

Abhishek Gupta

EECS Department, UC Berkeley
Berkeley
CA 94720

abhigupta@eecs.berkeley.edu
+1(510)-857-4588

<https://people.eecs.berkeley.edu/~abhigupta/>

<https://scholar.google.com/citations?user=1wLVDP4AAAAJ&hl=en>

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, *GPA – 3.91*.

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*.

Research Interests

Deep RL and Robotics, *Real world robotic learning, reward specification in RL, continual learning, imitation learning, dexterous robotic manipulation, meta-reinforcement learning.*

Core Machine Learning, *Meta-learning, distribution shift, Bayesian deep learning.*

References

Pieter Abbeel, *Professor, UC Berkeley*, pabbeel@cs.berkeley.edu.

Sergey Levine, *Assistant Professor, UC Berkeley*, svlevine@eecs.berkeley.edu.

Jacob Andreas, *Assistant Professor, MIT*, jda@mit.edu.

Chelsea Finn, *Assistant Professor, Stanford University*, cbfinn@cs.stanford.edu.

Jeff Clune, *Associate Professor, University of British Columbia*, jclune@gmail.com.

Satinder Singh, *Toyota Professor of Artificial Intelligence, University of Michigan, Ann-Arbor*, baveja@umich.edu.

Teaching

- Fall 2017 **Graduate Student Instructor**, *CS285: Deep Reinforcement Learning*, Instructor: Prof. Sergey Levine.
- Spring 2016 **Graduate Student Instructor**, *CS188: Introduction to Artificial Intelligence*, [Outstanding Graduate Student Instructor Award](#), Instructors: Prof. Pieter Abbeel, Prof. Anca Dragan.
- Fall 2014 **Teaching Assistant**, *CS70: Discrete Math and Probability Theory*, Instructor: Prof. Umesh Vazirani.

Honors and Awards

- 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

- 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

Selected Publications(* denotes equal contribution)

- [1] 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](#))
- [2] **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
- [3] 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

Conference Publications(* denotes equal contribution)

- [1] **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
- [2] **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
- [3] 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](#))
- [4] Tianhe Yu, Saurabh Kumar, **Abhishek Gupta**, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. In *NeurIPS*, 2020, PDF
- [5] 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](#))
- [6] 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](#))
- [7] **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

- [8] 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
- [9] 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
- [10] 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
- [11] 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
- [12] Dibya Ghosh, **Abhishek Gupta**, and Sergey Levine. Learning actionable representations with goal-conditioned policies. In *Int. Conf. on Learning Representations (ICLR)*, 2019, PDF
- [13] 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
- [14] 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
- [15] **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
- [16] **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
- [17] 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
- [18] 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
- [19] 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

- [20] **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
- [21] 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
- [22] **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
- [23] 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
- [24] Siddharth Srivastava, Shlomo Zilberstein, **Abhishek Gupta**, Pieter Abbeel, and Stuart Russell. Tractability of planning with loops. In *AAAI conference on Artificial Intelligence*, 2015, PDF
- [25] Alex X 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
- [26] 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

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

- [1] Charles Sun, Coline Devin, Jędrzej Orbik, Glen Berseth, **Abhishek Gupta**, and Sergey Levine. Realmm:real world mobile manipulation with applications to room cleaning. *Under Review*, 2020
- [2] John D Co-Reyes, Suvansh Sanjeev, Glen Berseth, **Abhishek Gupta**, and Sergey Levine. Ecological reinforcement learning. In *ArXiv Preprint*, 2020, PDF
- [3] Marvin Zhang, Henrik Marklund, **Abhishek Gupta**, Sergey Levine, and Chelsea Finn. Adaptive risk minimization: A meta-learning approach for tackling group shift. In *Under Review*, 2020, PDF
- [4] Giulia Vezzani, **Abhishek Gupta**, Lorenzo Natale, and Pieter Abbeel. Learning latent state representation for speeding up exploration. In *ArXiv Preprint*, 2019, PDF
- [5] **Abhishek Gupta***, Benjamin Eysenbach*, Chelsea Finn, and Sergey Levine. Un-supervised 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
- [6] 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

- [7] Ashvin Nair*, **Abhishek Gupta***, Murtaza Dalal, and Sergey Levine. Accelerating online reinforcement learning with offline datasets. In *Under Review*, 2021, PDF

Research Mentorship

- Undergraduate Kevin Li, Aaron Rovinsky, Tony Zhao, 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
- Graduate Aviral Kumar, John D Co-Reyes, Michael B. Chang, Giulia Vezzani, Kelvin Xu, Qiyang Li, Olivia Watkins

Grant Writing and Funding Awards

- 2020 Co-author on NSF proposal on Bootstrapping Natural Feedback for Reinforcement Learning with Pieter Abbeel (UCB), Jacob Andreas (MIT), Noah Goodman (Stanford)
- 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

- 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, ICML, ICLR, ICRA, IROS, JMLR, CoRL, RA-L

Outreach and Inclusion

- 2019 Mentor with BAIR Mentorship Program
- 2018 Talk at Thayimane Children's home
- 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)