pipeline; flight-tested a revised controller retaining PX4's velocity loop with acceleration feedforward, achieving stable hover in a motion-capture cage.

PX4-Gym – Learning-Based Quadrotor Control and Navigation Jun 2026 – Present
Isaac Lab, PyTorch, PPO, ROS2, PX4 *[code]*

- Extended a Torch-native PX4-style control pipeline in Isaac Lab for the Starling 2 Max, integrating vehicle and motor dynamics with PPO training at position, velocity, attitude, and body-rate control levels.
- Integrated depth/RGB perception and collision representations for obstacle navigation; built ROS2 policy-inference interfaces for PX4 offboard deployment.

Random Agent – Model-Based RL for Humanoid Locomotion Feb – May 2026
PyTorch, GRU World Model, MPPO, MuJoCo/MJLab *[code]*

- Developed a model-based reinforcement learning pipeline for Unitree G1 locomotion, updating a GRU world model from newly collected simulation data to predict robot dynamics from state–action histories.
- Implemented model-based PPO (MPPO), alternating dynamics-model updates with policy optimization on imagined rollouts; used replay-buffer histories to initialize predictions for velocity-command learning experiments.

Hierarchical Loco-Manipulation – Unitree G1 Sep – Dec 2025
Navigation; Isaac Lab, PPO *[code]*

- Developed the navigation component of a perception-driven humanoid system, connecting scene understanding and target following with downstream whole-body control.
- Integrated navigation with a PPO walk-and-reach policy trained in Isaac Lab using a distance-gated locomotion-to-reaching reward; the whole-body policy was evaluated in simulation.

Teaching Experience

Carnegie Mellon University

Teaching Assistant, ML-AI for Engineers *Fall 2026*
Support students in applying ML and AI methods to engineering problems through office hours and assignment feedback.

Teaching Assistant, ML-AI Project for Engineers *Spring 2026*
Mentored project teams on experiment design, model implementation, evaluation, and technical communication.

Teaching Assistant, Trustworthy AI *Fall 2025*
Helped students connect robustness, uncertainty, fairness, and evaluation concepts to practical ML workflows.

Technical Skills

Programming: Python, C++.

Learning and simulation: PyTorch, Isaac Lab, MuJoCo, Gymnasium; reinforcement learning, residual dynamics modeling, sim-to-real transfer.

Robotics and systems: ROS2, PX4, Micro XRCE-DDS, LiDAR SLAM, motion planning; Linux, Git, Docker.