Sarah Dean

https://sdean.website sdean@cornell.edu

Research Interests

I study the interplay between optimization, machine learning, and dynamics in real-world systems with the goal of understanding the fundamentals of data-driven methods for control and decision-making. My research is grounded in collaborative projects ranging from robotics to recommendation systems.

Academic Positions

Assistant Professor, Department of Computer Science Jan 2022 – present
Cornell University, Ithaca, NY.

Postdoctoral Scholar, Paul G. Allen School of Computer Science & Engineering Aug 2021 – Dec 2021
University of Washington, Seattle, WA.
Advised by Prof. Jamie Morgenstern.

Education

University of California, Berkeley
Ph.D., Electrical Engineering and Computer Science, August 2021.
Thesis: Reliable Machine Learning in Feedback Systems, advised by Prof. Benjamin Recht.
M.S., Electrical Engineering and Computer Science, May 2019.

University of Pennsylvania
B.S.E., Electrical Engineering and Mathematics, May 2016.

Honors and Awards

Best Paper Finalist, *Conference on Robot Learning*	2020
Best Paper Award, *NeurIPS Joint Workshop on AI for Social Good*	2019
Best Paper Award, *International Conference of Machine Learning*	2018
Best Student Paper in Imaging Systems, *OSA Imaging Applied Optics Congress*	2018
Tong Leong Lim Pre-Doctoral Prize, *UC Berkeley EECS Department*	2018
Atwater Kent Prize in Electrical Engineering, *University of Pennsylvania*	2016
Albert P. Godsho Engineering Prize, *University of Pennsylvania*	2016
Hugo Otto Wolf Memorial Prize, *University of Pennsylvania*	2016
E. Stuart Eichert, Jr. Memorial Prize for Electrical Engineering, *University of Pennsylvania*	2015
Good Teaching Award, *UPenn Math Department*	2015

Grants and Fellowships

Cornell Atkinson 2030 Project Fast Grant, *Expanding Weather Data Coverage*	2024
NSF OAC Grant (co-PI), *Frameworks: arXiv as an accessible large-scale open research platform*	2024
NSF CCF Medium Grant (lead PI), *Machine Learning Markets: Dynamics, Competition, and Interventions*	2023
Research Award, *LinkedIn*	2023
Bias and Transparency in AI Award, *Mozilla Technology Fund*	2023
Gift for Recommendations with Long-Term Strategic Objectives, *Wayfair*	2022
Research Gift, *Meta*	2022
Center for Longterm Cybersecurity Project Grant, *UC Berkeley*	2020
Social Science Matrix Research Grant, *UC Berkeley*	2019
Center for Longterm Cybersecurity Seed Grant, *UC Berkeley*	2019
NSF Graduate Research Fellowship	2016
Berkeley Fellowship, *UC Berkeley*	2016
Tau Beta Pi Fellowship	2016

Teaching

Instructor, *Cornell University CS Department.*
- Machine Learning in Feedback Systems, Fall 2022 and 2023.
- Introduction to Reinforcement Learning, Spring 2022-2024.

Graduate Student Instructor, *University of California, Berkeley EECS Department.*
- EECS Anti-Racism and Social Justice Course Development, Fall 2020.
- Statistical Learning Theory, Fall 2019.
- Introduction to Machine Learning, Fall 2018.

Teaching Assistant, *John's Hopkins Center for Talented Youth at Skidmore College.*
- Electrical Engineering, Summer 2016.

Teaching Assistant, *University of Pennsylvania ESE Department.*
- Digital Audio Basics, Spring 2014, 2016.
- Introduction to Electrical and Systems Engineering. Fall 2013, 2014, 2015.

Teaching Assistant, *University of Pennsylvania Math Department.*
- Integral Calculus, Spring 2016.
- Multivariate Calculus, Fall 2014, Spring 2015.

Tutor, *University of Pennsylvania.*
- Multivariate Calculus, Spring 2013, Fall 2013, Spring 2014.
- Linear Algebra and Differential Equations, Fall 2013, Spring 2014.

Internships

Research Intern at Canopy — Summer 2019
Explored concepts relating to user agency and developed a computationally efficient audit of model "reachability."

Infrastructure Quality Engineer Intern at Palantir — Summer 2015
Created an automated test suite for a data sharing product; wrote regression tests for a front end web form product.

Service and Leadership

Publications Chair for L4DC (2023). **Organizing Committee** for Workshop on Decision Making for Information Retrieval and Recommender Systems at WWW (2023) and Recommendation Ecosystems Workshop: Modeling, Optimization, and Incentive Design at AAAI (2024).

