Dhruv Metha Ramesh

New Brunswick, NJ | dhruv.metha@gmail.com | dhruvmetha.github.io
github.com/dhruvmetha | linkedin.com/in/dhruvmetha | dhruvmetha.github.io/scholar

Robot learning for contact-rich manipulation, sim-to-real reinforcement learning, legged navigation, and planning among movable obstacles.

Education

Ph.D. in Computer Science, Rutgers University, New Brunswick, NJ 2023 – expected May 2028
Advisors: Abdeslam Boularias and Kostas E. Bekris. GPA: 3.92/4.0.
Joined Rutgers in 2021 through the M.S. program and transferred into the Ph.D. program.

B.E. in Computer Science and Engineering, B.M.S. College of Engineering, Bengaluru, India 2015 – 2019
GPA: 8.8/10.0.

Research Experience

Ph.D. Candidate, Rutgers University, New Brunswick, NJ

SPAR: learned push ranking for navigation among movable obstacles *First author* *arXiv 2026*

- Trained a transformer with regression and ranking losses to score push actions for a mobile robot that must clear blocked paths, using about 1.3 million labeled outcomes from exhaustive two-push searches in MuJoCo.
- Cut median simulator calls 7.0× and planning time 3.8× against random search on 300 two-push problems. On 600 problems that need more pushes, raised success within 100 simulator calls from 20.2% to 50.7% without retraining.
- On a real differential-drive robot, the learned policy reached the goal in 37 of 40 trials without simulating candidate pushes, against 27 of 40 for random search and 2 of 40 for a planner that ignores movable objects.

Sim-to-real tight insertion with model-based control and residual RL *Co-first author* *ICRA 2025*

- Built the IsaacGym simulation setup for training a residual RL policy with PPO, a sparse success reward and a pose-noise curriculum, on top of a potential-field controller for plug insertion at tolerances under 1 mm.
- Led the real-robot development and deployment on a KUKA LBR iiwa14 with an RGB-D pose tracker, moving the policy zero-shot from a simulated Franka Panda. On five household objects unseen in training, including 2- and 3-prong chargers and a HAN connector, our method completed 49 of 50 insertions.
- Under 5 mm / 5° pose noise in simulation, our method raised insertion success from 3.28% with the controller alone to 96.25%, and trained in 2 to 3 hours on one GPU, against 8 to 10 hours reported for IndustReal.

PROBE: obstacle estimation from a quadruped's proprioception *First author* *ICRA 2025*

- Built a transformer that estimates the size, pose and movability of obstacles from a quadruped's joint positions, velocities, torques and pose alone, with no vision, including obstacles hidden behind a box the robot pushes.
- Trained locomotion and navigation policies with PPO in 4,096 parallel IsaacGym environments with domain randomization, built a 180K-trajectory dataset, and tested on 3,000 simulated scenes and 45 real trials with a Unitree Go1.

Reach-avoid probability estimation for robot controllers *Co-first author* *Under review*

- Two papers under review on reach-avoid probability estimation for black-box robot controllers under stochastic dynamics, with statistical guarantees.

Visuotactile insertion with a compliant hand, with Tel Aviv University *Co-author* *RA-L 2025*

- Built the simulator and the hardware deployment infrastructure for a transformer policy that combines tactile images and point clouds for insertion with a three-finger compliant hand, transferred zero-shot to a real arm.

KRAFT: kinodynamic replanning over approximate models *Co-author* *arXiv 2024*

- Built the test environments and ran the experiments, in MuJoCo and on a real MuSHR car, for a method that pairs sampling-based kinodynamic replanning with feedback control over models tuned by system identification.

Mono-STAR: single-camera scene tracking and reconstruction *Co-author* *ICRA 2023*

- Trained the optical flow model, using depth, and developed the baselines for a real-time single-camera 3D reconstruction system that handles fast motion, non-rigid deformation and topology change.

