Berk Ustun

Curriculum Vitae — April 2025

berk@ucsd.edu www.berkustun.com

| Academic Positions | University of California, San Diego Assistant Professor, School of Computing Information & Data Sciences | 2021 – Present |
|--------------------------------|---|--|
| | Google Deepmind Visiting Faculty, Responsible AI & Human-Centered Technology | 2020 – 202 I |
| | Harvard University Postdoctoral Fellow, Center for Research on Computation & Society | 2017 - 2020 |
| Research Interests | Areas: Machine Learning, Optimization, Human-Centered Design Topics: Safety, Interpretability, Fairness, Personalization, Governance Domains: Healthcare, Consumer Finance, Criminal Justice, Physical Sciences | |
| Education | Massachusetts Institute of Technology Ph.D. in Electrical Engineering & Computer Science Thesis: <u>Simple Linear Classifiers via Discrete Optimization</u> Committee: Cynthia Rudin, Leslie Kaebling, Stefanie Jegelka | 2012 – 2017 |
| | S.M. in Computation for Design & Optimization Thesis: <u>MCMC for Importance Sampling in Stochastic Programming</u> | 2010-2012 |
| | University of California, Berkeley B.S. in Industrial Engineering and Operations Research B.A. in Economics | 2005 – 2009 |
| Selected Awards & Honors | National Academy of Sciences – Kavli Fellow Amazon Research Award INFORMS Innovative Applications in Analytics Award Invenia Labs Social Environmental & Economic Challenges Award – Finalist INFORMS Computing Society Best Student Paper Prize INFORMS Wagner Award for Excellence in Operations Research Practice – Finalist INFORMS Innovative Applications in Analytics Award MIT Presidential Fellowship | 2022 2023 2019 2018 2017 2017 2016 2012 |
| Publications | See 🕿 <u>Google Scholar</u> or 🎓 <u>Semantic Scholar</u> for a complete list. | |
| * Equal 36. Contribution | Feature Responsiveness Scores: Model-Agnostic Explanations for Recourse Harry Cheon, Anneke Wernerfelt, Sorelle Friedler, Berk Ustun ICLR – International Conference on Learning Representations, 2025 | |
| 35. | <u>Regretful Decisions under Label Noise</u> Sujay Nagaraj, Yang Liu, Flavio Calmon, Berk Ustun ICLR – International Conference on Learning Representations, 2025 | |
| 34. | Learning with Temporal Label Noise Sujay Nagaraj, Walter Gerych, Sana Tonekaboni, Anna Goldenberg, Berk Ustun, Thomas Hartvigsen ICLR – International Conference on Learning Representations, 2025 | |
| 33. | Concept Bottleneck Large Language Models Chung-En Sun, Tuomas Oikarinen, Berk Ustun, Lily Weng ICLR – International Conference on Learning Representations, 2025 | |

