Sarvesh Prasanth Venkatesan

(213) 796-9207 | sarveshp@usc.edu | sarveshv219.github.io | linkedin.com/in/sarvesh-prasanth-v

EDUCATION

University of Southern California — Los Angeles, CA
M.S. in Electrical and Computer Engineering (Machine Learning) — *2025 – 2027*

Vellore Institute of Technology — Vellore, India
B.Tech. in Electronics and Communication Engineering — *2019 – 2023*

EXPERIENCE

PlusAI — Aug 2026 – Present
RL Planner Research Intern — *Santa Clara, CA*

- Developing an **RL safety layer** to improve planner trajectories in long-tail scenarios, trained online in closed-loop simulation.
- Built an event mining pipeline to extract planner failure scenarios from driving logs as training data for the safety layer.

Dynamic Robotics and Control Laboratory | University of Southern California — Dec 2025 – Present
Graduate Researcher | Advisor: Prof. Quan Nguyen — *Los Angeles, CA*

VLA Fine-Tuning for Humanoids

- Post-training **π0.5** for humanoid whole-body control via **SFT** followed by **RL fine-tuning** on long-horizon, multi-step tasks.

Visual Humanoid Loco-Manipulation

- Co-developed **ViBe**, a post-training framework adapting a frozen whole-body tracker to perceptive control using a pretrained vision backbone and LoRA adapters trained with **PPO**; zero-shot sim-to-real on Unitree G1 across four tasks.
- Built a **CLIP**-conditioned **vision-language motion generator** that produces task-directed reference motions for the ViBe tracker.
- Benchmarked teacher–student distillation against direct visual RL adaptation; deployed policies onboard via ONNX Runtime on Jetson Orin, reducing inference latency by ~40%.

Stochastic Robotics Lab | Indian Institute of Science — Jul 2024 – Aug 2025
Project Associate | Advisor: Prof. Shishir Kolathaya — *Bengaluru, India*

- Designed **GRoQ-LoCO**, a **recurrent-transformer generalist** locomotion policy that uses offline imitation learning to fuse diverse quadruped skills, achieving zero-shot transfer across 5+ unseen quadruped morphologies.
- Implemented **PPO-based RL locomotion** policies for the 70 kg Stoch5 quadruped; reduced **sim-to-real** gap via novel Ground Reaction Force and adaptive Raibert-inspired rewards, improving push recovery by **75–98%** on 30° low-traction inclines. [Videos]
- Implemented visual RL policies for locomotion on challenging terrains. [Code]

Mobile Robotics Lab | Indian Institute of Science — Sep 2023 – Jun 2024
Project Associate | Advisor: Prof. Debashish Ghose — *Bengaluru, India*

- Developed a centralized **UAV-UGV multi-agent** system for large-scale outdoor inspection, with the UAV executing coverage path planning and the UGV performing autonomous ground-level verification. [Video]
- Built onboard autonomy stacks on ROS2 using cuVSLAM for VIO, LiDAR-based obstacle avoidance, and multi-agent mission coordination in GPS-denied unstructured environments.
- Collaborated on **LocalizeSORT**, a multi-object tracking framework with world-coordinate association, achieving +15% HOTA and +30% prediction accuracy. [Paper]

PUBLICATIONS

- Lokesh Krishna, **Sarvesh Venkatesan**, An Zhang, Quan Nguyen. ViBe: Visual Behavior Adaptation for Perceptive Humanoid Whole-Body Control. *arXiv.*
- Narayanan PP, **Sarvesh Prasanth Venkatesan**, Srinivas Kantha Reddy, Shishir Kolathaya. GRoQ-LoCO: Generalist and Robot-agnostic Quadruped Locomotion Control using Offline Datasets. *arXiv.*
- Srihari Vemuru, Kumar Ankit, **Sarvesh Prasanth Venkatesan**, Debashish Ghose, Shishir Kolathaya. Yield Prediction and Counting Using World Coordinates. *In ASME IMECE-India 2026 Proceedings. ASME.*

PROJECTS

Manipulation VLA (Github)

- Trained a VLA model using a frozen LLaVA backbone with a **DiT**-based action head for multi-object manipulation in simulation.
- Benchmarked diffusion policy against flow matching on LIBERO, with flow matching achieving a higher task success rate and 2.5x inference speedup across diverse manipulation tasks.

Humanoid Locomotion (Github)

- Developed a scandot-based visual RL humanoid locomotion policy for terrain-adaptive walking. [Video]
- Trained a humanoid locomotion policy fusing egocentric depth and proprioception via **cross-modal transformer attention**, replacing privileged height scans.

Quadruped Motion Planning (Github)

- Implemented RRT* and optimization-based planners for quadruped navigation in MuJoCo, with collision checking and cost functions accounting for quadruped foothold constraints. [Video]
- Learned locomotion skills including stair climbing, slope traversal, and dynamic walking using PPO with terrain curriculum, integrated with global motion plans for hierarchical end-to-end autonomous navigation.

SKILLS

Programming: Python, C++.
ML: PyTorch, JAX, Reinforcement Learning, Imitation Learning, Diffusion, Flow Matching, VLA, sim-to-real, ONNX.
Robotics: ROS2, MoveIt, Isaac Sim, IsaacLab, MuJoCo, mjlab, OpenCV, PX4, AlpaSim, CARLA.
Tools: Docker, Git, AWS, WandB, Fusion360, LaTeX, Ubuntu Linux.