Abhishek Gupta

University of Washington
Seattle
WA 98195

abhgupta@cs.washington.edu
+1(510)-857-4588
https://homes.cs.washington.edu/~abhgupta/
https://scholar.google.com/citations?user=1wLVDP4AAAAJ&hl=en

Employment

2022–present **Assistant Professor**, *Paul G Allen School of Computer Science, University of Washington*, Seattle.

2025–present **Visiting Scholar**, *Toyota Research Institute*.

2024–2025 **Visiting Professor**, *NVIDIA*.

2021–2022 **Postdoctoral Scholar**, *Massachusetts Institute of Technology*.
Collaborators: Russ Tedrake, Pulkit Agrawal

Education

2015–2021 **Ph.D., Computer Science**, *University of California*, Berkeley.
Advisors: Pieter Abbeel, Sergey Levine

2011–2015 **B.S., Electrical Engineering and Computer Sciences**, *University of California*, Berkeley.

Experience

Academic

2015–present **Graduate Student Researcher**, *Berkeley AI Research (BAIR) Lab*, UC Berkeley, Advisors: Pieter Abbeel and Sergey Levine.
Research Area: Deep Reinforcement Learning, Robotics

2014–2015 **Undergraduate Researcher**, UC Berkeley, Advisor: Pieter Abbeel.
Research Areas: Task and Motion Planning, Apprenticeship Learning

Professional

Summer 2019 **Research Intern**, *Google Brain Robotics*, Hosts: Karol Hausman, Vikash Kumar.
Worked with the Google Brain Robotics team to solve long horizon manipulation problems

Summer 2013 **Software Development Engineering Intern**, *Motorola Solutions Inc.*.

Summer 2012 **Web Development Engineering Intern**, *EdX Berkeley*.

Teaching

Winter 2026 **Instructor**, *CSE478:Autonomous Robotics*, University of Washington.

Winter 2025 **Instructor**, *CSE478:Autonomous Robotics*, University of Washington.

Fall 2024 **Instructor**, *CSE579: Intelligent Control and Optimization*, University of Washington.

Spring 2024 **Instructor**, *CSE542: Reinforcement Learning*, University of Washington.

Winter 2024 **Instructor**, *CSE478: Autonomous Robotics*, University of Washington.

Spring 2023 **Instructor**, *CSE571: Probabilistic Robotics*, University of Washington.

Winter 2023 **Instructor**, *CSE599G: Deep Robotic Learning Learning*, University of Washington.

Fall 2017 **Graduate Student Instructor**, *CS285: Deep Reinforcement Learning*, Instructor: Prof. Sergey Levine, UC Berkeley.

Spring 2016 **Graduate Student Instructor**, *CS188: Introduction to Artificial Intelligence*, Outstanding Graduate Student Instructor Award, Instructors: Prof. Pieter Abbeel, Prof. Anca Dragan, UC Berkeley.

Fall 2014 **Teaching Assistant**, *CS70: Discrete Math and Probability Theory*, Instructor: Prof. Umesh Vazirani, UC Berkeley.