- 32. Discrimination Exposed? On the Reliability of Explanations for Discrimination Detection Julian Skirzynski, David Danks, Berk Ustun FAccT – ACM Conference on Fairness, Accountability, and Transparency, 2025
- 31. <u>Classification with Conceptual Safeguards</u> Hailey Joren, Charles T. Marx, Berk Ustun ICLR – International Conference on Learning Representations, 2024
- 30. Prediction without Preclusion: Recourse Verification with Reachable Sets Avni Kothari*, Bogdan Kulynych*, Lily Weng, Berk Ustun ICLR – International Conference on Learning Representations, 2024
 Spotlight Paper at ICLR24
- 29. <u>Operationalizing the Search for Less Discriminatory Alternatives in Fair Lending</u> Talia Gillis, Vitaly Mersault, **Berk Ustun** FAccT – ACM Conference on Fairness, Accountability, and Transparency, 2024
- 28. <u>Guidance for Unbiased Predictive Information for Healthcare Decision-Making and Equity</u> David Kent, Keren Ladin, John Cuddeback, Obidiugwu Duru, Sharad Goel, William Harvey, Jinny Park, Jessica Paulus, Joyce Sackey, Richard Sharp, Ewout Steyerberg, **Berk Ustun**, David van Klaveren, Saul Weingart NPJ – Nature Digital Medicine, 2024
- 27. <u>Stabilizing Recommendations by Rank-Preserving Fine-Tuning</u> Sejoon Oh, **Berk Ustun**, Julian McAuley, Srijan Kumar **TKDD** – *Transactions on Knowledge Discovery from Data*, 2024
- 26. <u>Providing Fair Recourse over Plausible Groups</u> Jayanth Yetukuri, Ian Hardy, Yevgeniy Vorobeychik, Berk Ustun, Yang Liu AAAI – Association for the Advancement of Artificial Intelligence, 2024
 Coral Presentation at AAAI24
- 25. <u>Participatory Personalization in Classification</u> Hailey Joren, Chirag Nagpal, Katherine Heller, Berk Ustun NeurIPS – Neural Information Processing Systems, 2023 ■ Spotlight Paper at NeurIPS23
- 24. When Personalization Harms Performance: Reconsidering the Use of Group Attributes in Prediction Vinith Suriyakumar, Marzyeh Ghassemi*, Berk Ustun* ICML – International Conference on Machine Learning, 2023 ■ Oral Presentation at ICML23
- 23. Predictive Multiplicity in Probabilistic Classification Jamelle Watson-Daniels, David Parkes, Berk Ustun AAAI – Association for the Advancement of Artificial Intelligence, 2023 ■ Oral Presentation at AAAI23
- 22. <u>Optimized Risk Score to Predict Mortality in Patients with Cardiogenic Shock in the Cardiac ICU</u> Eric Yamga, Sreekar Mantena, Darin Rosen, Emily Bucholz, Robert Yeh, Leo Celi, **Berk Ustun**, Neel Butala *Journal of the American Heart Association*, 2023
- 21. <u>Algorithmic Censoring in Dynamic Learning Systems</u> Jennifer Chien, Margaret Roberts, **Berk Ustun** EAAMO – ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization, 2023
- 20. Optimized Short-Forms of the Cognitive Distortions Questionnaire

Amanda S. Morrison, **Berk Ustun**, Arielle Horenstein, Simona Kaplan, Irismar Reis de Oliveira, Sedat Batmaz, James J. Gross, Ekaterina Sadikova, Curt Hemanny, Pedro Pires, Philippe Goldin, Ronald C. Kessler, Richard G. Heimberg *Journal of Anxiety Disorders*, 2022

- 19. <u>On the Epistemic Limits of Personalized Prediction</u> Lucas Monteiro^{*}, Carol Long^{*}, **Berk Ustun**, Flavio du Pin Calmon **NeurIPS** – *Neural Information Processing Systems*, 2022
- Rank List Sensitivity of Recommender Systems to Interaction Perturbations Sejoon Oh, Berk Ustun, Julian McAuley, Srijan Kumar CIKM – Conference on Information and Knowledge Management, 2022
- 17. <u>Development & Validation of a Model to Predict PTSD and Major Depression 3 Months after a Collision</u> Hannah Ziobrowski, Chris Kennedy, **Berk Ustun**, & 55 Other Authors *JAMA Psychiatry*, 2021
- 16. <u>Learning Optimal Predictive Checklists</u> Haoran Zhang, Quaid Morris, Berk Ustun*, Marzyeh Ghassemi* NeurIPS – Neural Information Processing Systems, 2021
- 15. <u>Predictive Multiplicity in Classification</u> Charles T. Marx, Flavio du Pin Calmon, **Berk Ustun** ICML – International Conference on Machine Learning, 2020
- 14. Learning Optimized Risk Scores Berk Ustun, Cynthia Rudin JMLR – Journal of Machine Learning Research, 2019
 Winner of the 2017 INFORMS Computing Society Best Student Paper Prize
 Part of Winning Entry for the 2019 INFORMS Innovative Applications in Analytics Award
- 13. Actionable Recourse in Linear Classification