Area Chair for NeuRIPS, ICML, and L4DC. **Conference reviewer** for ALT, ACC, CDC, ICML, ITCS, L4DC, and NeurIPS. **Journal reviewer** for IEEE-TAC, JMLR, SIMODS, and Springer Machine Learning.

Co-founder of Graduates for Engaged and Extended Scholarship in Computing and Engineering (geesegraduates.org), a cross-disciplinary group that aims to give graduate students a constructive place to reflect on issues of society and technology and **organizer** of affiliated panel and speaker events.

Women in Computer Science and Engineering lunch coordinator, 2018. **WITI@UC Women in Tech Symposium** planning committee, 2019.

Volunteer mentor for students in elementary school (Bay Area Scientists in Schools, 2017), middle school (Be A Scientist, 2016), high school (CalMentors, 2020), and college (BAIR Undergraduate Mentoring Program, 2017).

Publications

Preprints

1. *Initializing Services in Interactive ML Systems for Diverse Users.* arXiv:2312.11846
   Avinandan Bose, Mihaela Curmei, Daniel L Jiang, Jamie Morgenstern, Sarah Dean, Lillian J Ratliff, Maryam Fazel.
2. *Do Offline Metrics Predict Online Performance in Recommender Systems?* arXiv:2011.07931.
   Karl Krauth, Sarah Dean, Alex Zhao, Wenshuo Guo, Mihaela Curmei, Benjamin Recht, and Michael I. Jordan.

Conference Papers

1. *Strategic Usage in a Multi-Learner Setting.*
   International Conference on Artificial Intelligence and Statistics (AISTATS), 2024.
   Eliot Seo Shekhtman, Sarah Dean.
2. *Emergent segmentation from participation dynamics and multi-learner retraining.*
   International Conference on Artificial Intelligence and Statistics (AISTATS), 2024.
   Sarah Dean, Mihaela Curmei, Lillian J. Ratliff, Jamie Morgenstern, Maryam Fazel.
3. *Ranking with Long-Term Constraints*
   ACM International Conference on Web Search and Data Mining (WSDM), 2024.
   Kianté Brantley, Zhichong Fang, Sarah Dean, Thorsten Joachims.
4. *Online Convex Optimization with Unbounded Memory.*
   Advances in Neural Information Processing Systems (NeurIPS), 2023.
   Raunak Kumar, Sarah Dean, Robert D. Kleinberg.

5. *Reward Reports for Reinforcement Learning.*
   AAAI/ACM Conference on AI, Ethics, and Society (AIES), 2023.
   Thomas Krendl Gilbert, Sarah Dean, Tom Zick, Nathan Lambert, Aaron Snoswell, Soham Mehta.
6. *Perception-Based Sampled-Data Optimization of Dynamical Systems .*
   IFAC World Congress, 2023.
   Liliaokeawawa Cothren, Gianluca Bianchin, Sarah Dean, Emiliano Dall'Anese.
7. *Modeling Content Creator Incentives on Algorithm-Curated Platforms.*
   International Conference on Learning Representations, 2023.
   Jiri Hron, Karl Krauth, Michael I. Jordan, Niki Kilbertus, Sarah Dean.
8. *Preference Dynamics Under Personalized Recommendations.*
   ACM Conference on Economics and Computation, 2022.
   Sarah Dean and Jamie Morgenstern.
9. *Towards Robust Data-Driven Control Synthesis for Nonlinear Systems with Actuation Uncertainty.*
   IEEE Conference on Decision and Control (CDC), 2021.
   Andrew J. Taylor, Victor D. Dorobantu, Sarah Dean, Benjamin Recht, Yisong Yue, and Aaron D. Ames.
10. *Quantifying Availability and Discovery in Recommender Systems via Stochastic Reachability.*
    International Conference on Machine Learning (ICML), 2021.
    Mihaela Curmei, Sarah Dean, and Benjamin Recht.
11. *Certainty-Equivalent Perception-Based Control.*
    Learning for Dynamics and Control (L4DC), 2021.
    Sarah Dean and Benjamin Recht.
12. *AI Development for the Public Interest: From Abstraction Traps to Sociotechnical Risks.*
    IEEE International Symposium on Technology and Society (ISTAS), 2020.
    McKane Andrus, Sarah Dean, Thomas Krendl Gilbert, Nathan Lambert, and Tom Zick.
13. *Guaranteeing Safety of Learned Perception Modules via Measurement-Robust Control Barrier Functions.*
    Conference on Robot Learning (CoRL), 2020.
    Sarah Dean, Andrew Taylor, Ryan Cosner, Benjamin Recht, and Aaron Ames.
14. *Balancing Competing Objectives with Noisy Data: Score-Based Classifiers for Welfare-Aware Machine Learning.*
    International Conference on Machine Learning (ICML), 2020.
    Esther Rolf, Max Simchowitz, Sarah Dean, Lydia T. Liu, Daniel Bjorkegren, Moritz Hardt, and Joshua Blumenstock.
15. *Robust Guarantees for Perception-Based Control.*
    Learning for Dynamics and Control (L4DC), 2020.
    Sarah Dean, Nikolai Matni, Benjamin Recht, and Vickie Ye.
16. *Recommendations and User Agency: The Reachability of Collaboratively-Filtered Information.*
    Conference on Fairness, Accountability, and Transparency (FAccT), 2020.
    Sarah Dean, Sarah Rich, and Benjamin Recht.
17. *Safely Learning to Control the Constrained Linear Quadratic Regulator.*
    American Controls Conference (ACC), 2019.
    Sarah Dean, Stephen Tu, Nikolai Matni, and Benjamin Recht.
18. *Regret Bounds for Robust Adaptive Control of the Linear Quadratic Regulator.*
    Advances in Neural Information Processing Systems (NeurIPS), 2018.
    Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu.
19. *Delayed Impact of Fair Machine Learning.*
    International Conference on Machine Learning (ICML), 2018.
    Lydia T. Liu, Sarah Dean, Esther Rolf, Max Simchowitz, and Moritz Hardt.