Honors and Awards

2025 Runners up for the Madrona Prize at UW

2025 TRI Research Award

2025 Nominated for Best Paper Award at CoRL 2025

2025 Students (Chuning Zhu, Yunchu Zhang) win the Amazon AI Fellowship

2025 Best paper at ICML 2025 workshop on Building Physically Plausible World Models

2025 Student (Patrick Yin) awarded NSF Graduate Fellowship

2024 IEEE RAS Early Career Academic Award 2025

2024 Amazon Science Hub Research Award

2024 Finalist for Best Vision Paper at ICRA 2024

2023 Best Paper Award at CoRL 2023 Workshop on Learning Effective Abstractions for Planning

2023 Selected as a Toyota Research Institute Young Faculty Investigator

2023 Finalist for Best Systems Paper at RSS 2023

2023 Student (Chuning Zhu) awarded Amazon Science Hub Fellowship

2019 UC Berkeley Nominee for Google PhD Fellowship

2019 Best Paper Runners Up at the NeurIPS 2019 Meta-Learning Workshop

2018 Finalist for Best Reviewer at Conference on Robot Learning (CoRL) 2018

2018 Best Paper Award at the ICML 2018 Exploration in RL Workshop

2018 Best Paper Runners Up at the ICML 2018 LLARLA Workshop

2016 National Science Foundation Graduate Research Fellowship

2016 National Defense Science and Engineering Graduate Fellowship (*declined*)

2016 UC Berkeley Outstanding GSI Award

2015 EECS Berkeley Departmental Fellowship

2011 Edward Kraft Award for Freshmen

2009 KVPY Fellowship, Govt. of India

Talks and Presentations

2025 Invited Talk at Princeton Seminar

2025 Invited Lecture at ETH Zurich Robot Learning Class

2025 Invited Talk at Stanford SVL

2025 Invited Talk at Maven Robotics

2025 Invited Talk at RembeRL Workshop at CoRL

2025 Invited Talk at Cornell Robotics Seminar

2025 Invited Talk at Stanford Vision and Learning Lab

2025 Invited Talk at Deep RL Class at Sharif Institute of Technology

2025 Invited Panelist at NSF Workshop on Reinforcement Learning

2025 Invited Talk at MLSS in Arequipa, Peru (Upcoming)

2025 Invited Talk at Foster School AI Summit (Upcoming)

2025 Invited Lecture at EE 543 University of Washington (Upcoming)

2024 Invited Talk at USC Control Seminar

2024 Invited Talk at MSR Asia Seminar Series

2024 Invited Talk at NAE Frontiers of Engineering Symposium

2024 Invited Talk at Cornell Robotics Seminar

2024 Invited Talk at IEEE New Era World Leaders AI Summit

2024 Invited Talk at RL Beyond Rewards Workshop at RLC 2024

2024 Invited Tutorial at International Computer Vision Summer School (ICVSS)

2024 Invited Talk at Hitachi Research Forum

2024 Invited Talk at Princeton Conference of Information Science and Systems (CISS)

2023 Invited Talk at RSS Workshop on Robotics and AI: The Future of Industrial Assembly Tasks

2023 Invited Talk at RSS Workshop on Learning for Dexterous Manipulation

2023 Invited Talk at Seattle Mind and Machines Meeting

2023 Invited Guest lecture at CMU course 16-884: Deep Learning for Robotics

2023 Invited talk at UIUC Robotics Seminar

2022 Invited talk at NVIDIA Robotics

2022 Invited Guest lecture at MIT course 6.S898: Deep Learning

2022 Invited Guest lecture at MIT course 6.8200: Computational Sensorimotor Learning

2022 Invited talk at International Computer Vision Summer School (ICVSS)

2021 Invited talk at Cornell, University of Maryland, Universite de Montreal, University of North Carolina - Chapel Hill, University of Washington, Georgia Institute of Technology, University of Texas at Austin

2020 Invited talk at MIT CSAIL: Learning Systems for Dexterous Manipulation

2020 Guest lecture at NYU Deep RL course on Offline and Meta RL

2020 Presented at CMU RI Reading Group on Ingredients of Real World Robotic Reinforcement Learning

2020 Invited talk at ICLR Workshop: Beyond tabula rasa in reinforcement learning

2020 Invited talk at International Computer Vision Summer School(ICVSS) (Postponed)

2019 Talk at Google Alphabots seminar

2019 Invited talk at ReWORK Deep Reinforcement Learning Summit

2019 Guest lecture in EE 106 B - Introduction to Robotics at UC Berkeley

2019 Invited talk at IROS Workshop on "Same Goal, Different Approaches to Robotic Manipulation"

2018 Spotlight talk at NeurIPS 2018

2018 Invited talk at Uber AI Symposium

2018 Invited talk at OpenAI

2018 Invited talk at ReWORK Deep Robotic Learning Summit

2018 Talk at BAIR/BDD Seminar: Unsupervised Meta RL

2018 Contributed talk at ICML ERL workshop

2018 Contributed talk at ICML LLARLA workshop

2017 Invited talk at Google Brain on Multi Task and Multi Robot Transfer

Pre-prints and Working Papers(* denotes equal contribution)

[106] Sarvesh Patil, Mitsuhiko Nakamoto, Shashwat Saxena, Manan Agarwal, Giri Anantharaman, Cleah Winston, Jesse Zhang, Chaoyi Pan, Douglas Chen, Nai-Chieh Huang, Zeynep Temel, Oliver Kroemer, Hongkai Dai, Sergey Levine, **Abhishek Gupta**, Paarth Shah, and Max Simchowitz. Sample efficient full-finetuning of generative control policies. In submission, 2026, PDF

[105] Sriyash Poddar, Jacob Krantz, Matthew Chang, Xavier Puig, Roozbeh Mottaghi, Natasha Jaques, and **Abhishek Gupta**. Asymmetric on-policy distillation learns in-context exploration. In submission, 2026

[104] Chuning Zhu, Paarth Shah, Jose Barreiros, Krishnan Srinivasan, and **Abhishek Gupta**. Latent memory palace: Reasoning as autoregressive variational inference. In submission, 2026

[103] Arhan Jain, Mingtong Zhang, Kanav Arora, William Chen, Marcel Torne, Muhammad Zubair Irshad, Sergey Zakharov, Yue Wang, Sergey Levine, Chelsea Finn, Wei-Chiu Ma, Dhruv Shah, **Abhishek Gupta**, and Karl Pertsch. Polaris: Scalable real-to-sim evaluations for generalist robot policies. arXiv preprint, 2026, PDF

[102] Anthony Liang, Yigit Korkmaz, Jiahui Zhang, Minyoung Hwang, Abrar Anwar, Sidhant Kaushik, Aditya Shah, Alex S. Huang, Luke Zettlemoyer, Dieter Fox, Yu Xiang, Anqi Li, Andreea Bobu, **Abhishek Gupta**, Stephen Tu, Erdem Biyik, and Jesse Zhang. Robometer: Scaling general-purpose robotic reward models via trajectory comparisons. In submission, 2026

[101] Matthew M. Hong, Jesse Zhang, Anusha Nagabandi, and **Abhishek Gupta**. Turning the dial: Diffusion timestep-modulated rl for tunable exploration. In submission, 2026

[100] Jacob Levy, Tyler Westenbroek, Kevin Huang, Fernando Palafox, Patrick Yin, Shayegan Omidshafiei, Dong-Ki Kim, **Abhishek Gupta**, and David Fridovich-Keil. Simulation distillation: Pretraining world models in simulation for rapid real-world adaptation. In submission, 2026

[99] Jacob Berg, Chuning Zhu, Yanda Bao, Ishan Durugkar, and **Abhishek Gupta**. Semantic world models, 2025, PDF

[98] Siddhant Haldar, Lars Johannsmeier, Lerrel Pinto, **Abhishek Gupta**, Dieter Fox, Yashraj Narang, and Ajay Mandlekar. Point bridge: 3d representations for cross domain policy learning. In *arXiv*, 2026, PDF