Berk Ustun, Alexander Spangher, Yang Liu FAccT – ACM Conference on Fairness, Accountability, and Transparency, 2019 Featured in Wired

- <u>Developing an Optimal Short-Form of the PTSD Checklist for DSM-5</u> Kelly Zuromski, Berk Ustun, Irving Hwang, Terry Keane, Brian Marx, Murray Stein, Robert Ursano, Ronald C. Kessler Depression and Anxiety, 2019
- II. Fairness without Harm: Decoupled Classifiers with Preference Guarantees Berk Ustun, Yang Liu, David C. Parkes ICML – International Conference on Machine Learning, 2019
- <u>Repairing without Retraining: Avoiding Disparate Impact with Counterfactual Distributions</u> Hao Wang, Berk Ustun, Flavio du Pin Calmon ICML – International Conference on Machine Learning, 2019
- 9. On the Direction of Discrimination: an Information-Theoretic Analysis of Disparate Impact Hao Wang, Berk Ustun, Flavio du Pin Calmon ISIT – IEEE International Symposium on Information Theory, 2018
- 8. Optimized Risk Scores
 Berk Ustun, Cynthia Rudin
 KDD ACM SIGKDD Conference on Knowledge Discovery and Data Mining, 2017
 Winner of the 2017 INFORMS Computing Society Best Student Paper Prize
- <u>Optimized Scoring Systems: Towards Trust in ML for Healthcare and Criminal Justice</u> Cynthia Rudin, Berk Ustun *Interfaces*, 2018
 Finalist for 2018 INFORMS Wagner Award for Excellence in Operations Research Practice

- Association of an EEG-Based Risk Score With Seizure Probability in Hospitalized Patients
 Aaron Struck, Berk Ustun, Andres Rodriguez Ruiz, Jong Woo Lee, Suzette LaRoche, Lawrence Hirsch, Emily Gilmore,
 Jan Vlachy, Hiba Arif Haider, Cynthia Rudin, Brandon Westover
 JAMA Neurology, 2017
 Part of Winning Entry for the 2019 INFORMS Innovative Applications in Analytics Award
- 5. The World Health Organization Adult ADHD Self-Report Screening Scale for DSM-5

Berk Ustun, Lenard Adler, Cynthia Rudin, Stephen Faraone, Thomas Spencer, Patricia Berglund, Michael Gruber, Ronald C. Kessler JAMA Psychiatry, 2017 ■ Featured in NPR

- 4. Supersparse Linear Integer Models for Optimized Medical Scoring Systems
 Berk Ustun, Cynthia Rudin
 MLJ Machine Learning, 2015
 Part of Winning Entry for the 2016 INFORMS Innovative Applications in Analytics Award
- Interpretable Classification Models for Recidivism Prediction
 Jiaming Zeng*, Berk Ustun*, Cynthia Rudin
 Journal of the Royal Statistical Society: Series A, 2016
 Winner of the ASA Undergraduate Statistics Research Competition
- <u>Clinical Prediction Models for Sleep Apnea: Superiority of Medical History over Symptoms</u> Berk Ustun, Brandon Westover, Cynthia Rudin, Matt Bianchi Journal of Clinical Sleep Medicine, 2016
- 1. <u>Importance Sampling in Stochastic Programming: A Markov Chain Monte Carlo Approach</u> Panos Parpas, **Berk Ustun**, Mort Webster, Quan Kha Tran *INFORMS Journal of Computing*, 2015

PREPRINTS 5. <u>Selective Preference Aggregation</u> Shreyas Kadekodi^{*}, Hayden McTavish^{*}, **Berk Ustun** *In Submission*, 2025

- 4. <u>Understanding Fixed Predictions via Confined Regions</u> Connor Lawless, Tsui-Wei Wang, **Berk Ustun**, Madeleine Udell *In Submission*, 2025
- 3. <u>Mapping Enhancer-Gene Regulatory Interactions from Single-Cell Data</u>

