Sourav Selvaraj

Worcester, MA | +1 508-373-6943 | sselvaraj@wpi.edu | github.com/souravselvaraj | linkedin.com/in/souravselvaraj

Research interests: contact-rich control and learning for legged robots — MPC/WBC, force-controlled loco-manipulation, humanoid locomotion, sim-to-real validation.

Education

Worcester Polytechnic Institute (GPA: 4.0/4.0) **Aug 2025 – May 2027 (expected)**
Master of Science in Robotics Engineering *Worcester, MA*

- **Coursework:** Robot Control, Motion Planning, Optimal Control, Legged Robotics, Robot Learning

National Institute of Technology, Tiruchirappalli **Dec 2021 – May 2025**
B.Tech in Instrumentation and Control Engineering *Tiruchirappalli, India*

- **Coursework:** Control Systems, Robotics, Signal Processing, Machine Learning, Embedded Systems

Research Experience

Graduate Researcher — ReSCoT-Q: Reference-Guided Surface-Frame Control **Aug 2025 – Present**
Advisor: Dr. Mahdi Agheli *ALMaS Research Group, WPI*

- Developed **ReSCoT-Q**, a surface-frame MPC and weighted whole-body controller (**OCS2** SQP, 100 Hz) for force-controlled tool sliding on a **Unitree Go1** with a body-mounted tool and force/torque sensor, resolving contact into normal-force regulation, tangential-slide tracking, and cross-track stabilization.
- Designed a confidence-gated estimator fusing a nominal surface prior with online contact-wrench and tool-motion measurements, and a contact-stability supervisor gating sliding during contact acquisition, force recovery, and gait-induced dropouts.
- Demonstrated **109.9 s** of sustained hardware contact spanning a 91.3° arc of a cylindrical test surface, with **0.60 N** normal-force RMS and **18.8 mm** tracking RMS error at a 10 N target, despite a $\sim$10% radius error in the surface prior.

Graduate Researcher — GaitNet: Learning-Based Acyclic Footstep Planning **Nov 2025 – May 2026**
Advisor: Dr. Mahdi Agheli *ALMaS Research Group, WPI*

- Co-developed (with O. Sullivan) a hybrid learning–control footstep planner for Unitree Go1 acyclic locomotion over irregular terrain, combining leg-specific foothold cost-map prediction with reinforcement-learning-based gait selection and model-based quadruped control.
- Owned the simulation deployment stack: ROS integration of the planner with OCS2 perceptive MPC (foot-placement and collision constraints) and GPU elevation mapping from onboard depth sensing.
- Implemented terrain-aware footstep candidate generation using robot state, commanded velocity, and raycast heightmaps, enabling real-time filtering and ranking of feasible leg-selection, foot-displacement, and swing-duration actions.
- Evaluated GaitNet in NVIDIA Isaac Lab across commanded velocities and missing-step terrain difficulties, achieving 69.4% survival (episodes completed without falling into terrain gaps) versus 25.6% for a single-leg motion-planner baseline.

Undergraduate Researcher — TOTO, Autonomous Mobile Robot **Sep 2024 – May 2025**
Advisor: Dr. D. Ezhilarasi *NIT Trichy, India*

- Built an autonomous mobile robot combining **ICP-based SLAM** with an **Iterated Kalman Filter** fusing wheel-encoder and RPLiDAR measurements for real-time mapping and localization.
- Achieved $\pm$2% velocity-tracking error on uneven terrain with closed-loop speed control.

Research Intern — Bio-mimetic Autonomous Underwater Vehicle (BAUV) **Mar 2023 – Sep 2023**
Advisor: Dr. Asokan Thondiyath *IIT Madras, India*

- Designed the Jetson Orin sensor-interface PCB (KiCad) and implemented **CAN** communication to Maxon BLDC propulsion drives for a tuna-inspired biomimetic AUV.

Manuscripts

[1] **S. Selvaraj** and M. Agheli, "ReSCoT-Q: Reference-Guided Surface-Frame Control for Quadruped Robots," first-author manuscript in preparation, targeting *IEEE Robotics and Automation Letters (RA-L)*.

[2] O. Sullivan, **S. Selvaraj**, and M. Agheli, "GaitNet: Learning-Based Acyclic Footstep Planning for Dynamic Quadruped Locomotion," manuscript in preparation, targeting *IEEE ICRA*.

