

Aviral Kumar

2121 Berkeley Way West, 8th Floor, Berkeley, CA

✉ aviralk@berkeley.edu • 🌐 [aviralkumar2907.github.io](https://github.com/aviralkumar2907)

Google Scholar: <https://scholar.google.com/citations?user=zBUwaGkAAAAJhl=en>

Education

Ph.D. in Computer Science (August 2018 - August 2023)
University of California, Berkeley, USA
Advisor: Prof. Sergey Levine

B.Tech. with Honors in Computer Science and Engineering (July 2014 - August 2018)
Indian Institute of Technology Bombay, Mumbai, India
Advisor: Prof. Sunita Sarawagi

Research Experience

Google Brain, Mountain View. *Student Researcher* June 2020 - March 2023
Working on large-scale applications of offline RL and data-driven design with several collaborators.

Google Brain, Toronto. *Research Intern* Summer 2017 & 2018
Worked on generative models for graphs and hierarchical reinforcement learning with Geoff Hinton.

IIT Bombay. *Undergraduate Researcher* July 2017 - May 2018
Worked on calibration of neural networks, and NLP with Prof. Sunita Sarawagi

Institute of Science and Technology (IST), Austria May - July, 2016
Worked on formal verification of hybrid systems using RRTs with Prof. Tom Henzinger

Awards and Fellowships

- MIT Technology Review 35 Innovators under 35 2023
Named a semi-finalist thus far (top 100 globally across all areas of technology). Final list to come out in August 2023.
- C.V. & Daulat Ramamoorthy Distinguished Research Award, UC Berkeley [details] 2022
Given to a Ph.D. student in EECS at UC Berkeley for outstanding contributions to a new area of research in computer science.
- Apple Scholars in AI/ML PhD Fellowship 2022
- Facebook PhD Fellowship in ML (3/2100 awarded in ML worldwide) [Declined] 2021
- Berkeley EECS Departmental Fellowship 2018
- Institute Academic Prize, IIT Bombay 2014-15
- KVPY Fellowship with All India Rank 4 2012-13

Publications And Pre-Prints

Publications.....

(* indicates Equal Contribution)

- Offline Q-Learning on Diverse, Multi-Task Data Both Scales and Generalizes** [Paper] [Blog]
Aviral Kumar, Rishabh Agarwal, Xinyang Geng, George Tucker, Sergey Levine
Best Paper Runner-Up at Deep Reinforcement Learning Workshop, NeurIPS 2022
Oral Presentation at Deep Reinforcement Learning Workshop, NeurIPS 2022 (top 6/196 submissions)
International Conference on Learning Representations (ICLR), 2023 (**Notable Top 5%**) (Acceptance: 1.6%)
- Value-Based Deep Reinforcement Learning Requires Explicit Regularization** [Paper] [Talk]
Aviral Kumar, Rishabh Agarwal, Aaron Courville, Tengyu Ma, George Tucker, Sergey Levine
Oral Presentation at the Deep Reinforcement Learning Workshop at NeurIPS 2021 (given to 6/140 papers).
Oral Presentation at the Overparameterization Workshop at ICML 2021 (given to 6/45 papers).
Oral Presentation at RL for Real Life Workshop at ICML 2021 (12/90 papers).
International Conference on Learning Representations (ICLR), 2022 (**Spotlight**) (Acceptance Rate: 5.2%)
- A Workflow for Offline Model-Free Robotic Reinforcement Learning** [Paper] [Talk]
Aviral Kumar*, Anikait Singh*, Stephen Tian, Chelsea Finn, Sergey Levine
Conference on Robot Learning (CoRL), 2021 (**Oral Presentation**). (Acceptance rate: 6.5%)