Maya Sheth, Wei-Lin Qiu, Rosa Ma, Andreas Gschwind, Evelyn Jagoda, Anthony Tan, Hjörleifur Einarsson, Bram Gorissen, Danilo Dubocanin, Christopher McGinnis, Dulguun Amgalan, Ansuman Satpathy, Thouis Jones, Lars Steinmetz, Anshul Kundaje, **Berk Ustun**, Jesse Engreitz, Robin Andersson *In Submission*, 2025

- 2. <u>Confirmation Effects with Interpretable Models</u> Julian Skirzynski, Elena Glassman, **Berk Ustun** *In Preparation*, 2025
- Learning from Model Databases: Map to Train, Filter to Constrain, Cast to Explain Berk Ustun, Ryan Hammonds, Tynan Seltzer, Margo Seltzer, Cynthia Rudin In Preparation, 2025

| MENTEDENIC H-:1 DID States OLD States OLICED CSE | D |
|---|-------------------|
| Lulian Shimmali: DID States (2005) CSE | 202I - PRESENT |
| Shravaa Kadakadi. DhD Student @ UCSD CSE | 2022 – PRESENT |
| Harmy Chaop, MS Student @ UCSD CSE | 2023 - PRESENT |
| Marry Cheon. MS Student @ UCSD CSE | 2024 - PRESENT |
| Meredith Stewart. Mis Student @ UCSD CSE | 2024 – PRESENT |
| Former Students & Mentees | |
| Jamelle Watson-Daniels. PhD Student @ Harvard SEAS \rightarrow Research Scientist @ Meta A | 1 2020-2024 |
| Hao Wang . PhD Student @ Harvard SEAS \rightarrow Research Scientist @ IBM Research | 2017-2020 |
| Avni Kothari. MS Student @ UCSD CSE \rightarrow Data Scientist @ UCSF | 2022 - 2023 |
| Vinith Suriyakumar. MS Student @ University of Toronto \rightarrow PhD Student @ MIT CSAII | L 2020-2022 |
| Haoran Zhang. MS Student @ University of Toronto \rightarrow PhD Student @ MIT CSAIL | 2020 — 202 I |
| Alexander Spangher. MS Student at Columbia \rightarrow PhD Student @ USC CS | 2017-2019 |
| Charles T. Marx. BS Student @ Haverford College \rightarrow PhD Student @ Stanford CS | 2019 - 2020 |
| Jiaming Zeng. BS Student at MIT \rightarrow PhD Student @ Stanford MSE | 2014-2016 |
| RESEARCH NSF CISE Core Program FUNDING RI: Foundations of Recourse Verification in Machine Learning Co-PIs: Berk Ustun, Lily Weng. | 2023 - 2026 |
| Total Funds/PI Share: \$1.2M/\$600K | |
| NSF Program on Fairness in Artificial Intelligence FAI: Foundations of Fair AI in Medicine: Ensuring the Fairness of Patient Attribute Co-PIs: Flavio du Pin Calmon, Elena Glassman, Berk Ustun . Total Funds/PI Share: \$1M/\$300K | 2021 – 2024 es |
| Amazon Research Award | 2023 - 2024 |
| Participatory Personalization for Machine Learning | 5 1 |
| PI: Berk Ustun. | |
| Total Funds: \$120K | |
| Harvard Data Science Initiative Trust in Science Award | 2021 - 2022 |
| Identifying & Correcting Confirmation Bias in Model Selection | |
| Co-PIs: Elena Glassman, Berk Ustun. | |
| Total Funds/PI Share: \$25K / \$25K | |
| NIH Bridge to Artificial Intelligence | 2022 - 2026 |
| A FAIR Bridge2AI Center | |
| Co-Investigator – Ethics Core. | |
| Total Funds/PI Share: \$4.5M/\$212.5K | |
| NSF Convergence Accelerator | 2021 - 2023 |
| AI & Community-Driven Wildland Fire Innovation via WIFIRE Commons | |
| Senior Personnel. | |
| Total Funds/PI Share: \$5M/\$100K | |
| TEACHING DSC140A – Probabilistic Modeling & Machine Learning | 2022 – Present |
| DSC190/291 – Machine Learning Competitions | 2024 – Present |
| DSC291 – Interpretability & Explainability in Machine Learning | 202 i – Present |
| DSC180A/B – Data Science Capstone | 202 I – Present |
| INVITED University of Zürich Law and Economics Seminar | 2025 |
| TALKS ETH Zürich Law and Economics Colloquium | 2025 |
| UC Berkeley Workshop on Individualized Decision-Making | 2024 |