[97] Pranav Teegavarapu, Sriyash Poddar, Rico Qi, Natasha Jaques, and **Abhishek Gupta**. Grounding robot learning from interventions in real humans. In submission, 2026

[96] Mateo Guaman Castro, Sidharth Rajagopal, Daniel Gorbatov, Matt Schmittle, Rohan Baijal, Octi Zhang, Rosario Scalise, Sidharth Talia, Emma Romig, Celso de Melo, Byron Boots, **Abhishek Gupta**, and Anqi Li. Vamos: A hierarchical vision-language-action model for capability-modulated and steerable navigation. In *arXiv*, 2025, PDF

[95] Nicklas Hansen, Iretiayo Akinola, Yijie Guo, Jie Xu, Bingjie Tang, Hao Su, Xiaolong Wang, **Abhishek Gupta**, Dieter Fox, and Yashraj Narang. Generalizable robotic insertion with world models. In *In Submission*, 2025

[94] Henri Fung, Jack Lowry, Boling Yang, Thomas Kaminsky, Joshua Smith, Maya Cakmak, and **Abhishek Gupta**. Swiperl: A reinforcement learning system for grasping in clutter through non-prehensile pre-grasp motion. In *In Submission*, 2024

[93] John D Co-Reyes, Suvansh Sanjeev, Glen Berseth, **Abhishek Gupta**, and Sergey Levine. Ecological reinforcement learning. In *ArXiv Preprint*, 2020, PDF

[92] Giulia Vezzani, **Abhishek Gupta**, Lorenzo Natale, and Pieter Abbeel. Learning latent state representation for speeding up exploration. In *ArXiv Preprint*, 2019, PDF

[91] **Abhishek Gupta***, Benjamin Eysenbach*, Chelsea Finn, and Sergey Levine. Unsupervised meta-learning for reinforcement learning. In *ArXiv Preprint*, 2018(Best paper runners up at Lifelong Learning a Reinforcement Learning Approach (LLARLA) workshop at ICML 2018), PDF

[90] Tuomas Haarnoja, Aurick Zhou, Kristian Hartikainen, George Tucker, Sehoon Ha, Jie Tan, Vikash Kumar, Henry Zhu, **Abhishek Gupta**, Pieter Abbeel, and Sergey Levine. Soft actor-critic algorithms and applications. In *ArXiv Preprint*, 2018, PDF

[89] Ashvin Nair*, **Abhishek Gupta***, Murtaza Dalal, and Sergey Levine. Accelerating online reinforcement learning with offline datasets. In *ArXiv Preprint*, 2021, PDF

Conference Publications(* denotes equal contribution)

[88] Quinn Pfeifer, Ethan Pronovost, Paarth Shah, Khimya Khetarpal, Siddhartha Srinivasa, and **Abhishek Gupta**. Difference-aware retrieval polices for imitation learning. In *ICLR*, 2026

[87] Entong Su, Tyler Westenbroek, Anusha Nagabandi, and **Abhishek Gupta**. Rfs: Reinforcement learning with residual flow steering for dexterous manipulation. In *ICLR*, 2026, PDF

[86] Patrick Yin, Tyler Westenbroek, Zhengyu Zhang, Ignacio Dagnino, Eeshani Shilamkar, Numfor Mbiziwo-Tiapo, Simran Bagaria, Xinlei Liu, Galen Mullins, Andrey Kolobov, and **Abhishek Gupta**. Emergent dexterity via diverse resets and large-scale reinforcement learning. In *ICLR*, 2026, PDF

[85] Jesse Zhang, Marius Memmel, Kevin Kim, Dieter Fox, Jesse Thomason, Fabio Ramos, Erdem Bıyık, **Abhishek Gupta**, and Anqi Li. Peek: Guiding and minimal image representations for zero-shot generalization of robot manipulation policies. In *ICRA*, 2026, PDF

[84] Kevin Huang, Rosario Scalise, Cleah Winston, Yunchu Zhang Ayush Agrawal, Rohan Baijal, Markus Grotz, Byron Boots, Benjamin Burchfiel, Hongkai Dai, Masha Itkina, Paarth Shah, and **Abhishek Gupta**. Using non-expert data to robustify imitation learning via offline reinforcement learning. In *ICRA*, 2026, PDF

[83] Yijie Guo, Iretiayo Akinola, Lars Johannsmeier, Hugo Hadfield, **Abhishek Gupta**, and Yashraj Narang. Using non-expert data to robustify imitation learning via offline reinforcement learning. In *ICRA*, 2026

[82] Bingjie Tang*, Iretiayo Akinola, Jie Xu, Bowen Wen, Dieter Fox, Gaurav Sukhatme, Fabio Ramos, **Abhishek Gupta**, and Yashraj Narang. Refinery: Active fine-tuning and deployment-time optimization for contact-rich policies. In *ICRA*, 2026

[81] Pranav Atreya, Karl Pertsch, Tony Lee, Moo Jin Kim, Arhan Jain, Artur Kuramshin, Clemens Eppner, Cyrus Neary, Edward Hu, Fabio Ramos, et al. Roboarena: Distributed real-world evaluation of generalist robot policies. In *Proceedings of the Conference on Robot Learning (CoRL 2025)*, 2025, PDF

[80] Yunchu Zhang, Shubham Mittal, Zhengyu Zhang, Liyiming Ke, Siddhartha Srinivasa, and Abhishek Gupta. Atk: Automatic task-driven keypoint selection for robust policy learning. In *CoRL*, 2025, PDF

