

DHRUV SHAH

SENIOR RESEARCH SCIENTIST @ GOOGLE DEEPMIND

CONTACT INFORMATION	Google DeepMind 1600 Amphitheatre Parkway Mountain View, CA USA 94043	Webpage: cs.berkeley.edu/~shah E-Mail: shah@cs.berkeley.edu Google Scholar
EDUCATION	University of California, Berkeley 2019 – 2024 <i>M.S. & Ph.D. in Electrical Engineering & Computer Science</i> Advisor: Prof. Sergey Levine; GPA: 4.0/4.0 Indian Institute of Technology, Bombay 2015 – 2019 <i>B.Tech. (with Honors) in Electrical Engineering; GPA: 9.54/10</i>	
HONORS AND AWARDS	Microsoft Future Leader in Robotics & AI 2024 Best Conference Paper Award × 2, Intl. Conference on Robotics & Automation (ICRA) 2024 Best Student Paper Award (Finalist) × 2, –”– 2024 Best Paper Award in Cognitive Robotics (Finalist) , –”– 2024 Best Paper Award in Robot Manipulation (Finalist) , –”– 2024 Best Systems Paper Award (Finalist) , Robotics: Science and Systems (RSS) 2022 Berkeley Fellowship, UC Berkeley (<0.2% of graduate applicants) 2019–24 National Academy of Engineering Award (INAE), India ×2 2019 & 2018	
REFEREED PUBLICATIONS	[1] SELFIE: Autonomous Self-Improvement with RL for Vision-Based Navigation around People <i>Conference on Robot Learning (CoRL) 2024 (Oral Presentation, 4.3%)</i> Noriaki Hirose, <i>Dhruv Shah</i> , Ajay Sridhar, Kyle Stachowicz, Sergey Levine [2] LeLaN: Learning A Language-conditioned Navigation Policy from In-the-Wild Video <i>Conference on Robot Learning (CoRL) 2024</i> Noriaki Hirose, Catherine Glossop, Ajay Sridhar, <i>Dhruv Shah</i> , Oier Mees, Sergey Levine [3] Mobility VLA: Multimodal Instruction Navigation with Long-Context VLMs and Topological Graphs <i>Conference on Robot Learning (CoRL) 2024</i> Zhuo Xu, et al. ... (22 authors) [4] Extreme Cross-Embodiment Learning for Manipulation and Navigation <i>Robotics: Science and Systems (RSS) 2024</i> <i>Berkeley DeepDrive Workshop 2024 (Invited Talk)</i> J. Yang, C. Glossop, A. Bhorkar, <i>Dhruv Shah</i> , Q. Vuong, C. Finn, D. Sadigh, S. Levine [5] GOAT: GO to Any Thing <i>Robotics: Science and Systems (RSS) 2024</i> T. Gervet [†] , M. Chang [†] , M. Khanna [†] , S. Yenamandra [†] , <i>Dhruv Shah</i> , T. Min, C. Paxton, D. Batra, R. Mottaghi, D. S. Chaplot, J. Malik	