4. **Data-Driven Offline Optimization for Architecting Hardware Accelerators** [Paper] [Talk] [Blog]
Aviral Kumar*, Amir Yazdanbakhsh*, Milad Hashemi, Kevin Swersky, Sergey Levine
International Conference on Learning Representations (ICLR), 2022. (Acceptance rate: 32.2%)
5. **Confidence-Conditioned Value Functions for Offline Reinforcement Learning** [Paper]
Joey Hong, Aviral Kumar, Sergey Levine
International Conference on Learning Representations (ICLR), 2023 (Notable top 5%) (Acceptance: 1.6%)
6. **How to Leverage Unlabeled Data in Offline Reinforcement Learning** [Paper] [Talk]
Tianhe Yu*, Aviral Kumar*, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Sergey Levine
International Conference on Machine Learning (ICML), 2022. (Acceptance Rate: 21.9%)
7. **Should I Run Offline Reinforcement Learning or Behavioral Cloning?** [Paper] [Blog]
Aviral Kumar*, Joey Hong*, Anikait Singh, Sergey Levine
International Conference on Learning Representations (ICLR), 2022. (Acceptance rate: 32.2%)
8. **Design-Bench: Benchmarks for Data-Driven Offline Model-Based Optimization** [Paper] [GitHub]
Brandon Trabucco*, Xinyang Geng*, Aviral Kumar, Sergey Levine.
International Conference on Machine Learning (ICML), 2022. (Acceptance Rate: 21.9%)
9. **Conservative Data-Sharing for Multi-Task Offline Reinforcement Learning** [Paper]
Tianhe Yu*, Aviral Kumar*, Yevgen Chebotar, Karol Hausman, Sergey Levine, Chelsea Finn
Neural Information Processing Systems (NeurIPS), 2021. (Acceptance Rate: 26.0%)
10. **COMBO: Conservative Offline Model-Based Policy Optimization** [Paper]
Tianhe Yu*, Aviral Kumar*, Rafail Rafailov, Aravind Rajeswaran, Sergey Levine, Chelsea Finn
Neural Information Processing Systems (NeurIPS), 2021. (Acceptance Rate: 26.0%)
11. **Why Generalization in RL is Difficult: Epistemic POMDP and Implicit Partial Observability**
Dibya Ghosh*, Jad Rahme* Aviral Kumar, Amy Zhang, Ryan Adams, Sergey Levine [Paper]
Neural Information Processing Systems (NeurIPS), 2021. (Acceptance Rate: 26.0%)
12. **Conservative Objective Models for Effective Offline Model-Based Optimization** [Paper] [Blog]
Brandon Trabucco*, Aviral Kumar*, Xinyang Geng, Sergey Levine . (* Equal Contribution)
International Conference on Machine Learning (ICML), 2021. (Acceptance rate: 21.4%)
13. **Implicit Under-Parameterization Inhibits Data-Efficient Deep Reinforcement Learning** [Paper]
Aviral Kumar*, Rishabh Agarwal*, Dibya Ghosh, Sergey Levine
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
14. **OPAL: Offline Primitive Discovery for Accelerating Offline Reinforcement Learning** [Paper]
Anurag Ajay, Aviral Kumar, Pulkit Agarwal, Sergey Levine, Ofir Nachum
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
15. **Conservative Safety Critics for Exploration** [Paper]
Homanga Bharadhwaj, Aviral Kumar, Nick Rhinehart, Sergey Levine, F. Shkurti, Animesh Garg
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
16. **Benchmarks for Deep Off-Policy Evaluation** [Paper]
J. Fu, M. Norouzi, O. Nachum, G. Tucker, Z. Wang, A. Novikov, M. Yang, M. R. Zhang, Y. Chen, A. Kumar,
C. Paduraru, S. Levine, T. Paine
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
17. **COG: Connecting New Skills to Past Experience with Offline Reinforcement Learning**
Avi Singh, Albert Yu, Jonathan Yang, Jesse Zhang, Aviral Kumar, Sergey Levine
Conference on Robotic Learning (CoRL) 2020.
18. **DisCor: Corrective Feedback in Reinforcement Learning via Distribution Correction**
Aviral Kumar, Abhishek Gupta, Sergey Levine [Paper] [Blog] [Talk]
Neural Information Processing Systems (NeurIPS), 2020 (Spotlight). (Acceptance rate: 2.96%)
19. **Conservative Q-Learning for Offline Reinforcement Learning**
Aviral Kumar, Aurick Zhou, George Tucker, Sergey Levine [Paper] [Blog]
Neural Information Processing Systems (NeurIPS), 2020. (Acceptance rate: 20.0%)
20. **Model Inversion Networks for Model-Based Optimization**
Aviral Kumar, Sergey Levine [Paper]
Neural Information Processing Systems (NeurIPS), 2020. (Acceptance rate: 20.0%)
21. **One Solution is Not All You Need: Few-Shot Extrapolation via Structured MaxEnt RL**

