

Sorelle A. Friedler
Shibulal Family Professor of Computer Science

Department of Computer Science, Haverford College
 370 Lancaster Ave., Haverford, PA 19041
 sorelle@cs.haverford.edu
 sorelle.friedler.net

RESEARCH INTERESTS

The **fairness, accountability, and transparency of machine learning**, in contexts from criminal justice to materials discovery and AI policy. Key papers include work on disparate impact in machine learning, accelerating materials discovery with interpretable machine learning, and understanding the sociotechnical context of fair machine learning.

APPOINTMENTS**Haverford College**

Shibulal Family Professor of Computer Science	July 2024 - present
Shibulal Family Associate Professor of Computer Science	July 2022 - June 2024
Associate Professor of Computer Science	July 2020 - June 2024
Assistant Professor of Computer Science	July 2014 - June 2020
Visiting Assistant Professor of Computer Science	July 2012 - June 2014

The Brookings Institution

Nonresident Senior Fellow	October 2024 - present
---------------------------	------------------------

Association for Computing Machinery (ACM)

Chair, U.S. Technology Policy Committee	August 2025 - present
---	-----------------------

White House Office of Science and Technology Policy

Assistant Director for Data and Democracy	July 2021 - December 2022
---	---------------------------

Data & Society Research Institute

Senior Policy Fellow	June 2023 - December 2023
Affiliate	Sept. 2016 - July 2021, 2024 - present
Fellow	Sept. 2015 - Sept. 2016

Alphabet, Inc. (formerly Google, Inc.)

Software Engineer, Search Infrastructure (Google) and X (formerly Google [x])	August 2010 - June 2012
---	-------------------------

EDUCATION

University of Maryland, College Park, MD	Ph.D.	August 2005 - August 2010
Computer Science.		
Thesis title: <i>Geometric Algorithms for Objects in Motion</i> . Advisor: David M. Mount.		

University of Maryland, College Park, MD	M.S.	August 2005 - December 2007
Computer Science.		

Swarthmore College, Swarthmore, PA	B.A.	August 2000 - May 2004
Computer Science. Minor: Mathematics.		

GRANTS

- Ford Foundation** 2025 - 2026
Core support for leading researchers to provide technical support for AI civil rights audits. Danaé Metaxa and Sorelle Friedler. \$200,000.
- NSF IIS-1955321** 2020 - 2025
III: Medium: Collaborative Research: Evaluating and Maximizing Fairness in Information Flow on Networks. Aaron Clauset, Sorelle Friedler, and Blair Sullivan, originally also including Suresh Venkatasubramanian and Carlos Scheidegger. \$995,908. (Haverford portion: \$128,670).
- DARPA Synergistic Discovery and Design (SD2)** 2018 - 2023
TA2+TA3: Discovering Reactions and Uncovering Mechanisms of Hybrid Organohalide Perovskite Formation. Joshua Schrier, Sorelle Friedler, and Alexander Norquist. \$7.4M.
- The Ford Foundation** 2021 - 2022
OSTP Data and Democracy Support. Sorelle Friedler. 185K.
- The MacArthur Foundation** 2021 - 2022
OSTP Data and Democracy Support. Sorelle Friedler. 189K.
- Mozilla Responsible Computer Science Challenge** 2019 - 2021
Responsible Problem Solving: Focusing on the societal consequences of design choices in data structures and algorithms. Sorelle Friedler and Kathi Fisler, originally also including Suresh Venkatasubramanian and Seny Kamara. \$312,000 (Haverford portion: \$83,524).
- NSF DMR-1709351** 2017 - 2021
CDS&E: D3SC: The Dark Reaction Project: A machine-learning approach to exploring structural diversity in solid state synthesis. Joshua Schrier, Sorelle Friedler, and Alexander Norquist. \$645,288.
- LinkedIn Data Access Award** 2018 - 2020
Gaining access to hard-to-reach and disadvantaged populations via controlled interventions in the economic graph. Suresh Venkatasubramanian, danah boyd, and Sorelle Friedler. Non-monetary data access grant.
- NSF IIS-1633387** 2016 - 2020
BIGDATA: Collaborative Research: F: Algorithmic Fairness: A Systemic and Foundational Treatment of Nondiscriminatory Data Mining. Suresh Venkatasubramanian, danah boyd, and Sorelle Friedler. \$953,432 (Haverford portion: \$172,742).
- Knight News Challenge Prototype Fund** 2016
Could your data discriminate? Sorelle Friedler, Wilneida Negron, Surya Mattu, Suresh Venkatasubramanian. \$35,000.
- Data & Society Research Institute Fellow** 2015 - 2016
Preventing Discrimination in Machine Learning: from theory to law and policy. \$10,000.
- NSF DMR-1307801** 2013 - 2016
The Dark Reaction Project: a machine learning approach to materials discovery. Joshua Schrier, Alexander Norquist, and Sorelle Friedler. \$299,998.