Journal Articles

1. *Axes for Sociotechnical Inquiry in AI Research.*
   IEEE Transactions on Technology and Society, 2021.
   Sarah Dean, Thomas Krendl Gilbert, Nathan Lambert, and Tom Zick.
2. *High-throughput fluorescence microscopy using multi-frame motion deblurring.*
   Biomedical Optics Express, 2020.
   Zachary Phillips, Sarah Dean, Laura Waller, and Benjamin Recht.

3. *On the Sample Complexity of the Linear Quadratic Regulator.*
   Foundations of Computational Mathematics, 2019.
   Sarah Dean, Horia Mania, Nikolai Matni, Benjamin Recht, and Stephen Tu.

Whitepapers

1. *Choices, Risks, and Reward Reports: Charting Public Policy for Reinforcement Learning Systems.*
   Center for Long-Term Cybersecurity Whitepaper Series, 2022.
   Thomas Krendl Gilbert, Sarah Dean, Tom Zick, Nathan Lambert.

Workshop Papers

1. *Recommender Systems as Dynamical Systems: Interactions with Viewers and Creators.*
   Workshop on Recommendation Ecosystems: Modeling, Optimization and Incentive Design at AAAI, 2024.
   Sarah Dean, Evan Dong, Meena Jagadeesan, Liu Leqi.
2. *To ask or not to ask: Robot-assisted bite acquisition with human-in-the-loop contextual bandits.*
   Out-of-Distribution Generalization in Robotics at Conference on Robot Learning (CoRL), 2023.
   Rohan Banerjee, Sarah Dean, Tapomayukh Bhattacharjee.
3. *Random Features Approximation for Fast Data-Driven Control.*
   Gaussian Processes, Spatiotemporal Modeling, and Decision-making Systems NeuRIPS 2022.
   Kimia Kazemian and Sarah Dean.
4. *Cross-Dataset Propensity Estimation for Debiasing Recommender Systems.*
   Workshop on Distribution Shifts: Connecting Methods and Applications at NeurIPS 2022.
   Fengyu Li and Sarah Dean.
5. *Engineering a Safer Recommender System.*
   Responsible Decision Making in Dynamic Environments Workshop at ICML 2022.
   Liu Leqi and Sarah Dean.
6. *Reward Reports for Reinforcement Learning.*
   Responsible Decision Making in Dynamic Environments Workshop at ICML 2022.
   Thomas Krendl Gilbert, Sarah Dean, Tom Zick, Nathan Lambert, Aaron Snoswell.
7. *Designing Recommender Systems with Reachability in Mind.*
   Participatory Approaches to Machine Learning Workshop at ICML 2020.
   Sarah Dean, Mihaela Curmei, and Benjamin Recht.
8. *Balancing Competing Objectives for Welfare-Aware Machine Learning with Imperfect Data.*
   AI for Social Good Workshop at NeurIPS 2019.
   Esther Rolf, Max Simchowitz, Sarah Dean, Lydia T. Liu, Daniel Bjorkegren, Moritz Hardt, and Joshua Blumenstock.
9. *Optimal Path and Illumination Design for Multiframe Motion Deblurring.*
   OSA Imaging and Applied Optics Congress 2018.
   Sarah Dean, Zachary Phillips, Laura Waller, and Benjamin Recht.
10. *A Broader View on Bias in Automated Decision-Making: Reflecting on Epistemology and Dynamics.* Workshop on fairness, accountability, and transparency in machine learning. (FAT/ML), 2018.
    Roel Dobbe, Sarah Dean, Thomas Gilbert, and Nitin Kohli.