- [6] **NoMaD: Goal Masked Diffusion Policies for Navigation and Exploration**
International Conference on Robotics and Automation (ICRA) 2024
Best Conference Paper Award (0.05%)
Best Student Paper Award (Finalist, 0.2%)
Best Paper Award in Cognitive Robotics (Finalist, 0.1%)
NeurIPS 2023 Workshop on Foundation Models for Decision-Making (Oral Presentation)
CoRL 2023 Workshop on Pre-Training for Robot Learning (Oral Presentation)
 Ajay Sridhar, *Dhruv Shah*, Catherine Glossop, Sergey Levine
- [7] **Open X-Embodiment: Robotic Learning Datasets and RT-X Models**
International Conference on Robotics and Automation (ICRA) 2024
Best Conference Paper Award (0.05%)
Best Student Paper Award (Finalist, 0.2%)
Best Paper Award in Robot Manipulation (Finalist, 0.1%)
CoRL 2023 Workshop Towards Generalist Robots (Oral Presentation)
 Open X-Embodiment Collaboration, et al. ... (150+ authors)
- [8] **Grounded Decoding: Guiding Text Generation with Grounded Models for Robot Control**
Advances in Neural Information Processing Systems (NeurIPS) 2023
 W. Huang, F. Xia, *Dhruv Shah*, D. Driess, A. Zeng, Y. Lu, P. Florence, I. Mordatch, S. Levine, K. Hausman, B. Ichter
- [9] **SACSoN: Scalable Autonomous Data Collection for Social Navigation**
IEEE Robotics and Automation Letters (RA-L) 2023
Conference on Robot Learning (CoRL) 2023 (Live Demo)
IROS 2023 Workshop on Social Robot Navigation (Spotlight Presentation)
 Noriaki Hirose, *Dhruv Shah*, Ajay Sridhar, Sergey Levine
- [10] **ViNT: A Foundation Model for Visual Navigation**
Conference on Robot Learning (CoRL) 2023 (Oral Presentation & Live Demo, 6.6%)
BayLearn Machine Learning Symposium 2023 (Oral Presentation, <8%)
Dhruv Shah[†], A. Sridhar[†], N. Dashora[†], K. Stachowicz, K. Black, N. Hirose, S. Levine
- [11] **Navigation with Large Language Models: Semantic Guesswork as a Heuristic for Planning**
Conference on Robot Learning (CoRL) 2023
Dhruv Shah[†], Michael Equi[†], Blazej Osinski, Fei Xia, Brian Ichter, Sergey Levine
- [12] **FastRLAP: A System for Learning High-Speed Driving via Deep RL and Autonomous Practicing**
Conference on Robot Learning (CoRL) 2023
 Kyle Stachowicz[†], *Dhruv Shah*[†], Arjun Bhorkar[†], Ilya Kostrikov, Sergey Levine
- [13] **HomeRobot: An Open Source Software Stack for Mobile Manipulation Research**
AAAI Fall Symposium: Unifying Representations for Robot Application Dev. 2023
 C. Paxton, A. Wang, B. Shah, B. Matulevich, *Dhruv Shah*, K. Yadav, S. Ramakrishnan, S. Yenamandra, Y. Bisk
- [14] **GNM: A General Navigation Model to Drive Any Robot**
International Conference on Robotics and Automation (ICRA) 2023
Dhruv Shah[†], Ajay Sridhar[†], Arjun Bhorkar, Noriaki Hirose, Sergey Levine
- [15] **ExAug: Robot-Conditioned Navigation Policies via Geometric Experience**

Augmentation

International Conference on Robotics and Automation (ICRA) 2023

Noriaki Hirose, *Dhruv Shah*, Ajay Sridhar, Sergey Levine

- [16] **Learning Robotic Navigation from Experience: Principles, Methods, and Recent Results**
Philosophical Transactions of the Royal Society of London: B 2022 (**Invited Paper**)
Sergey Levine, *Dhruv Shah*
- [17] **Offline Reinforcement Learning for Visual Navigation**
Conference on Robot Learning (CoRL) 2022 (**Oral Presentation, 65%**)
Dhruv Shah[†], A. Bhorkar[†], H. Leen, I. Kostrikov, N. Rhinehart, S. Levine
- [18] **LM-Nav: Robotic Navigation with Large Pre-Trained Models of Language, Vision, and Action**
Conference on Robot Learning (CoRL) 2022
BayLearn Machine Learning Symposium 2022 (**Oral Presentation, <8%**)
Dhruv Shah[†], Blazej Osinski[†], Brian Ichter, Sergey Levine
- [19] **ViKiNG: Vision-Based Kilometer-Scale Navigation with Geographic Hints**
Robotics: Science and Systems (RSS) 2022 (**Oral Presentation**)
Best Systems Paper Award (Finalist, <2%)
Dhruv Shah, Sergey Levine
- [20] **Value Function Spaces: Skill-Centric State Abstractions for Long-Horizon Reasoning**
International Conference on Learning Representations (ICLR) 2022
Dhruv Shah, Peng Xu, Yao Lu, Ted Xiao, Alex Toshev, Sergey Levine, Brian Ichter
- [21] **Hybrid Imitative Planning with Geometric and Predictive Costs for Off-road Environments**
International Conference on Robotics and Automation (ICRA) 2022
N. Dashora[†], D. Shin[†], *Dhruv Shah*, H. Leopold, D. Fan, A. Agha, N. Rhinehart, S. Levine
- [22] **Rapid Exploration for Open-World Navigation with Latent Goal Models**
Conference on Robot Learning (CoRL) 2021 (**Oral Presentation, 65%**)
ICLR 2021 Workshop on Never-Ending Reinforcement Learning (**Oral Presentation**)
Dhruv Shah, Benjamin Eysenbach, Nicholas Rhinehart, Sergey Levine
- [23] **ViNG: Learning Open-World Navigation with Visual Goals**
International Conference on Robotics and Automation (ICRA) 2021
Dhruv Shah, Benjamin Eysenbach, Gregory Kahn, Nicholas Rhinehart, Sergey Levine
- [24] **Aerial Manipulation Using Hybrid Force and Position NMPC Applied to Aerial Writing**
Robotics: Science and Systems (RSS) 2020
D. Tzoumanikas, F. Graule, Q. Yan, *Dhruv Shah*, M. Popovic, S. Leutenegger
- [25] **The Ingredients of Real World Robotic Reinforcement Learning**
International Conference on Learning Representations (ICLR) 2020 (**Spotlight Presentation, 4.1%**)
H. Zhu[†], J. Yu[†], A. Gupta[†], *Dhruv Shah*, K. Hartikainen, A. Singh, V. Kumar, S. Levine
- [26] **Swarm Aggregation without Communication and Global Positioning**
IEEE Robotics and Automation Letters (RA-L) 2019
International Conference on Robotics and Automation (ICRA) 2019
Dhruv Shah, Leena Vachhani