AWARDS

- Chace/Parker Teaching Award** 2019 - 2020

The Chace/Parker Annual Endowed Haverford College Teaching Award is used to honor the outstanding teaching contributions of a Haverford College professor during a given year. The award of \$3,000 is given to only one professor a year.

Ann G. Wylie Dissertation Fellowship

2009 - 2010

Provides tuition, stipend, and health insurance for one semester. Awarded to “outstanding students working on the final stages of their dissertations.”

AT&T Labs Fellowship Program

2006 - 2009

Provides tuition, stipend, health insurance, and conference funds for 3 years. Awarded to 5 “outstanding under-represented minority and women students” chosen from a national pool.

Verizon Fellowship

2006 - 2007

Monetary award for outstanding academic achievement.

Graduate School Fellow, University of Maryland

2005 - 2007

Monetary award for “academic merit, intellectual ability, and the student’s potential to make a unique contribution to the diversity of the educational experience on this campus.”

PAPERS *Author orderings are mostly alphabetical or with faculty last and students first.
Haverford undergraduate co-authors are denoted with a *.*

White House Reports

[WH] The White House Office of Science and Technology Policy. Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People. Oct. 4, 2022.

- Alondra Nelson, Sorelle Friedler, and Ami Fields-Meyer. Blueprint for an AI Bill of Rights: A Vision for Protecting Our Civil Rights in the Algorithmic Age. White House Office of Science and Technology Policy. Oct. 4, 2022.
- The White House. FACT SHEET: Biden-Harris Administration Announces Key Actions to Advance Tech Accountability and Protect the Rights of the American Public. Oct. 4, 2022.

In-progress Manuscripts

[M1] Alex Crane, Sorelle A. Friedler, Mihir Patel*, and Blair Sullivan. Equalizing Closeness Centralities via Edge Additions. May 9, 2025.

[M2] Kate Barnes, Mia Ellis-Einhorn*, Nayera Hasan*, Mohammad Fanous*, Blair Sullivan, Sorelle Friedler, and Aaron Clauset. Intervention addresses demographic and prestige disparities in academic coauthorship network. Spring 2025.

[M3] Dennis Robert Windham, Caroline J Wendt, Alex Crane, Sorelle A Friedler, Blair D Sullivan, Aaron Clauset. Fast algorithms to improve fair information access in networks. arXiv:2409.03127, Sept. 4, 2024.

Peer-reviewed Publications

[1] Sorelle A. Friedler and Andrew Selbst. The OMB Artificial Intelligence Memoranda. *Berkeley Technology Law Journal*, Forthcoming, 2025.

- [2] Grace Proebsting*, Oghenefejiro Isaacs Anigboro*, Charlie M. Crawford*, Danaé Metaxa, Sorelle A. Friedler. Identity-related Speech Suppression in Generative AI Content Moderation. *ACM Conference on Equity and Access in Algorithms, Mechanisms, and Optimization (EAAMO)*, 2025. (Acceptance rate: 30%)

