Daksh Adhar

+1(412)287-0335 | adhardaksh@gmail.com | linkedin.com/in/daksh-adhar | a-daksh.github.io

EDUCATION

Carnegie Mellon University – M.S. in Robotic Systems Development | GPA: 4.05/4 — August 2024 – May 2026
IIT Guwahati – B.Tech. in Engineering Physics | GPA: 8.27/10.00 — November 2020 – July 2024

Courses: Deep RL, Deep Learning, Learning for 3D Vision, Computer Vision, Optimal Control, Robot Mobility, Autonomy

SKILLS

Python, C/C++, JAX, PyTorch, TensorRT, MuJoCo, Isaac Gym, PyBullet, ROS 2, MoveIt2, Gazebo, Docker, Kubernetes

PROFESSIONAL EXPERIENCE

Foundation Robotics — San Francisco, CA, USA
AI Engineer, World Model Team — July 2026 – Present

- Developing a latent predictive **world-action model** coupling future visual latents with flow-based whole-body actions
- Fine-tuned the model on **~40 min** of teleop data to output whole-body actions, reaching **70% success** over 50 runs
- Built a pipeline converting **500 hrs** of heterogeneous human motion data into whole-body controller training labels
- Deployed policies on the humanoid via on-device **TensorRT** inference, owning the checkpoint-to-robot release path

1X Technologies — Palo Alto, CA, USA
AI Resident, Reinforcement Learning Team — May 2025 – August 2025

- Trained RL policies for **in-hand manipulation** on NEO hands, with domain randomization for sim-to-real transfer
- Designed **sim-to-sim** benchmarking metrics across Isaac Gym and MuJoCo to evaluate policy robustness
- Developed a **ROS 2 Humble C++** controller to deploy evaluated policies in real-time simulation and teleoperation

RESEARCH EXPERIENCE

MaxLab CMU — Pittsburgh, PA, USA
Research Assistant, Prof. Max Simchowitz — January 2026 – May 2026

- Built a MuJoCo benchmark for generative dexterous policies on a unified **22-DoF UR5e + Allegro** arm-hand system
- Built a Leap Motion teleoperation pipeline with QP-based arm IK and **neural finger retargeting** for data collection
- Trained flow-matching behavior cloning policies on **~10 min of robot data** as the base policy for online RL fine-tuning
- Raised success on pick-and-place task from **30% to 64%** with RL post-training in **~1.3k online episodes** (~7 hrs)

Biomimetic Robotics & Artificial Intelligence Laboratory — Guwahati, India
Research Assistant, Prof. Shyamanta M. Hazarika — January 2023 – May 2024

- Built a testing framework in **PyBullet** and **OpenAI Gym** to train RL policies for bionic hand power grasping
- Formulated reward functions and used **Soft Actor-Critic** algorithm to enable grasp-and-lift of deformable objects
- Applied **domain randomization** for sim-to-sim, achieving 38% slip reduction and 14% decrease in deformation

PROJECTS

Reinforcement Learning Post-training of a Flow-based VLA | Code | Report — January 2026 – May 2026

- Implemented **RL Token** on **SmolVLA**, compressing frozen VLM embeddings into one token for a **TD3** actor-critic
- Validated the token with **PCA** and **t-SNE**, where last-step embeddings separate success from failure episodes
- Raised **Meta World** peg-insert task success from **72% to 88%** over a BC baseline with no extra demonstrations

Hierarchical VLM for Vision-Language Navigation on Mobile Robots | Code | Report — September 2025 – October 2025

- Built an Input Handler to parse and route natural **language queries** for numerical, object reference, and navigation tasks
- Implemented **A*** exploration with loop-minimizing heuristics, achieving 82% map generation success across 17 scenes
- Enabled **spatial reasoning** via 3D-to-2D projection for object color association and geometric predicate prompting

AR-assisted Robotic Total Knee Arthroplasty | Website — October 2024 – December 2025

- Developed a surgical robot for total knee replacement, achieving **2 mm and 2°** accuracy in drilling surgical pins
- Utilized **DINOv3** for bone cluster detection, **SAM** for segmentation, and **ICP** for real-time surgical plan registration
- Built a **ROS 2/MoveIt2** collision-aware drilling planner with an **Apple Vision Pro** AR interface for full system control

PUBLICATIONS

CMU Vision-Language-Autonomy Challenge (Presentation) — October 2025
IROS 2025; AI Meets Autonomy: Vision Language and Autonomous Systems Workshop — IROS Workshop

Grasp force optimization as a Bilinear Matrix inequality problem: A Deep-learning approach — December 2023
6th National Conference on Multidisciplinary Design, Analysis and Optimization — arXiv/2312.05034

Reinforcement Learning-Based Bionic Reflex Control for Anthropomorphic Robotic Grasping — November 2023
Journal Paper at Cambridge University Press, Robotica — Cambridge/Robotica