- [27] **Projection Design for Compressive Source Separation using Mean Errors and Cross-Validation**
International Conference on Image Processing (ICIP) 2019
 Dhruv Shah, Ajit Rajwade
- [28] **Designing Constrained Projections for Compressed Sensing: Mean Errors and Anomalies with Coherence**
Global Conference on Signal and Information Processing (GlobalSIP) 2018
 Dhruv Shah[†], Alankar Kotwal[†], Ajit Rajwade
[†] Equal Contribution

PRE-PRINTS

- [29] **Vision Language Models are In-Context Value Learners**
arXiv 2024
 Y. Ma, J. Hejna, A. Wahid, C. Fu, Dhruv Shah, J. Liang, Z. Xu, S. Kirmani, P. Xu, D. Driess, T. Xiao, J. Tompson, O. Bastani, D. Jayaraman, W. Yu, T. Zhang, D. Sadigh, F. Xia
- [30] **STEER: Flexible Robotic Manipulation via Dense Language Grounding**
arXiv 2024
 L. Smith, A. Irpan, M. G. Arenas, S. Kirmani, D. Kalashnikov, Dhruv Shah, T. Xiao
- [31] **Traversability-Aware Legged Navigation by Learning from Real-World Visual Data**
arXiv 2024
 H. Zhang, Z. Li, X. Zeng, L. Smith, K. Stachowicz, Dhruv Shah, L. Yue, Z. Song, W. Xia, S. Levine, K. Sreenath, Y. Liu
- [32] **Gen2Act: Human Video Generation in Novel Scenarios enables Generalizable Robot Manipulation**
arXiv 2024
 H. Bharadhwaj, D. Dwibedi, A. Gupta, S. Tulsiani, C. Doersch, T. Xiao, Dhruv Shah, F. Xia, D. Sadigh, S. Kirmani

INVITED TALKS

- Guiding Robotic Planning with Large Pre-Trained Models**
 Invited Speaker, VLM3 Workshop @ ICRA 2024 May 2024
- The Foundation Model Path to Open-World Robots**
 Microsoft Invited Speaker, University of Maryland April 2024
 Department Seminar, Columbia University April 2024
 –”, Purdue University April 2024
 –”, Massachusetts Institute of Technology March 2024
 –”, Princeton University March 2024
 –”, University of California, San Diego March 2024
 –”, University of California, Los Angeles February 2024
 –”, University of California, Berkeley February 2024
 –”, University of Michigan February 2024
- Learning General-Purpose Robot Navigation**
 Invited Speaker, ML4AD Workshop @ NeurIPS 2023 December 2023
 AirLab Seminar, Carnegie Mellon University November 2023
 Bay Area Robotics Symposium October 2023
 MILA Robot Learning Seminar, Universite de Montreal September 2023
 Bay Area Machine Learning Symposium October 2023
 Seminar Series, Vayu Robotics July 2023