[79] Andrew Wagenmaker, Mitsuhiko Nakamoto, Yunchu Zhang, Seohong Park, Waleed Yagoub, Anusha Nagabandi, Abhishek Gupta, and Sergey Levine. Steering your diffusion policy with latent space reinforcement learning. In *CoRL*, 2025 (Oral), PDF(Nominated for Best Paper at CoRL 2025)

[78] Marcel Torne Villasevil, Arhan Jain, Jiayi Yuan, Vidyaaranya Macha, Lars Lien Ankile, Anthony Simeonov, Pulkit Agrawal, and **Abhishek Gupta**. Robot learning with super-linear scaling. In *RSS*, 2025, PDF

[77] Chuning Zhu, Raymond Yu, Siyuan Feng, Benjamin Burchfiel, Paarth Shah, and **Abhishek Gupta**. Unified world models: Coupling video and action diffusion for pretraining on large robotic datasets. In *RSS*, 2025PDF(Best Paper at ICML 2025 workshop on Building Physically Plausible World Models)

[76] Hongchi Xia, Entong Su, Marius Memmel, Arhan Jain, Raymond Yu, Numfor Mbiziwo-Tiapo, Ali Farhadi, **Abhishek Gupta**, Shenlong Wang, and Wei-Chiu Ma. Drawer: Digital reconstruction and articulation with environment realism. In *CVPR*, 2025, PDF

[75] Michael Murray, **Abhishek Gupta**, and Maya Cakmak. Teaching robots with show and tell: Using foundation models to synthesize robot policies from language and visual demonstration. In *CoRL*, 2024, PDF

[74] Zoey Chen, Zhao Mandi, Homanga Bharadhwaj, Mohit Sharma, Shuran Song, **Abhishek Gupta**, and Vikash Kumar. Semantically controllable augmentations for generalizable robot learning. In *IJRR*, 2025, PDF

[73] Patrick Yin, Tyler Westenbroek, Simran Bagaria, Kevin Huang, Ching-An Cheng, Andrey Kolobov, and **Abhishek Gupta**. Rapidly adapting policies to the real-world via simulation-guided fine-tuning. In *ICLR*, 2025, PDF

[72] Marius Memmel, Jacob Berg, Bingqing Chen, **Abhishek Gupta**, and Jonathan Francis. Strap: Robot sub-trajectory retrieval for augmented policy learning. In *ICLR*, 2025, PDF

[71] Yi Li, Yuquan Deng, Jesse Zhang, Joel Jang, Marius Memmel, Caelan Reed Garrett, Fabio Ramos, Dieter Fox, Anqi Li, **Abhishek Gupta**, and Ankit Goyal. Hamster: Hierarchical action models for open-world robot manipulation. In *ICLR*, 2025, PDF

[70] Yijie Guo, Bingjie Tang, Iretiayo Akinola, Dieter Fox, **Abhishek Gupta**, and Yashraj Narang. Srsa: Skill retrieval and adaptation for robotic assembly tasks. In *ICLR*, 2025 (Spotlight), PDF

[69] Abhay Deshpande, Liyiming Ke, Quinn Pfeifer, **Abhishek Gupta**, and Siddhartha Srinivasa. Data efficient behavior cloning for fine manipulation via continuity-based corrective labels. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2024, PDF

[68] Karthikeya Vemuri, Alan Wu, Arnav Thareja, Zoey Chen, Ian Good, Jeffrey Lipton, and **Abhishek Gupta**. Duolingo: Dynamics utilization for online translation of actions. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2025

[67] Yancheng Liang, Daphne Chen, **Abhishek Gupta**, Simon Shaolei Du*, and Natasha Jaques*. Learning to cooperate with humans using generative agents. In *NeurIPS*, 2024, PDF

[66] Andrew Wagenmaker, Kevin Huang, Liyiming Ke, Byron Boots, Kevin Jamieson, and **Abhishek Gupta**. Overcoming the sim-to-real gap: Leveraging simulation to learn to explore for real-world rl. In *NeurIPS*, 2024, PDF

[65] DROID dataset team. Droid: A large-scale in-the-wild robot manipulation dataset. In *Robotics: Science and Systems (RSS)*, 2024, PDF

[64] Chuning Zhu, Xinqi Wang, Tyler Han, Simon Du, and **Abhishek Gupta**. Distributional successor features enable zero-shot policy optimization. In *NeurIPS*, 2024, PDF

[63] Sriyash Poddar, Yanming Wan, Hamish Ivison, **Abhishek Gupta**, and Natasha Jaques. Personalizing reinforcement learning from human feedback with variational preference learning. In *NeurIPS*, 2024, PDF(Best Paper Finalist at NeurIPS Workshop on Pluralistic Alignment, Spotlight at NeurIPS 2024)

[62] Marcel Torne, Anthony Simeonov, Zechu Li, April Chan, Tao Chen, **Abhishek Gupta***, and Pulkit Agrawal*. Reconciling reality through simulation: A real-to-sim-to-real approach for robust manipulation. In *Robotics: Science and Systems (RSS)*, 2024, PDF

[61] Qiuyu Chen, Aaron Walsman, Marius Memmel, Alex Fang, Karthikeya Vemuri, Alan Wu, Dieter Fox, and **Abhishek Gupta**. Urdformer: A pipeline for constructing articulated simulation environments from real-world images. In *Robotics: Science and Systems (RSS)*, 2024, Website