Selected Research Projects

Code and demonstration videos: github.com/souravselvaraj

Emergent Ice-Skating Locomotion for Humanoids via Blade-Contact RL [video] **Jul 2026 – Present**

- Training a **Unitree G1** to ice-skate on passive 3 mm knife-edge blades with PPO in MuJoCo — no motion capture or imitation data: skating emerges from **anisotropic blade–ice contact** (elliptic friction cones) and a reward separating forward glide from penalized lateral slip.
- Engineered for hardware transfer: **asymmetric actor–critic** with privileged blade velocities (critic-only), a staged 0.15–1.2 m/s speed curriculum, ungated slip penalties closing reward-hacking exploits, and domain randomization over PD gains, inertia, payload, actuator delay, and blade mounting.
- In simulation, current policies track a 0.9 m/s command within **0.09 m/s** at **94% single-support**, holding lateral blade slip to **2–4 cm/s** and gliding up to 70% of stance.

Human-Like Bimanual Lifting for Humanoids via Residual AMP–PPO [code] **May 2026 – Jul 2026**

- Trained a **Unitree G1** to squat, grasp, and carry objects bimanually in MuJoCo via a **residual policy over a frozen locomotion actor**, with an **AMP discriminator** on human mocap as a capped style regularizer so imitation never outpays task completion.
- Stabilized training with **success-adaptive curricula** annealing a virtual object assist, an asymmetric actor–critic with privileged contact forces, noisy object-state observations, and domain randomization over object mass, friction, CoM, and base pushes.
- Closed successive **reward-hacking exploits** via gated posture penalties; converged policies consistently hold objects at goal height, centered within **3 cm**, with human-like posture (torso pitch inside the mocap prior's 2–13° range).

Conditional Flow Matching Policies for Robot Manipulation [code] **Apr 2026 – May 2026**

- Implemented a **Conditional Flow Matching (CFM)** visuomotor policy from scratch in **robomimic** — Diffusion Policy's receding-horizon architecture with iterative denoising replaced by a learned ODE velocity field integrated in **10 Euler steps**.
- Achieved **26.7 ms** per action chunk at inference, an **11**× speedup over the default 100-step DDPM Diffusion Policy on an identical 65M-parameter backbone; ran solver and inference-step ablations (Euler/midpoint, 10–50 steps) against Diffusion Policy baselines.
- Benchmarked on **5 robomimic tasks** (3 seeds, 50 rollouts each), matching or beating **BC-RNN** references on 4/5: **100%** lift, **94%** can, **84% vs. 71%** two-arm transport, **64%** tool_hang.

Acrobatic Quadrotor Flight using Geometric Control and L1 Adaptive Augmentation [code] **Oct 2025 – Dec 2025**

- Implemented an **SE(3)** geometric controller with **L1 adaptive** augmentation on a **Crazyflie 2.1** in ROS 2 simulation (hover, zero-radius flip, recovery), achieving **2.2 cm** position-tracking RMSE and rejecting up to 94% of force-step disturbances.

Autonomous Rover — ISRO Robotics Challenge 2024 **Dec 2022 – Aug 2024**

- Led a 10-member team building a 6-wheeled rocker-bogie autonomous rover with **ORB-SLAM** terrain mapping (Intel RealSense D435i), kinodynamic path planning, and custom position/speed-control electronics.

Additional projects: vision-based imitation learning — Behavior Cloning vs. Diffusion Policy on robosuite Pick-and-Place (code); task-and-motion planning — PDDL + OMPL block assembly on a Franka Panda.

Teaching Experience

Teaching Assistant — Legged Robotics (Graduate) **Summer 2026**
Instructor: Dr. Mahdi Agheli *Worcester Polytechnic Institute*

Laboratory Curriculum Developer **2026 – Present**
Robotics Engineering Department *Worcester Polytechnic Institute*

- Designing hands-on lab modules for the *Legged Robotics* (Dr. Agheli) and *Parallel Robotics* (Dr. McCann) courses.

Academic Service & Leadership

Reviewer **Jan – Feb 2026**
IEEE Transactions on Automation Science and Engineering (T-ASE)

- Reviewed one manuscript on data-driven predictive control for robotics automation.

Treasurer & Team Lead **Aug 2024 – May 2025**
Robotics and Machine Intelligence Club, NIT Trichy

- Led the club's ISRO Robotics Challenge team; organized robotics workshops, seminars, and competitions.

Technical Skills

- **Robotics Control:** MPC, Whole-Body Control, OCS2, Trajectory Optimization
- **Robot Learning:** Reinforcement Learning (PPO, AMP), Behavior Cloning, Diffusion Policy, PyTorch, robomimic
- **Simulation & Middleware:** MuJoCo, NVIDIA Isaac Lab, Gazebo, ROS/ROS2
- **Hardware:** Unitree Go1/G1, Force/Torque Sensing, Franka Panda, Crazyflie, Jetson Orin, Intel RealSense, CAN
- **Programming:** Python, C/C++, MATLAB, Docker