Invited Talks

- *Feedback and Dynamics in Machine Learning Systems*, Brookings Virtual Seminar on AI, Economics, and Public Policy, July 2023.
- *User Dynamics in Machine Learning Systems*, Princeton Networks & Cognition Workshop, June 2023.
- *On Uniform Error Bounds and Guarantees for Perception-Based Control*, ACC Workshop on Safe & Robust Learning for Perception-based Planning and Control, May 2023.
- *Online Convex Optimization with Unbounded Memory*, ACC Workshop on Online Optimization Methods for Data-driven Feedback Control, May 2023.
- *Preference and Participation Dynamics in Learning Systems*
  - L4DC Keynote, June 2022.
  - Cornell AI Seminar, September 2022.
  - NYU Math and Data Seminar, February 2023.
  - Cornell Econometrics Reading Group, May 2023.
  - Princeton Networks & Cognition Workshop, June 2023.

  - Brookings Virtual Seminar on AI, Economics, and Public Policy, July 2023.
- *Feedback, Dynamics, and Safety in Machine Learning Systems*, NCCR Symposium on Socially responsible Automation, October 2022.
- *Data-driven Control and Decision-making in Feedback Systems*, Cornell CAM Colloquium, January 2022.
- *Towards Certifiably Safe Nonlinear Control with Sensor and Dynamics Uncertainties*
  - UCSD Dynamic Systems & Controls Seminar, December 2021.
  - CISS Invited Session on Safe Reinforcement Learning, March 2022.
  - Minisymposium on Learning from scarce data at SIAM Conference on Mathematics of Data Science, September 2022.
- *Quantifying Availability and Discovery in Recommender Systems via Reachability*, Cornell AI Seminar, September 2021.
- *Reliable Machine Learning in Feedback Systems*
  - Robotics Institute Seminar at Carnegie Mellon University, April 2021.
  - CS Department Colloquium at Princeton University, March 2021.
  - CS Seminar at Brown University, March 2021.
  - Allen School Colloquium at University of Washington, March 2021.
  - ECE Seminar at University of Michigan, March 2021.
  - CS Colloquium at NYU, March 2021.
  - ESE Spring Colloquium at University of Pennsylvania, March 2021.
  - ECE Seminar at University of Wisconsin at Madison, March 2021.
  - CS Seminar at Northeastern University, February 2021.
  - ECE Seminar at Cornell Tech, February 2021.
  - EECS Seminar at Massachusetts Institute of Technology, February 2021.
  - CSE Colloquium at University of Minnesota, February 2021.
  - MINDS Symposium on the Foundations of Data Science at Johns Hopkins University, February 2021.
  - CS Seminar at University of Chicago, February 2021.
  - CS Lecture at University of Texas at Austin, February 2021.
  - MS&E Seminar at Stanford University, January 2021.
  - CS Colloquium at Cornell University, January 2021.
  - Frontiers in Computing and Mathematical Sciences at California Institute of Technology, January 2021.
- *On the Sample Complexity of the Linear Quadratic Regulator*, RL Theory Virtual Seminar, May 2020.
- *Safe and Robust Perception-Based Control*
  - Stanford Robotics and Autonomous Systems Seminar, February 2020.
  - CDS Seminar at California Institute of Technology, February 2020.
- *Delayed Impact of Fair Machine Learning*, Sister Conferences Track at the International Joint Conferences on Artificial Intelligence, August 2019.
- *Guarantees for Learning-Enabled Control*, Interplay between Control, Optimization, and Machine Learning Workshop at the American Controls Conference, July 2019.
- *Safely Learning to Control the Linear Quadratic Regulator*, CITRIS/CPAR Control Theory and Automation Symposium, April 2019.