[60] Zhang-Wei Hong, Aviral Kumar, Sathwik Karnik, Abhishek Bhandwaldar, Akash Srivastava, Joni Pajarinen, Romain Laroche, **Abhishek Gupta**, and Pulkit Agrawal. Beyond uniform sampling: Offline reinforcement learning with imbalanced datasets. In *NeurIPS*, 2023, PDF

[59] Vikash Kumar, Rutav Shah, Gaoyue Zhou, Vincent Moens, Vittorio Caggiano, **Abhishek Gupta**, and Aravind Rajeswaran. Robohive: A unified framework for robot learning. In *NeurIPS datasets and benchmarks*, 2023, PDF

[58] Michael Murray, **Abhishek Gupta**, and Maya Cakmak. Learning to grasp in clutter with interactive visual failure prediction. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2024

[57] Liyiming Ke, Yunchu Zhang, Abhay Deshpande, Siddhartha Srinivasa, and **Abhishek Gupta**. Ccil: Continuity-based data augmentation for corrective imitation learning. In *Int. Conf. on Learning Representations (ICLR)*, 2024, PDF

[56] Athul Paul Jacob, **Abhishek Gupta**, and Jacob Andreas. Modeling boundedly rational agents with latent inference budgets. In *Int. Conf. on Learning Representations (ICLR)*, 2024, PDF

[55] Zhaoyi Zhou, Chuning Zhu, Runlong Zhou, Qiwen Cui, Abhishek Gupta, **Abhishek Gupta**, and Simon Shaolei Du. Free from bellman completeness: Trajectory stitching via model-based return-conditioned supervised learning. In *Int. Conf. on Learning Representations (ICLR)*, 2024, PDF(Oral at NeurIPS Workshop on Foundation Models for Decision Making)

[54] Marius Memmel, Andrew Wagenmaker, Chuning Zhu, Dieter Fox, Abhishek Gupta, and **Abhishek Gupta**. Asid: Active exploration for system identification and reconstruction in robotic manipulation. In *Int. Conf. on Learning Representations (ICLR)*, 2024, PDF(Oral Presentation)

[53] Meenal Parakh, Alisha Fong, Anthony Simeonov, Tao Chen, **Abhishek Gupta**, and Pulkit Agrawal. Lifelong robot learning with human assisted language planners. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2024, PDF

[52] Daniel Yang, Jacob Berg, Davin Tjia, Dima Damen, Pulkit Agrawal, and **Abhishek Gupta**. Rank2reward: Learning shaped reward functions from passive video. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2024

[51] Zichen Zhang, Yunshuang Li, Osbert Bastani, Abhishek Gupta, **Abhishek Gupta**, Dinesh Jayaraman, Yecheng Jason Ma, and Luca Weihs. Universal visual decomposer: Long-horizon manipulation made easy. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2024, PDF(Best Paper at Learning Abstractions for Planning Workshop at CoRL 2023), (Oral at NeurIPS Workshop on Foundation Models for Decision Making)

[50] Jianlan Luo, Zheyuan Hu, Charles Xu, You Liang Tan, Jacob Berg, Archit Sharma, Stefan Schaal, Chelsea Finn, **Abhishek Gupta**, and Sergey Levine. Serl: A software suite for sample-efficient robotic reinforcement learning. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2024, PDF

[49] Marcel Torne, Max Balsells, Zihan Wang, Samedh Desai, Tao Chen, Pulkit Agrawal, and **Abhishek Gupta**. Breadcrumbs to the goal: Goal-conditioned exploration from human-in-the-loop feedback. *NeurIPS*, 2023, PDF

[48] Boyuan Chen, Chuning Zhu, Pulkit Agrawal, Kaiqing Zhang, and **Abhishek Gupta**. Self-supervised reinforcement learning that transfers using random features. *NeurIPS*, 2023, PDF

[47] Chuning Zhu, Max Simchowitz, Siri Gadipudi, and **Abhishek Gupta**. Repo: Resilient model-based reinforcement learning by regularizing posterior predictability. *NeurIPS*, 2023, PDF(Spotlight)

[46] Max Balsells I Pamies, Marcel Torne Villasevil, Zihan Wang, Samedh Desai, Pulkit Agrawal, and **Abhishek Gupta**. Autonomous robotic reinforcement learning with asynchronous human feedback. *CoRL*, 2023, PDF

[45] H.J. Terry Suh, Glen Chou, Hongkai Dai, Lujie Yang, **Abhishek Gupta**, and Russ Tedrake. Fighting uncertainty with gradients: Offline reinforcement learning via diffusion score matching. *CoRL*, 2023, PDF

[44] Zheyuan Hu, Aaron Rovinsky, Jianlan Luo, Vikash Kumar, **Abhishek Gupta**, and Sergey Levine. Reboot: Reuse data for bootstrapping efficient real-world dexterous manipulation. *CoRL*, 2023, PDF

[43] Max Simchowitz, Kaiqing Zhang, and **Abhishek Gupta**. Tackling combinatorial distribution shift: A matrix completion perspective. *COLT*, 2023, PDF

[42] Aviv Netanyahu*, **Abhishek Gupta***, Max Simchowitz, Kaiqing Zhang, and Pulkit Agrawal. Learning to extrapolate: A transductive approach. In *Int. Conf. on Learning Representations (ICLR)*, 2023, PDF