| BIRS Workshop on Bridging Prediction and Intervention Problems in Social Systems | 2024 |
|--|------|
| Harvard University, School of Engineering and Applied Sciences | 2024 |
| Massachusetts Ethical ML in Healthcare Seminar | 2024 |
| Information Theory & Applications Workshop | 2024 |
| Stanford University Statistics Seminar | 2023 |
| Mathematical Sciences Research Institute | 2023 |
| UCSD Institute for Practical Ethics Seminar Series | 2023 |
| UCSF Bakar Computational Health Sciences Institute | 2023 |
| École Polytechnique Fédérale de Lausanne | 2023 |
| U.S. General Services Administration | 2023 |
| Information Theory & Applications Workshop | 2023 |
| NIH Bridge2AI - AI-READI Group | 2023 |
| Conference on the Mathematical Theory of Deep Neural Networks – Keynote | 2022 |
| Kavli Frontiers of Science Symposium | 2022 |
| ICML Workshop on Algorithmic Recourse – Keynote | 2021 |
| NeurIPS Workshop on Fair AI in Financial Services | 2020 |
| ICML Workshop on Participatory Approaches in ML | 2020 |
| Google AI | 2020 |
| Brown University CS Colloquium | 2020 |
| University of Washington CS Colloquium | 2020 |
| University of Maryland CS Colloquium | 2020 |
| University of Chicago CS Colloquium | 2020 |
| University of British Columbia CS Colloquium | 2020 |
| Northeastern University CS Colloquium | 2020 |
| University of Toronto Mechanical and Industrial Engineering Seminar | 2020 |
| UCSD Halicioğlu Data Science Institute | 2020 |
| Tufts University CS Colloquium | 2020 |
| Yale University CS Colloquium | 2020 |
| Georgetown CS Colloquium | 2020 |
| Duke AI+Health Seminar Series | 2020 |
| CCC Workshop on Economics & Fairness | 2019 |
| Harvard Economics & Computation Seminar Series | 2019 |
| ACM Conference on Fairness, Accountability, and Transparency | 2019 |
| INFORMS Annual Meeting | 2019 |
| NeurIPS Workshop on Human-Centric ML | 2019 |
| Data & Society Meeting on Fair ML in Health | 2019 |
| MIT Lincoln Labe AL Seminer Series | 2019 |
| Analysis Consulting Group. Data Science Seminar Series | 2019 |
| | 2019 |
| Harvard Data Science Conference | 2018 |
| Lornell University OKIE Department | 2018 |
| Tratvaru CACo Conoquium | 2018 |
| INFORMS Annual Meeting | 2017 |
| Data & Society Kesearch Institute | 2017 |
| INFORMIS Annual Iviceting | 2016 |
| UP EININ UTIHIINOIOgy UOIIOquium | 2016 |