	ARL DCIST PI Meeting, University of Pennsylvania	June 2023
	Intuitive Interfaces for Learning from Offline Data	
	Bay Area Robotics Symposium	October 2022
	Scientific Speaker Series, Wayve	September 2022
	Kilometer-Scale Navigation with Geographic Hints	
	ML Seminar, Toyota Research Institute	March 2022
	RACER Seminar, NASA Jet Propulsion Laboratory	March 2022
	Berkeley Deep Drive Seminar, UC Berkeley	February 2022
	Skill-Centric State Abstractions for Planning	
	Google Brain/DeepMind Open Research Talks	November 2021
	Learning to Explore Open-World Environments	
	Google Brain/DeepMind Open Research Talks	November 2021
PRESS COVERAGE	The Human v/s Robot Earth Rover Challenge @ IROS 2024	October 2024
	WIRED, VentureBeat, Morningstar, MarketWatch	
	Mobility VLA: Multimodal Instruction Following with Long-Context VLMs	July 2024
	TechCrunch, TechXplore, MarkTechPost, TeqnoVerse (UAE)	
	GOAT: GO to Any Thing	November 2023
	MarkTechPost, ITinAI (Singapore)	
	Open X-Embodiment: Robotic Learning Datasets and RTX Models	October 2023
	MIT Tech Review, IEEE Spectrum, VentureBeat, Tech Times, Synced Review (Canada), TechForge (UK), Analytics India Magazine (India)	
	FastRLAP: A System for Learning High-Speed Driving	May 2023
	TechXplore, SyncedReview (Canada), MarkTechPost, TechEBlog	
	GNM: A General Navigation Model to Drive Any Robot	December 2022
	MarkTechPost	
	LM-Nav: Robotic Navigation with Large, Pre-Trained Models	August 2022
	Two Minute Papers, Utmel (Hong Kong)	
	ViKiNG: Kilometer-Scale Exploration in the Real World	March 2022
	IEEE Spectrum, ZDNet, Wevolver (Netherlands)	
	DARPA RACER (JPL/UC Berkeley/MIT/GeorgiaTech)	January 2022
	IEEE Spectrum, Caltech News, DARPA News, The Defense Post	
	RECON: Rapid Exploration with Latent Goal Models	December 2021
	RSIP Vision (Israel)	
BLOG POSTS	Scaling up Learning Across Many Different Robot Types	October 2023
	Google DeepMind Blog	
	Extracting Skill-Centric State Abstractions from Value Functions	April 2022
	Google AI Blog	
	Learning to Explore the Real World with a Ground Robot	November 2021
	Berkeley AI Research (BAIR) Blog	
	The Ingredients of Real World Robotic Reinforcement Learning	April 2020
	Berkeley AI Research (BAIR) Blog	

TEACHING
EXPERIENCE

Guest Lecturer

Cornell University (CS 6758) Fall 2024
Princeton University (ECE 518) Fall 2024
University of Michigan (EECS 598) Spring 2024

Student Instructor

University of California, Berkeley
CS 182/282A: Deep Neural Networks Spring 2023
CS 285: Deep Reinforcement Learning Fall 2021
Indian Institute of Technology, Bombay
CS 101: Introduction to Programming 2016, 2019
MA 207: Partial Differential Equations 2018

ADVISING

Undergraduate Students

Ajay Sridhar (2022–24, BS @ UC Berkeley; **NSF GRFP, CRA Finalist**) → *PhD @ Stanford CS*
Nitish Dashora (2020–23, BS @ UC Berkeley; **NSF GRFP, Astronaut Sc.**) → *PhD @ MIT EECS*
Michael Equi (2022–23, BS @ UC Berkeley) → *Research Eng. @ Physical Intelligence*
Hrish Leen (2022–24, BS/MS @ UC Berkeley) → *PhD @ Georgia Tech Robotics*
Arjun Bhorkar (2021–, BS/MS @ UC Berkeley; **Siebel Scholar**) → *Research Eng. @ Bloomberg*

External Collaborators

Active: Duy Nguyen (2024–, PhD @ Princeton), Lihan Zha (2024–, PhD @ Princeton), Jianzhu Yao (2024–, PhD @ Princeton), Catherine Glossop (2023–, PhD @ Berkeley), Kyle Stachowicz (2022–, PhD @ Berkeley), Noriaki Hirose (2022–, Researcher @ Toyota Japan)

Past: Jonathan Yang (2023–24, PhD @ Stanford), Hongbo Zhang (2023–24, PhD @ CUHK), Blazej Osinski (2022–23, PhD @ Warsaw)

Interns @ Google DeepMind

2024: Homanga Bharadhwaj (CMU), Joey Hejna (Stanford), Rohan Sinha (Stanford), Laura Smith (UC Berkeley), Jonathan Yang (Stanford)

SERVICE

Conference Organization

Area Chair, International Conference on Robotics & Automation (ICRA) 2024 – Present

Workshop Organization

3rd Workshop on Language and Robot Learning @ CoRL 2024
Morphology-Aware Policy and Design Learning Workshop @ CoRL 2024
The Earth Rover Challenge @ IROS 2024
6th Workshop on Robot Learning @ NeurIPS 2023 (**Lead Organizer**)
2nd Workshop on Language and Robot Learning @ CoRL 2023 (**Lead Organizer**)
2nd Workshop on Learning from Diverse, Offline Data @ ICRA 2023
1st Workshop on Language and Robot Learning @ CoRL 2022 (**Lead Organizer**)
1st Workshop on Learning from Diverse, Offline Data @ RSS 2022

Peer Review

Robotics — CoRL, RSS, RA-L, ICRA, T-RO, AuRo, IROS, ISRR, Humanoids, IJRR
Machine Learning — ICLR, NeurIPS, ICML
Computer Vision — T-PAMI