[41] Sameer Pai, Tao Chen, Megha Tippur, Edward H. Adelson, **Abhishek Gupta***, and Pulkit Agrawal*. Tactofind: A tactile only system for object retrieval. *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2023, PDF

[40] Yunchu Zhang, Liyiming Ke, Abhay Deshpande, **Abhishek Gupta**, and Siddhartha S. Srinivasa. Cherry-picking with reinforcement learning. *Robotics: Science and Systems (RSS)*, 2023, PDF

[39] Yuqing Du, Olivia Watkins, Zihan Wang, Cédric Colas, Trevor Darrell, Pieter Abbeel, **Abhishek Gupta**, and Jacob Andreas. Guiding pretraining in reinforcement learning with large language models. *ICML*, 2023, PDF

[38] Zoey Qiuyu Chen, Sho Kiami, **Abhishek Gupta***, and Vikash Kumar*. Genaug: Retargeting behaviors to unseen situations via generative augmentation. *Robotics: Science and Systems (RSS)*, 2023, PDF(Nominated for Best Systems Paper at RSS 2023)

[37] Kelvin Xu, Zheyuan Hu, Ria Doshi, Aaron Rovinsky, Vikash Kumar, **Abhishek Gupta**, and Sergey Levine. Dexterous manipulation from images: Autonomous real-world RL via substep guidance. *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2023, PDF

[36] **Abhishek Gupta**, Corey Lynch, Brandon Kinman, Garrett Peake, Sergey Levine, and Karol Hausman. Demonstration-bootstrapped autonomous practicing via multi-task reinforcement learning. *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2023, PDF

[35] **Abhishek Gupta***, Aldo Pacchiano*, Yuexiang Zhai, Sham M. Kakade, and Sergey Levine. Unpacking reward shaping: Understanding the benefits of reward engineering on sample complexity. In *NeurIPS*, 2022, PDF

[34] Anurag Ajay*, **Abhishek Gupta***, Dibya Ghosh, Sergey Levine, and Pulkit Agrawal. Distributionally adaptive meta reinforcement learning. In *NeurIPS*, 2022, PDF

[33] Zoey Qiuyu Chen, Karl Van Wyk, Yu-Wei Chao, Wei Yang, Arsalan Mousavian, **Abhishek Gupta**, and Dieter Fox. Learning robust real-world dexterous grasping policies via implicit shape augmentation. In *Conference on Robot Learning (CoRL)*, 2022, PDF

[32] Archit Sharma, Kelvin Xu, Nikhil Sardana, **Abhishek Gupta**, Karol Hausman, Sergey Levine, and Chelsea Finn. Autonomous reinforcement learning: Formalism and benchmarking. In *Int. Conf. on Learning Representations (ICLR)*, 2022, PDF

[31] Charles Sun, Coline Devin, Jedrzej Orbik, Glen Berseth, **Abhishek Gupta**, and Sergey Levine. Fully autonomous real-world reinforcement learning with applications to mobile manipulation. *CoRL*, 2021, PDF

[30] Marvin Zhang, Henrik Marklund, **Abhishek Gupta**, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift. In *NeurIPS*, 2021, PDF

[29] Kate Rakelly, **Abhishek Gupta**, Carlos Florensa, and Sergey Levine. Which mutual-information representation learning objectives are sufficient for control? In *NeurIPS*, 2021, PDF

[28] Archit Sharma, **Abhishek Gupta**, Sergey Levine, Karol Hausman, and Chelsea Finn. Autonomous reinforcement learning via subgoal curricula. In *NeurIPS*, 2021, PDF

[27] Olivia Watkins, Trevor Darrell, Pieter Abbeel, Jacob Andreas, and **Abhishek Gupta**. Teachable reinforcement learning via advice distillation. In *NeurIPS*, 2021, PDF

[26] **Abhishek Gupta***, Kevin Li*, Vitchyr Pong, Ashwin Reddy, Aurick Zhou, Justin Yu, and Sergey Levine. Mural: Meta-learning uncertainty-aware rewards for outcome-driven reinforcement learning. In *ICML*, 2021

[25] **Abhishek Gupta***, Justin Yu*, Tony Zhao*, Vikash Kumar*, Kelvin Xu, Thomas Devlin, Aaron Rovinsky, and Sergey Levine. Reset-free reinforcement learning via multi-task learning: Learning dexterous manipulation behaviors without human intervention. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2021

[24] Dibya Ghosh*, **Abhishek Gupta***, Ashwin Reddy, Justin Fu, Benjamin Eysenbach, Coline Devin, and Sergey Levine. Learning to reach goals via iterated supervised learning. In *Int. Conf. on Learning Representations (ICLR)*, 2021, PDF(Oral at ICLR 2021)

[23] Tianhe Yu, Saurabh Kumar, **Abhishek Gupta**, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *NeurIPS*, 2020, PDF

[22] Aviral Kumar, **Abhishek Gupta**, and Sergey Levine. Discor: Corrective feedback in reinforcement learning via distribution correction. In *NeurIPS*, 2020, PDF(Spotlight at NeurIPS 2.9% acceptance)

[21] Henry Zhu*, Justin Yu*, **Abhishek Gupta***, Dhruv Shah, Kristian Hartikainen, Avi Singh, Vikash Kumar, and Sergey Levine. The ingredients of real-world robotic reinforcement learning. In *Int. Conf. on Learning Representations (ICLR)*, 2020, PDF (Spotlight at ICLR 2020)

[20] **Abhishek Gupta**, Vikash Kumar, Corey Lynch, Sergey Levine, and Karol Hausman. Relay policy learning: Solving long-horizon tasks via imitation and reinforcement learning. In *Conference on Robot Learning (CoRL)*, 2019, PDF

