

Aviral Kumar

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

✉ aviralk@berkeley.edu • 🌐 aviralkumar2907.github.io

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

Education

Ph.D. in Computer Science (August 2018 - Present)

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

Other Research Experience

Google Brain, Mountain View. *Student Researcher* June 2020 - Summer 2022
Working on applications of offline RL, black-box model-based optimization in chip design and related problems.

Google Brain, Toronto. *Research Intern* May - August, 2018

Worked on generative models for graphs with Jimmy Ba, Kevin Swersky and Geoff Hinton.

Google Brain, Toronto. *Research Intern* May - July, 2017

Worked on hierarchical reinforcement learning with Kevin Swersky and Geoff Hinton.

IIT Bombay. *Undergraduate Researcher* July 2017 - May 2018

Worked on calibration of neural networks, and NMT models 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

Academic Achievements/Fellowships

- Apple Scholars in AI/ML PhD Fellowship 2022
- Facebook PhD Fellowship in ML, 2021 (3/2100 applicants awarded worldwide) [Declined]
- Berkeley EECS Departmental Fellowship 2018
- Institute Academic Prize, IIT Bombay 2014-15
- KVPY Fellowship with All India Rank 4 2012-13

Publications

Publications.....

(* indicates Equal Contribution)

- Value-Based Deep Reinforcement Learning Requires Explicit Regularization**
Aviral Kumar, Rishabh Agarwal, Aaron Courville, Tengyu Ma, George Tucker, Sergey Levine
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 Presentation)
[Paper] [Talk] (Acceptance Rate: 5.2%)
- A Workflow for Offline Model-Free Robotic Reinforcement Learning**
Aviral Kumar*, Anikait Singh*, Stephen Tian, Chelsea Finn, Sergey Levine [Paper] [Talk]
Conference on Robot Learning (CoRL), 2021 (*Oral* Presentation). (Acceptance rate: 6.5%)
- Data-Driven Optimization for Architecting Hardware Accelerators**
Aviral Kumar, Amir Yazdanbaksh, Milad Hashemi, Kevin Swersky, Sergey Levine [Paper] [Talk] [Blog]
International Conference on Learning Representations (ICLR), 2022. (Acceptance rate: 32.2%)

4. **Should I Run Offline Reinforcement Learning or Behavioral Cloning?** [Paper]
Aviral Kumar, Joey Hong, Anikait Singh, Sergey Levine
International Conference on Learning Representations (ICLR), 2022. (Acceptance rate: 32.2%)
5. **Conservative Data-Sharing for Multi-Task Offline Reinforcement Learning**
Tianhe Yu*, Aviral Kumar*, Yevgen Chebotar, Karol Hausman, Sergey Levine, Chelsea Finn [Paper]
Neural Information Processing Systems (NeurIPS), 2021. (Acceptance Rate: 26.0%)
6. **COMBO: Conservative Offline Model-Based Policy Optimization**
Tianhe Yu*, Aviral Kumar*, Rafail Rafailov, Aravind Rajeswaran, Sergey Levine, Chelsea Finn [Paper]
Neural Information Processing Systems (NeurIPS), 2021. (Acceptance Rate: 26.0%)
7. **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%)
8. **Conservative Objective Models for Effective Offline Model-Based Optimization**
Brandon Trabucco*, Aviral Kumar*, Xinyang Geng, Sergey Levine . (* Equal Contribution) [Paper]
International Conference on Machine Learning (ICML), 2021. (Acceptance rate: 21.4%)
9. **Implicit Under-Parameterization Inhibits Data-Efficient Deep Reinforcement Learning**
Aviral Kumar, Rishabh Agarwal, Dibya Ghosh, Sergey Levine [Paper]
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
10. **OPAL: Offline Primitive Discovery for Accelerating Offline Reinforcement Learning**
Anurag Ajay, Aviral Kumar, Pulkit Agarwal, Sergey Levine, Ofir Nachum [Paper]
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
11. **Conservative Safety Critics for Exploration**
Homanga Bharadhwaj, Aviral Kumar, Nick Rhinehart, Sergey Levine, F. Shkurti, Animesh Garg [Paper]
International Conference on Learning Representations (ICLR), 2021. (Acceptance rate: 28.7%)
12. **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.
13. **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** Presentation). (Acceptance rate: 2.96%)
14. **Conservative Q-Learning for Offline Reinforcement Learning**
Aviral Kumar, Aurick Zhou, George Tucker, Sergey Levine [Paper]
Neural Information Processing Systems (NeurIPS), 2020. (Acceptance rate: 20.0%)
15. **Model Inversion Networks for Model-Based Optimization**
Aviral Kumar, Sergey Levine [Paper]
Neural Information Processing Systems (NeurIPS), 2020. (Acceptance rate: 20.0%)
16. **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%)
17. **Stabilizing Off-Policy Q-Learning via Bootstrapping Error Reduction**
Aviral Kumar, Justin Fu, George Tucker, Sergey Levine [Paper]
Neural Information Processing Systems (NeurIPS), 2019 (Acceptance rate: 21.1%)
18. **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%)
19. **Graph Normalizing Flows**
Jenny Liu*, Aviral Kumar*, Jimmy Ba, Jamie Kiros, Kevin Swersky [Paper]
Neural Information Processing Systems (NeurIPS), 2019 (Acceptance rate: 21.1%)
20. **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%*)