Publications

Conference and Journal Papers

[1] **Dhruv Metha Ramesh**, Aravind Sivaramakrishnan, Shreesh Keskar, Kostas E. Bekris, Jingjin Yu, Abdeslam Boularias. PROBE: Proprioceptive Obstacle Detection and Estimation while Navigating in Clutter. *IEEE International Conference on Robotics and Automation (ICRA), pp. 698–704, 2025.*
arXiv:2505.11848 | dhruvmetha.github.io/legged-probe | github.com/dhruvmetha/probe

[2] Isidoros Marougkas*, **Dhruv Metha Ramesh***, Joe H. Doerr, Edgar Granados, Aravind Sivaramakrishnan, Abdeslam

Boularias, Kostas E. Bekris. Integrating Model-based Control and RL for Sim2Real Transfer of Tight Insertion Policies. *IEEE International Conference on Robotics and Automation (ICRA), pp. 2102–2109, 2025.* arXiv:2505.11858 | dhruvmetha.github.io/insert

[3] Osher Azulay, **Dhruv Metha Ramesh**, Nimrod Curtis, Avishai Sintov. Visuotactile-Based Learning for Insertion with Compliant Hands. *IEEE Robotics and Automation Letters (RA-L), vol. 10, no. 4, pp. 4053–4060, 2025.* arXiv:2411.06408

[4] Haonan Chang, **Dhruv Metha Ramesh**, Shijie Geng, Yuqiu Gan, Abdeslam Boularias. Mono-STAR: Mono-camera Scene-level Tracking and Reconstruction. *IEEE International Conference on Robotics and Automation (ICRA), pp. 820–826, 2023.* arXiv:2301.13244

Preprints

[5] **Dhruv Metha Ramesh**, Tri-An Nguyen, Junshan Huang, Jingjin Yu, Kostas E. Bekris, Abdeslam Boularias. SPAR: Search-trained Push Action Ranker for Efficient Navigation Among Movable Obstacles. *arXiv preprint, 2026.* arXiv:2610.28362

[6] Aravind Sivaramakrishnan, Sumanth Tangirala, **Dhruv Metha Ramesh**, Edgar Granados, Kostas E. Bekris. KRAFT: Sampling-Based Kinodynamic Replanning and Feedback Control over Approximate, Identified Models of Vehicular Systems. *arXiv preprint, 2024.* arXiv:2409.11522

Under Review

[7] Sumanth Tangirala*, **Dhruv Metha Ramesh***, Ewerton R. Vieira, Abdeslam Boularias, Ying Hung, Kostas E. Bekris. Uncertainty-Aware Reach-Avoid Estimation with Statistical Guarantees.

[8] Sumanth Tangirala*, **Dhruv Metha Ramesh***, Ewerton R. Vieira, Abdeslam Boularias, Ying Hung, Kostas E. Bekris. GRAPE: Generative Reach-Avoid Probability Estimation for Black-Box Robot Controllers under Stochastic Dynamics.

*Equal contribution. Citations on Google Scholar.

Teaching

Rutgers University, Department of Computer Science, New Brunswick, NJ 2022 – present

- Instructor of record, CS460/560 Introduction to Computational Robotics, Fall 2026
- Instructor of record, CS440 Introduction to Artificial Intelligence, Fall 2024 and Summer 2026
- Teaching assistant, CS440 Introduction to Artificial Intelligence and CS562 Advanced Robotics
- Mentored undergraduate researchers Tri-An Nguyen, a SPAR co-author, and Riddhima Medatwal.

Industry Experience

Software Engineering Intern, then Software Engineer, Leanovate Solutions, Bengaluru, India 2018 – 2020

- Built a Bluetooth indoor navigation app and led ReactJS front-end work on a space management platform.

Skills

Software: Python, C++, PyTorch, IsaacGym, MuJoCo, ROS, Stable Baselines3
Methods: reinforcement learning, PPO, residual RL, curriculum learning, sim-to-real transfer, domain randomization, transformers, search-based planning
Robots and sensors: Unitree Go1, KUKA LBR iiwa14, MuSHR, OpenHand compliant hand, Robotiq 3-finger gripper, AllSight tactile sensors, Intel RealSense D435

Service

Reviewer for:

- Robotics: Science and Systems (RSS), 2023–2025
- IEEE International Conference on Robotics and Automation (ICRA), 2022–2026
- IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2022–2026
- Conference on Robot Learning (CoRL), 2024
- Conference on Neural Information Processing Systems (NeurIPS), 2024
- IEEE Robotics and Automation Letters (RA-L)