[19] Michael Ahn, Henry Zhu, Kristian Hartikainen, Hugo Ponte, **Abhishek Gupta**, Sergey Levine, and Vikash Kumar. Robel: Robotics benchmarks for learning with low-cost robots. In *Conference on Robot Learning (CoRL)*, 2019, PDF

[18] Russell Mendonca, **Abhishek Gupta**, Rosen Kralev, Pieter Abbeel, Sergey Levine, and Chelsea Finn. Guided meta-policy search. In *NeurIPS*, 2019 (Spotlight at NeurIPS 2.5% acceptance, Oral presentation at Workshop on Multi-Task and Lifelong Learning at ICML 2019), PDF

[17] Allan Jabri, Kyle Hsu, **Abhishek Gupta**, Benjamin Eysenbach, Sergey Levine, and Chelsea Finn. Unsupervised curricula for visual meta-reinforcement learning. In *NeurIPS*, 2019 (Spotlight at NeurIPS 2.5% acceptance)PDF

[16] John D. Co-Reyes, **Abhishek Gupta**, Suvansh Sanjeev, Nick Altieri, Jacob Andreas, John DeNero, Pieter Abbeel, and Sergey Levine. Guiding policies with language via meta-learning. In *Int. Conf. on Learning Representations (ICLR)*, 2019(Best paper runner up at Workshop on Meta-Learning at NeurIPS 2018), PDF

[15] Dibya Ghosh, **Abhishek Gupta**, and Sergey Levine. Learning actionable representations with goal-conditioned policies. In *Int. Conf. on Learning Representations (ICLR)*, 2019, PDF

[14] Benjamin Eysenbach, **Abhishek Gupta**, Julian Ibarz, and Sergey Levine. Diversity is all you need: Learning skills without a reward function. In *Int. Conf. on Learning Representations (ICLR)*, 2019, PDF

[13] Michael Chang, **Abhishek Gupta**, Thomas Griffiths, and Sergey Levine. Automatically composing representation transformations as a means for generalization. In *Int. Conf. on Learning Representations (ICLR)*, 2019, PDF

[12] **Abhishek Gupta***, Henry Zhu*, Aravind Rajeswaran, Sergey Levine, and Vikash Kumar. Dexterous manipulation with deep reinforcement learning: Efficient, general, and low-cost. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2019, PDF

[11] **Abhishek Gupta**, Russell Mendonca, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Meta-reinforcement learning of structured exploration strategies. In *NeurIPS*, 2018 (Best paper at Exploration in RL workshop at ICML 2018, Spotlight at NeurIPS 3.4% acceptance), PDF

[10] John D. Co-Reyes*, Yuxuan Liu*, **Abhishek Gupta***, Benjamin Eysenbach, Pieter Abbeel, and Sergey Levine. Self-consistent trajectory autoencoder: Hierarchical reinforcement learning with trajectory embeddings. In *ICML*, 2018, PDF

[9] YuXuan Liu*, **Abhishek Gupta***, Pieter Abbeel, and Sergey Levine. Imitation from observation: Learning to imitate behaviors from raw video via context translation. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2018, PDF

[8] Aravind Rajeswaran*, Vikash Kumar*, **Abhishek Gupta**, John Schulman, Emanuel Todorov, and Sergey Levine. Learning complex dexterous manipulation with deep reinforcement learning and demonstrations. In *Robotics: Science and Systems (RSS)*, 2018, PDF

[7] **Abhishek Gupta***, Coline Devin*, YuXuan Liu, Pieter Abbeel, and Sergey Levine. Learning invariant feature spaces to transfer skills with reinforcement learning. In *Int. Conf. on Learning Representations (ICLR)*, 2017, PDF

[6] Coline Devin*, **Abhishek Gupta***, Trevor Darrell, Pieter Abbeel, and Sergey Levine. Learning modular neural network policies for multi-task and multi-robot transfer. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2017, PDF

[5] **Abhishek Gupta**, Clemens Eppner, Sergey Levine, and Pieter Abbeel. Learning dexterous manipulation for a soft robotic hand from human demonstration. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2016, PDF

[4] Rohan Chitnis, Dylan Hadfield-Menell, **Abhishek Gupta**, Siddharth Srivastava, Edward Groshev, Christopher Lin, and Pieter Abbeel. Guided search for task and motion plans using learned heuristics. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2016, PDF

[3] Siddharth Srivastava, Shlomo Zilberstein, **Abhishek Gupta**, Pieter Abbeel, and Stuart Russell. Tractability of planning with loops. In *AAAI conference on Artifical Intelligence*, 2015, PDF

[2] Alex Lee, Henry Lu, **Abhishek Gupta**, Sergey Levine, and Pieter Abbeel. Learning force-based manipulation of deformable objects from multiple demonstrations. In *Proc. IEEE Int. Conf. Robotics and Automation (ICRA)*, 2015,PDF

[1] Alex X Lee, **Abhishek Gupta**, Henry Lu, Sergey Levine, and Pieter Abbeel. Learning from multiple demonstrations using trajectory-aware non-rigid registration with applications to deformable object manipulation. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, 2015, PDF

