

# Sunny Nitin Deshpande

Portfolio: sunnydeshpande.com Github: github.com/SunnyDeshpande

Work Authorization: F-1 Visa | CPT-eligible for internships

Email: sunnydeshpande9900@gmail.com

LinkedIn: sunny-deshpande

## EDUCATION

- University of Illinois Urbana-Champaign** Urbana, IL  
MEng in Autonomy and Robotics | GPA: 3.85/4.00  
Aug 2025 - Dec 2026  
*Coursework:* Reinforcement Learning, Advanced Computer Vision, AV Safe Autonomy, Humanoid Robotics
- Singapore University of Technology and Design** Singapore  
BEng, Robotics Focus, Minor in CS | SHARP Honours, Global Distinguished Scholarship  
Sep 2021 - May 2025  
*Capstone:* Led 7-person team to deploy an autonomous curved glass-cleaning robot in live trials at Singapore Changi Airport.

## TECHNICAL SKILLS

- Languages:** Python (Primary), C++, C, MATLAB, CUDA
- Machine Learning & Robot Learning:** Machine Learning (ML), Deep Learning, Reinforcement Learning (PPO, SAC), Hierarchical RL, Multi-Agent RL, Curriculum Learning, Domain Randomization, Sim-to-Real, Imitation Learning (Behavior Cloning, DAgger), VLA Models (NaVILA, OpenVLA, HumanVLA), VLMs (CLIP, Moondream2, Gemini)
- Autonomy Stack (Perception, Planning, Controls):** Computer Vision (CV), Visual Odometry, SLAM, Sensor Fusion (LiDAR + RGB-D), Point Cloud Processing, YOLO, TensorRT | Kalman Filter, A\*, RRT\*, Trajectory Optimization, MPC/MPP1 | PID, Pure Pursuit, Stanley, Admittance Control
- Tools & Hardware:** ROS/ROS2, PyTorch, Isaac Lab, Isaac Gym, MuJoCo, Nav2, MoveIt, Docker, Git, Linux | Unitree G1 & Go1, Booster K1, UR5e, Polaris GEM e4, Jetson AGX Orin, Ouster OS1-128, OAK-D, ZED2

## WORK EXPERIENCE

- Hyundai Motor Group Innovation Centre Singapore** Singapore  
Robotics Fleet Software Engineer Intern, Robotics Center  
May 2024 - Sep 2024
  - Overhauled fleet communication for **200+ production logistics AMRs** from polling-based REST to an event-driven publish/subscribe architecture, **cutting command-response latency by 55%** (332 ms  $\rightarrow$  151 ms) and enabling real-time fleet telemetry that was previously unobservable to operators.
  - Replaced stop-and-wait intersection coordination with **velocity-profiled trajectory blending** using Bezier smoothing and adaptive-lookahead pure pursuit, **increasing multi-robot intersection throughput by 40%** after a staged rollout across the production fleet.
  - Integrated a traffic-aware fleet path planner into the multi-threaded fleet management stack with corridor congestion scoring, ETA-based rerouting, and kinematic feasibility validation, **lifting completed missions by 10%** via canary deployment.
- Venti Technologies** Singapore  
AV Simulation Engineer Intern, Planning & Control Team  
Sep 2023 - Dec 2023
  - Extended the high-fidelity AV simulator from single-trailer to **multi-trailer articulated ego vehicles** (Autonomous Prime Movers), computing real-time hitch angles and per-trailer dynamics through a modular OOP class hierarchy that lets new trailers be added with **zero code changes**.
  - Built an **automated brake performance testing framework** executing parameterized scenarios to calibrate simulated vehicle dynamics; participated in company-wide port fleet trials, triaging operational logs in a pipeline adopted post-internship for regression testing.
- A\*STAR Institute for Infocomm Research (I2R)** Singapore  
Robotics & AI Research Intern, Perception Team  
May 2019 - Aug 2019
  - Designed an **end-to-end CNN visual navigator** on a Pioneer 3-DX that consumed RGB-D imagery plus a pose-difference vector and directly regressed velocity commands, **replacing the classical SLAM control loop with a single learned model**; trained and benchmarked 50+ architectures.
  - Built a **custom data pipeline** that auto-generated waypoints and synchronized RGB-D, odometry, and command streams via approximate-time filtering, **producing 36,520 training samples with zero manual labeling** by leveraging the SLAM stack as ground truth.
  - Delivered an **80x costmap update acceleration** (0.125 Hz  $\rightarrow$  10 Hz) via voxel-grid downsampling and statistical outlier removal upstream of the local planner, enabling real-time avoidance of overhanging obstacles invisible to the 2D scan.