- Saurabh Kumar, **Aviral Kumar**, Sergey Levine, Chelsea Finn
Neural Information Processing Systems (NeurIPS), 2020. (Acceptance rate: 20.0%)
22. **Stabilizing Off-Policy Q-Learning via Bootstrapping Error Reduction**
Aviral Kumar, Justin Fu, George Tucker, Sergey Levine [Paper] [Blog]
Neural Information Processing Systems (NeurIPS), 2019 (Acceptance rate: 21.1%)
 23. **Diagnosing Bottlenecks in Deep Q-Learning Algorithms**
Justin Fu*, **Aviral Kumar***, Matthew Soh, Sergey Levine [Paper]
International Conference on Machine Learning (ICML), 2019 (Acceptance rate: 22.5%)
 24. **Graph Normalizing Flows**
Jenny Liu*, **Aviral Kumar***, Jimmy Ba, Jamie Kiros, Kevin Swersky [Paper]
Neural Information Processing Systems (NeurIPS), 2019 (Acceptance rate: 21.1%)
 25. **Trainable Calibration Measures for Neural Networks from Kernel Mean Embeddings**
Aviral Kumar, Sunita Sarawagi, Ujjwal Jain [Paper]
International Conference on Machine Learning (ICML), 2018 (Acceptance rate: 24.9%)
 26. **Data-Driven Offline Decision-Making via Invariant Representation Learning** [Paper]
Han Qi*, Yi Su*, **Aviral Kumar***, Sergey Levine
Neural Information Processing Systems (NeurIPS), 2022. (Acceptance rate: 25.6%)
 27. **DASCO: Dual-Generator Offline Reinforcement Learning** [Paper]
Quan Vuong, **Aviral Kumar**, Sergey Levine, Yevgen Chebotar
Neural Information Processing Systems (NeurIPS), 2022. (Acceptance rate: 25.6%)
 28. **Don't Start from Scratch: Leveraging Prior Data to Automate Robotic Reinforcement Learning**
Homer Walke, Jonathan Yang, Albert Yu, **Aviral Kumar**, Jędrzej Orbik, Avi Singh, Sergey Levine
Conference on Robot Learning (CoRL), 2022.
 29. **Off-Policy Actor-Critic For Recommender Systems**
Minmin Chen, Can Xu, Vince Gatto, Devanshu Jain, **Aviral Kumar**, Ed Chi
ACM Conference on Recommender Systems, 2022.
 30. **Efficient Deep Reinforcement Learning Requires Regularizing Statistical Overfitting**
Qiyang Li, **Aviral Kumar**, Ilya Kostrikov, Sergey Levine
International Conference on Learning Representations (ICLR), 2023. (Acceptance Rate: 31.8%)

Tutorial Papers.....

31. **Offline Reinforcement Learning: Tutorial, Review and Perspectives on Open Problems**
Sergey Levine, **Aviral Kumar**, George Tucker, Justin Fu [Paper]

Pre-Prints.....

32. **Pre-Training for Robots: Offline RL Can Enable Learning New Tasks in a Handful of Trials** (2022)
Aviral Kumar*, Anikait Singh*, Frederik Ebert*, Yanlai Yang, Chelsea Finn, Sergey Levine [Paper]
Oral Presentation at the CoRL 2022 Workshop on Robot Adaptation in the Real World, December 2022.
33. **Latent Conservative Objective Models for Crystal Structure Prediction** (2023)
Han Qi, Stefano Rando, Xinyang Geng, Iku Ohama, **Aviral Kumar**, Sergey Levine.
34. **Cal-QL: Calibrated Offline RL Pre-Training for Efficient Online Fine-Tuning** (2023)
Mitsuhiko Nakamoto*, Simon Zhai*, Anikait Singh, Yi Ma, Chelsea Finn, **Aviral Kumar**, Sergey Levine.
35. **Offline RL on Realistic Datasets: Heteroskedasticity and Support Constraints** [Paper] (2022)
Anikait Singh*, **Aviral Kumar***, Quan Vuong, Yevgen Chebotar, Sergey Levine
36. **Effective Offline RL Needs Going Beyond Pessimism: Representations and Dist. Shift** (2022)
Xinyang Geng*, Kevin Li*, Abhishek Gupta, **Aviral Kumar**, Sergey Levine
ICML Workshop on Decision-Awareness in Reinforcement Learning, 2022.
37. **Data-Driven Optimization for Protein Design: Workflows, Algorithms and Metrics** (2022)
Sathvik Kolli, Amy Lu, Xinyang Geng, **Aviral Kumar**, Sergey Levine
ICLR Workshop on Machine Learning for Drug Discovery, 2022.
38. **D4RL: Datasets for Deep Data-Driven Reinforcement Learning** [Paper] [GitHub] (2020)
Justin Fu, **Aviral Kumar**, Ofir Nachum, George Tucker, Sergey Levine
39. **Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning** (2019)