Research Mentorship

Postdoctoral Tyler Westenbroek, Jesse Zhang

Graduate	Liyiming Ke (Advisor), Mateo Castro (Advisor), Yunchu Zhang (Advisor), Patrick Yin (Advisor), Sriyash Poddar (Advisor), Chuning Zhu (Advisor), Marius Memmel (Advisor), Qiuyu Chen (Advisor), Entong Su (Advisor), Arhan Jain (Advisor), Eric Cai (Advisor), Rohan Baijal, Rosario Scalise, Kevin Huang, Jack Lowry, Boling Yang, Michael Murray, Joel Jang, Daphne Chen, Zihan Wang, Aviral Kumar, John D Co-Reyes, Michael B. Chang, Giulia Vezzani, Kelvin Xu, Qiyang Li, Olivia Watkins, Andi Peng, Aviv Netanyahu, Boyuan Chen, Anurag Ajay, Tao Chen, Jesse Zhang, Nicklas Hansen, Bingjie Tang, Mitsuhiko Nakamoto, Harine Ravichandran
Undergraduate	Kevin Li, Aaron Rovinsky, Tony Zhao (now Ph.D student at Stanford), Ashwin Reddy, Justin Yu, Henry Zhu (now Ph.D student at Stanford University), Dibya Ghosh (now Ph.D student at UC Berkeley), Russell Mendonca (now Ph.D student at CMU), Xinyi Ren (now at Google), YuXuan Liu (now Ph.D student at UC Berkeley), Suvansh Sanjeev (now Ph.D student at CMU), Thomas Devlin, Siri Gadipudi, Shosuke Kiami, Vidyaaranya Macha, Kai Kwan Fung (now at Apple), Davin Tjia, Tejoram Vivekanandan, Alan Wu, Karthikeya Vemuri, Samedh Desai, Zheyuan Hu (now Ph.D. student at CMU), Sameer Pai, Meenal Parakh (now Ph.D. student at Princeton), Marcel Torne Villasevil (now Ph.D. student at Stanford), Max Balsells e Palmies(now masters student at ETH Zurich), Arhan Jain (now Ph.D. student at UW), Raymond Yu, Jacob Berg (now masters student at UW), Rico Qi, Pranav Teegavarapu, Quinn Pfeiffer, Simran Bagaria, Cleah Winston, Eeshani Shilamkar, Will Huey (Visitor), Matthew Hong (Visitor), Shubham Mittal, Aditya Shah
Research Scientist	Kanta Hamasaki (Hitachi), Numfor Tiapo
Alumni	Liyiming Ke (Physical Intelligence), Qiuyu Chen (Stealth Startup)

Grant Writing and Funding Awards

2025/26	Toyota Sponsored Research
2025/26	Hitachi Sponsored Research
2025/26	Boeing Sponsored Research
2025/26	Field AI Gift Award
2025/26	MSR Gift Award
2025/26	Amazon Sponsored Research Award
2024/25	Amazon Science Hub Research Award
2024/25	Sponsored Research Agreement with Army Research Lab
2024/25	Sponsored Research Agreement with Hitachi
2024/25	AI@UW Seed Grant from UW eScience Institute
2024/25	Part of TRI University Partnership 2.0 - Young Faculty Investigator
2022-24	Part of Amazon Science Hub Sponsored Project: Robotic Manipulation in Densely Packed Containers
2022	Awarded NSF RI medium award on Bootstrapping Natural Feedback for Reinforcement Learning with Jacob Andreas (MIT)

2020 Awarded BAIR commons project with Karol Hausman and Sergey Levine

2019 AWS Research Proposal accepted for $\$40,000\$$

2018 AWS Research Proposal accepted for $\$15,000\$$

Service

2025 Organizing Workshop on Generalist policies in the wild and Robo-Arena challenge at CoRL

2025 Organizing Real-to-Sim-to-Real Visuomotor Robot Learning Workshop and Challenges at IROS

2025 Organizing IROS Workshop on Multimodal Robot Learning in Physical Worlds

2025 Organizing CoRL Workshop on Generalizable Priors for Robot Manipulation and CoRL Workshop on Evaluation and Deployment Across the Robot Learning Lifecycle

2025 Participant at DARPA ISAT Study Group

2025 Organizing 2 workshops at RSS on Foundation Models in Robotics and Reasoning Models in Robot Learning

2024 NSF Panelist: TRAILBLAZER Program

2024 Organizing workshop on Adaptive Learning Systems at TTIC

2024 Special Session Chair for International Symposium on Artificial Intelligence and Mathematics (ISAIM)

2024 NSF Panelist: Foundational Research in Robotics (FRR)

2022–present Co-organizer, University of Washington Robotics Seminar

2022–present Area Chair for NeurIPS, AAAI, CVPR, IROS, ICRA, ICCV, ICML, RSS, AAAI special issue

2019 Co-organized workshop on Multi-Task and Lifelong Learning at ICML 2019

2019 Co-organized workshop on Structure and Priors in RL at ICLR 2019

2018 Outstanding Reviewer at CoRL 2018

2015–present Reviewer for NeurIPS, RSS, ICML, ICLR, ICRA, IJRR, IROS, JMLR, CoRL, RA-L, Science Robotics

Outreach and Inclusion

2023 Participant in UW Faculty Cohort Program

2023 Outreach talk for FTC Robotics Outreach

2023 Outreach talk at Seattle Robotics Society

2023– Leads UW robotics lab outreach program, partnering with AVELA and other student organizations

2019 Mentor with BAIR Mentorship Program

2016 Invited talk at Bay Area Teen Science Conference

2015-2016 Outreach co-ordinator for Robot Learning Lab

2013-2014 Mentor with Berkeley Engineers and Mentors (BEAM)