## SELECTED PROJECTS

- Hierarchical PPO for Dual Unitree G1 Humanoids | Isaac Lab** Oct 2025 - Present
  - Designed a **two-tier PPO hierarchy** for the 29-DOF Unitree G1: a low-level locomotion policy is trained through a **7-phase curriculum** (standing  $\rightarrow$  omnidirectional walking) across 30M simulation steps with 25+ shaped rewards and aggressive domain randomization on friction, mass, and actuator delay.
  - Stacked a goal-conditioned high-level navigation policy on top of the frozen locomotion primitives; two humanoids **converge from 10m apart to within 0.5m at 95% success**. Extending with a CNN height-scan terrain encoder and pursuing physical G1 access for sim-to-real transfer.
- NaVILA VLA Navigation on Physical Unitree Go1 | targeting CoRL 2026** Jan 2026 - Present
  - Deployed a frozen **8B Vision-Language-Action model** fully onboard a **physical Unitree Go1** quadruped with edge compute, mapping multi-view camera frames and a natural-language instruction directly to base-velocity commands, one of the first VLA navigation stacks running on-device on a real legged robot.
  - Adding a **token-level confidence-gated clarification layer** with a learned query token, fine-tuned through LoRA and DAgger, so the robot asks the human before acting under ambiguity, enabling **closed-loop conversational navigation** rather than blind instruction following.
- HumanVLA Loco-Manipulation with Unified CLIP Backbone** Oct 2025 - Dec 2025
  - Replaced HumanVLA's dual-encoder student (image + text) with a **frozen unified CLIP backbone** and small adapters, distilled via behavior cloning and DAgger across 615 simulation episodes, reaching **62.1% success on unseen humanoid loco-manipulation tasks** with reduced placement error.
  - Diagnosed a **catastrophic forgetting failure** where unfreezing CLIP collapsed performance to 0%; this ablation motivated the frozen-backbone strategy now used in the physical Go1 NaVILA deployment above.
- Contact-Gated Residual RL with OpenVLA: Sim-to-Real on Physical UR5e** Jan 2026 - May 2026
  - Trained a **SAC residual policy** on top of a frozen **OpenVLA** base for tight-tolerance USB-A insertion: the base model handles free-space approach, while a force-hysteresis gate hands off to a 100 Hz residual the moment contact is detected.
  - Pretrained the residual in simulation under a **cavity-tightening curriculum** with force/torque-noise, end-effector pose, and action-latency randomization; the force-noise-DR variant **transferred zero-shot to the physical UR5e**, completing real insertions without further tuning.
- AutoShield | ADAPT: Pedestrian Safety Stack on Polaris GEM e4** Oct 2025 - Present
  - Owned the full pedestrian perception-to-prediction pipeline running live on the physical Polaris GEM e4:** Ouster 128-channel LiDAR (DBSCAN + human-geometry filtering) fused with YOLOv11 stereo via sensor-strength weighting (LiDAR distance, camera bearing) feeds a **MID-style diffusion Transformer** (DDPM-trained, DDIM 10-step inference, joint multi-agent cross-attention) reaching **minFDE-20 of 0.529m at 5s** on Argoverse 2.
  - Validated in simulation a Torch-based **MPP1 motion planner** (K=600 rollouts, 3s horizon, 10 Hz, jointly optimizing steering and acceleration; **0.007m mean / 0.024m max** lateral error inside the 0.5m safety spec) and a **text-promptable goal-selection module** (YOLO-World / LangSAM open-vocabulary detectors fused with LiDAR returns to emit MPP1 goal poses from natural-language prompts); on-vehicle, AutoShield's earlier **3-state TTC safety FSM** (Stanley + PID, hard-brake override) achieved **91% success across diverse pedestrian scenarios** in field tests.

## PUBLICATIONS, AWARDS & LEADERSHIP

- Publications:** *Evaluating Visual Odometry Methods for Autonomous Driving in Rain*, IEEE CASE 2023 (benchmarked 7 VO/SLAM algorithms across 15+ rain sequences; proposed best-performing stereo variant). *Characterization of Focal EEG Signals: A Review*, Elsevier FGCS 2019 (LS-SVM over 52 nonlinear features, 87.93% accuracy on Bern-Barcelona).
- Awards & Recognition:** SUTD SHARP Honours (Singapore's University Research Programme); SUTD Global Distinguished Scholarship; SAUVC 2023 Finalist (*Software Lead*, SUTD AUV team); Best Design Award, Asia-Pacific Vex Robotics Championship.
- Leadership & Service:** Senior House Guardian, SUTD residential life (2022-25); Director of Student Relations, SUTD ROOT Student Government (2021-22); Infantry, Singapore Armed Forces, National Service (2019-21).