| | American Society of Criminology Annual Meeting | 2015 |
|------------|---|---------------------|
| | ASA Conference on Statistical Learning and Data Mining | 2014 |
| | MIT Big Data Initiative | 2014 |
| | INFORMS Annual Meeting | 2014 |
| | AAAI Conference on Artificial Intelligence | 2014 |
| | | 2013 |
| Academic | Advisory Boards | |
| Activities | Steering Committee – ACM Conference on Fairness, Accountability & Transparency | 2023 - 2024 |
| | AlgoSoc Advisory Board – Dutch Ministry of Education, Culture and Science | 2022 – Present |
| | Conference Leadership | |
| | Executive Committee – ACM Conference on Fairness, Accountability & Transparency | 2024 – Present |
| | Senior Area Chair – Neural Information Processing Systems | 2024 – Present |
| | Program Chair – ACM Conference on Fairness, Accountability & Transparency | 2023 |
| | Organizer – MSRI Workshop on Randomization, Fairness, and Neutrality | 2023 |
| | Organizer – FAT/ML Workshop | 2017, 2018 |
| | Grant Reviewing | |
| | National Science Foundation Panelist, CISE | 2019, 2023 |
| | Journal Reviewing | |
| | Journal of Machine Learning Research, Management Science, IEEE Transactions on Sign | al Processing, Sta- |
| | tistical Analysis & Data Mining, Artificial Intelligence, Information Sciences, Minds & M | achines, Big Data, |
| | Epidemiology, Nature Digital Medicine, Artificial Intelligence & Law, Journal of Quantit | ative Criminology. |
| | Conference Area Chairing | |
| | NeurIPS – Neural Information Processing Systems | 2023 |
| | ICML – International Conference on Machine Learning | 2025 |
| | ICLR – International Conference on Learning Representations | 2023, 2024 |
| | FAccT – ACM Conference on Fairness, Accountability & Transparency | 2020, 2022, 2024 |
| | Conference Reviewing | |
| | NeurIPS – Neural Information Processing Systems | 2018-2022 |
| | ICML – International Conference on Machine Learning | 2019, 2022 |
| | ICLR – International Conference on Learning Representations | 2020–202 I |
| | FAccT – ACM Conference on Fairness, Accountability and Transparency | 2018-2021 |
| | AISTATS – Artificial Intelligence and Statistics | 2019 |

| AISTATS – Artificial Intelligence and Statistics |
|---|
| AAAI – Association for the Advancement of Artificial Intelligence |
| HCOMP – AAAI Conference on Human Computation and Crowdsourcing |
| UAI – Conference on Uncertainty in Artificial Intelligence |
| Petal Technologies. New York, NY |

SELECTED Petal Technologies. New York, NY PROFESSIONAL Co-Founder & Technical Advisor

EXPERIENCE Spearheaded machine learning strategy to lend responsibly to consumers without credit history in the US. Petal is a fintech startup with over 100 employees and 300K customers.

Massachusetts General Hospital. Boston, MA

Departments of Neurology, MGH Lab of Computer Science, MGH Sleep Laboratory Built customized machine learning models to (i) predict seizure risk for patients in the ICU, (ii) predict the risk of hospital readmissions, and (iii) screen obstructive sleep apnea using electronic medical records.

Amazon. Seattle, WA Research Scientist Intern, IPC Buying Strategy Team

Developed algorithms to identify complementary products sold on Amazon.com. Proposed new inventory and transportation policies for complementary products that achieved major cost savings.

2019 2019 2018

2015-2024

2014 - 2020

2013

| Selected Press & Editorials | What does a Fair Algorithm Actually Look Like?Wired, October 2018.Big Data vs. the Credit GapPolitico, February 2018.Just Data. How algorithms go bad – and how they can be saved.Colloquy, August 2018.Do You Zone Out? Procrastinate? Might Be Adult ADHDNational Public Radio, April 2017.Good News for Screening for Adult ADHDInvited Commentary in JAMA Psychiatry, April 2017.A Novel Clinical Score to Assess Seizure RiskInvited Commentary in JAMA Neurology, October 2017.Big-Data or Slim-Data: Predictive Analytics Will Rule with WorldEditorial in Journal of Clinical Sleep Medicine, February 2016. |
|-----------------------------------|--|
| | |

PERSONALLanguages: Fluent in English, French and Turkish
Interests: Basketball, Gardening, Hiking, Cooking, Photography, RPGs, Hard Science Fiction