Tutorials, Invited Talks and Research Blog Posts

Tutorials

- **NeurIPS 2020 Tutorial on Offline Reinforcement Learning** (w/ S. Levine) [Materials] December 2020

Invited Talks

- **Recent Advances in Offline Reinforcement Learning**, *Stanford RL Forum* April 2023
- **Reinforcement Learning from Static Datasets: Algorithms, Analysis and Applications**
 - Department of Computer Science, University of Toronto February 2023
 - Computer Science Colloquium, Cornell University, Ithaca March 2023
 - EECS Department, University of Michigan at Ann Arbor March 2023
 - Department of Computer Science, Princeton University March 2023
 - Computer Science and Machine Learning Departments, Carnegie Mellon University April 2023
 - Cornell Tech CS Colloquium April 2023
- **Offline RL for Robot Pre-Training and Fine-Tuning**, *NeurIPS 2022 Real Robot Competition* Dec. 2022
- **Implementation Talk: Offline RL and CQL**, *NeurIPS 2022 Deep RL Workshop* Dec. 2022
- **How To Make Offline RL Scale and Generalize**, *CS 285 guest lecture at UC Berkeley* Nov. 2022
- **Pre-Training and Fine-Tuning Robots using Offline RL**, *UC Berkeley* Oct. 2022
- **When and How to Use Offline Reinforcement Learning?**, *Princeton University* April 2022
- **Tackling Optimization Challenges in Offline RL**
 - Facebook AI Research, Paris* April 2022
 - Google Apprenticeship Learning Summit (invited)* February 2022
- **Optimization and Workflows for Offline RL**, *LinkedIn Distinguished Speaker Series* March 2022
- **Offline Optimization for Architecting Hardware Accelerators**, *Deepmind & Waymo* November 2021
- **Understanding Value-Based Deep Reinforcement Learning**, *Google Brain RL Seminar* August 2021
- **Making Deep RL Easier to Use: Tackling Optimization & Tuning Challenges with Deep RL Methods**
 - General Motors RL/Planning Reading Group* August 2021
 - ML Collective: Deep Learning Classics and Trends* July 2021
- **Offline Reinforcement Learning: Emerging Trends & Algorithms**
 - iRobot Corporation* January 2021
 - University of Southern California (USC), RL Reading Group* January 2021
- **Offline Reinforcement Learning**
 - UC Berkeley, Host: Prof. Jiantao Jiao* September 2020
 - NYU RL Seminar* August 2020
- **Conservative Q-Learning for Offline Reinforcement Learning**, *Google Brain* July 2020
- **An Introduction to Modern Reinforcement Learning**, *INFO 251 guest lecture, UC Berkeley* April 2020
- **Model Inversion Networks**, *Google Brain, Mountain View* December 2019
- **Stabilizing Off-Policy Q-Learning**, *Google Brain, Mountain View* July 2019
- **Reversible Graph Nets: Discriminative and Generative Models**, *Google Brain, Toronto* August 2018
- **Hierarchical Reinforcement Learning with Substructure**, *Google Brain, Toronto* July 2017

Workshop Panels

- **Directions in Deep Reinforcement Learning**, *NeurIPS 2021 Deep RL Workshop* December 2021

Blog Posts

- **Data-Driven Deep Reinforcement Learning**, *Berkeley AI Research (BAIR) Blog* December 2019
- **Does On-Policy Data Collection Fix Errors in Reinforcement Learning?**, *BAIR Blog* March 2020
- **Reinforcement Learning is Supervised Learning on Optimized Data**, *BAIR Blog* October 2020
- **Offline RL: How Conservative Algorithms can Enable New Applications**, *BAIR Blog* December 2020
- **Designs from Data: Offline Optimization via Conservative Training**, *BAIR Blog* October 2021
- **Offline Optimization for Architecting Hardware Accelerators**, *Google AI Blog* March 2022