Tutorial Papers.....

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

Pre-Prints.....

22. **How to Leverage Unlabeled Data in Offline Reinforcement Learning** (2022)
Tianhe Yu*, **Aviral Kumar***, Yevgen Chebotar, Chelsea Finn, Karol Hausman, Sergey Levine [Paper]
23. **Design-Bench: Benchmarks for Data-Driven Offline Model-Based Optimization** (2022)
Brandon Trabucco*, Xinyang Geng*, **Aviral Kumar**, Sergey Levine. [Paper] [GitHub]
24. **D4RL: Datasets for Deep Data-Driven Reinforcement Learning** (2020)
Justin Fu, **Aviral Kumar**, Ofir Nachum, George Tucker, Sergey Levine [Paper] [GitHub]
25. **Advantage-Weighted Regression: Simple and Scalable Off-Policy Reinforcement Learning** (2019)
Xue Bin Peng, **Aviral Kumar**, Grace Zhang, Sergey Levine [Paper]

Tutorials, Invited Talks and Research Blog Posts

Tutorials.....

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

Invited Talks.....

- **How to Tune Value-Based Offline RL In Practice**, *LinkedIn Distinguished Speaker Series* March 2022
- **Tackling Optimization Challenges in Offline RL**, *Google Apprenticeship Learning Summit* Feb 2022
- **Offline Optimization for Architecting Hardware Accelerators**, *Deepmind & Waymo* November 2021
- **Understanding Mysteries in Value-Based Deep Reinforcement Learning**, *Google Brain* August 2021
- **Making Deep RL Easier to Use: Tackling Optimization and Tuning Challenges with Deep RL Methods**, *General Motors RL/Planning Reading Group* August 2021
- **Making Deep RL Easier to Use: Tackling Optimization and Tuning Challenges with Deep RL Methods**, *ML Collective: Deep Learning Classics and Trends* July 2021
- **Offline Reinforcement Learning: Emerging Trends & Algorithms**, *iRobot Corporation* January 2021
- **Offline Reinforcement Learning: Emerging Trends & Algorithms**, *USC* January 2021
- **Offline Reinforcement Learning**, *UC Berkeley, Host: Prof. Jiantao Jiao* September 2020
- **Offline Reinforcement Learning: Challenges and Algorithms**, *NYU RL Seminar* August 2020
- **Conservative Q-Learning for Offline Reinforcement Learning**, *Google Brain* July 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

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

Grant Proposals

- Co-Author for the Proposal on "Offline Reinforcement Learning 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.

Relevant Coursework

Machine Learning – Theory of Reinforcement Learning and Bandits (UCB), Reinforcement Learning and Empirical Processes (UCB), Optimization (UCB), Deep Reinforcement Learning (UCB), Information Theory (UCB), Computer Vision (UCB), Machine Learning for Systems (UCB), Probabilistic Graphical Models and Deep Learning, Foundations of Learning Agents (Multi-armed bandits), Information Retrieval.

Statistics – Probability Theory, Statistical Inference, Theoretical Statistics (UCB)

Professional Service

- **Workshop:** Organizer & co-founder for Workshop on Offline Reinforcement Learning at NeurIPS 2020, 2021.
- **Reviewer** for ICLR (2020-2022) [**top reviewer award at ICLR 2021**], ICML (2020, 2021, 2022), NeurIPS (2019-2021) [**top 50% of all reviewers at NeurIPS 2020**], CoRL (2021), 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
- **Guest Lecturer** on Reinforcement Learning at Applied Machine Learning class (INFO 251), UC Berkeley.

Teaching Experience

- **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 students' questions. Lecture slides are available [\[here\]](#) (lectures 15-17) and lecture recordings are available [\[here\]](#).
- Logic in Computer Science (CS228), IIT Bombay Spring 2017
- Discrete Structures (CS 207), IIT Bombay Fall 2016
- Calculus (MA 105), IIT Bombay Fall 2015