- **Should I use RL or Imitation Learning?**, *Berkeley AI Research (BAIR) Blog* April 2022
- **Pre-Training Generalist Agents via Offline Reinforcement Learning**, *Google AI Blog* February 2023

Grant Proposals

Co-Author for the Proposal on “*Offline RL for Energy-Efficient Power Grids*” with Sergey Levine, Natasha Jaques, Justin Fu (UC Berkeley) and J. Zico Kolter, Priya Donti (CMU). **Approved for funding** of \$250,000 under the C3.ai DTI Initiative to Advance AI for Energy and Climate Security, March 2021.

Research Mentoring

I have had the fortune of working with and mentoring amazing student collaborators.

Undergraduate & Masters Students

- Brandon Trabucco (2020-21; UC Berkeley → PhD student at CMU)
- Homanga Bharadhwaj (2020; MS at University of Toronto → PhD student at CMU)
- Anikait Singh (UC Berkeley, to graduate in 2023; CRA Undergraduate Researcher Finalist)
- Han Qi (UC Berkeley, to graduate in 2023; CRA nomination from UC Berkeley)
- Jason Wang (UC Berkeley, to graduate in 2023)
- Sathvik Kolli (5th year MS at UC Berkeley)
- Kevin Li (5th year MS at UC Berkeley)

PhD Students

- Saurabh Kumar (Stanford University; Fall 2019 - Spring 2020)
- Tianhe Yu (Stanford University; Fall 2020 - 2022)
- Dibya Ghosh (UC Berkeley; Fall 2020 - Summer 2021)
- Jad Rahme (Princeton University; Fall 2020 - Summer 2021)
- Amy Lu (UC Berkeley; Fall 2021 - Fall 2022)
- Joey Hong (UC Berkeley; Fall 2021 - Fall 2022)
- Qiyang Li (UC Berkeley; since Summer 2022)
- Mitsuhiko Nakamoto (UC Berkeley; since Fall 2022)
- Jakub Grudzien (UC Berkeley; since Fall 2022)
- Simon Zhai (UC Berkeley; since Fall 2022)

Professional Service

Workshop Organization

- NeurIPS Workshop on Offline Reinforcement Learning (New Orleans) Dec. 2022
- NeurIPS Workshop on Offline Reinforcement Learning (Virtual) Dec. 2021
- NeurIPS Workshop on Offline Reinforcement Learning (Virtual) Dec. 2020

Lead co-organizer for the workshop for three consecutive years: 2020, 2021 and 2022.

Conference and Journal Reviewing

- Area Chair for NeurIPS 2023
- Reviewer for ICLR (2020-2023) [**top reviewer award at ICLR 2021**], ICML (2020-2023), NeurIPS (2019-2022) [**top 50% of all reviewers at NeurIPS 2020**], CoRL (2021-2022), AISTATS (2022), NeurIPS Workshops Selection Committee (2021), Internal Journal for Robotics Research (IJRR) (2021), Transactions on Pattern Analysis and Machine Intelligence (TPAMI) (2019).
- Area Chair for Workshop on Meta Learning at NeurIPS 2020, 2021

Teaching Experience

- **Graduate Student Instructor** for Deep Reinforcement Learning (CS 285), UC Berkeley Fall 2020

Designed and delivered 3 lectures on topics in offline RL, RL theory and recent topics in RL algorithm design in addition to guiding students on final projects, designing homeworks and answering questions. Lecture slides are available [\[here\]](#) (lectures 15-17) and lecture recordings are available [\[here\]](#).

- **Graduate Student Instructor** for Introduction to Artificial Intelligence (CS 188), UC Berkeley Fall 2022
- **Teaching Assistant** for Logic in Computer Science (CS 228), IIT Bombay Spring 2017
- **Teaching Assistant** for Discrete Structures (CS 207), IIT Bombay Fall 2016
- **Teaching Assistant** for Calculus (MA 205), IIT Bombay Fall 2015

Service, Outreach and Inclusion

- Mentor with the BAIR Undergraduate Mentorship Program 2022
- UC Berkeley EECS PhD Admissions Committee 2019, 2020
- Mentor with the Department Academic Mentorship Program, IIT Bombay 2017
- Mentor with the National Service Scheme at IIT Bombay, India 2014-2015