★ **Best Paper Award**

- [3] Seung Hyun Cheon, Anneke Wernerfelt*, Sorelle A. Friedler, Berk Ustun. Feature Responsiveness Scores: model-agnostic explanations for recourse. *International Conference on Learning Representations (ICLR)*, 2025. (Acceptance rate: 32%)
- [4] Yaaseen Mahomed, Charlie Crawford*, Sanjana Gautam, Sorelle A. Friedler, Danaé Metaxa. Auditing GPT's Content Moderation Guardrails: Can ChatGPT Write Your Favorite TV Show? *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2024. (Acceptance rate: 18%)
- [5] Xiaorong Wang*, Clara Na, Emma Strubell, Sorelle Friedler, Sasha Luccioni. Energy and Carbon Considerations of Fine-Tuning BERT. *Conference on Empirical Methods in Natural Language Processing (EMNLP): Findings of EMNLP*, 2023. (Acceptance rate: 43.5%)
- [6] Mohsen Abbasi, Calvin Barrett*, Sorelle A. Friedler, Kristian Lum, Suresh Venkatasubramanian. Measuring and mitigating voting access disparities: a study of race and polling locations in Florida and North Carolina. *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2023. (Acceptance rate: 25%)
- [7] Ashkan Bashardoust, Sorelle A. Friedler, Carlos Scheidegger, Blair D. Sullivan and Suresh Venkatasubramanian. Reducing Access Disparities in Networks using Edge Augmentation. *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2023. (Acceptance rate: 25%)
- [8] Lydia Reader, Pegah Nokhiz, Cathleen Power, Neal Patwari, Suresh Venkatasubramanian, and Sorelle A. Friedler. Models for understanding and quantifying feedback in societal systems. *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2022. (Acceptance rate: 25%)
- [9] Venkateswaran Shekar, Gareth Nicholas*, Mansoor Ani Najeib, Margaret Zeile, Vincent Yu*, Xiaorong Wang*, Dylan Slack*, Zhi Li, Philip W. Nega, Emory Chan, Alexander J. Norquist, Joshua Schrier, and Sorelle A. Friedler. Active Meta-Learning for Predicting and Selecting Perovskite Crystallization Experiments. *The Journal of Chemical Physics*, Feb. 14, 2022.
- [10] I. Elizabeth Kumar, Carlos Scheidegger, Suresh Venkatasubramanian, and Sorelle A. Friedler. Shapley Residuals: Quantifying the limits of the Shapley value for explanations. In *Neural Information Processing Systems (NeurIPS)*, 2021. (Acceptance rate: 26%)
- [11] Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. The (im)possibility of fairness: different value systems require different mechanisms for fair decision making. *Communications of the ACM*, April, 2021.
- [12] I. Elizabeth Kumar, Suresh Venkatasubramanian, Carlos Scheidegger, and Sorelle A. Friedler. Problems with Shapley-value-based explanations as feature importance measures. In *International Conference on Machine Learning (ICML)*, 2020. (Acceptance rate: 22%)
- [13] Dylan Slack*, Sorelle A. Friedler, and Emile Givental*. Fairness Warnings and Fair-MAML: Learning Fairly with Minimal Data. In *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020. (Acceptance rate: 24%)

- [14] Xiwen Jia*, Oscar Huang*, Allyson Lynch*, Matthew Danielson*, Immaculate Lang'at*, Alexander Milder*, Aaron Ruby*, Hao Wang*, Sorelle A. Friedler, Alexander J. Norquist, and Joshua Schrier. Anthropogenic biases in chemical reaction data hinder exploratory inorganic synthesis. *Nature*, 573: 251–255, Sept. 12, 2019.
- [15] Ian Pendleton, Gary Cattabriga, Zhi Li, Mansoor Ani Najeeb, Sorelle Friedler, Alexander Norquist, Emory Chan, and Joshua Schrier. Experiment Specification, Capture and Laboratory Automation Technology (ESCALATE): A software pipeline for automated chemical experimentation and data management. *MRS Communications* (Special Issue Research Letter: Artificial Intelligence), 2019.
- [16] Harry Levin* and Sorelle A. Friedler. Automated Congressional Redistricting. *ACM Journal of Experimental Algorithmics*, 24.1 (2019): 1-10.
- [17] Charles Marx*, Richard Phillips*, Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. Disentangling Influence: Using disentangled representations to audit model predictions. In *Neural Information Processing Systems (NeurIPS)*, 2019. (Acceptance rate: 21%)
- [18] Benjamin Fish, Ashkan Bashardoust, danah boyd, Sorelle Friedler, Carlos Scheidegger and Suresh Venkatasubramanian. Gaps in Information Access in Social Networks. In *The Web Conference (WWW)*, 2019. (Acceptance rate: 18%)
- [19] Mohsen Abbasi, Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. Fairness in representation: Quantifying stereotyping as a representational harm. In *SIAM International Conference on Data Mining (SDM)*, 2019. (Acceptance rate: 23%)
- [20] Sorelle A. Friedler, Carlos Scheidegger, Suresh Venkatasubramanian, Sonam Choudhary, Evan P. Hamilton*, and Derek Roth*. A comparative study of fairness-enhancing interventions in machine learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 2019. (Acceptance rate: 24%)
- [21] Andrew Selbst, danah boyd, Sorelle A. Friedler, Suresh Venkatasubramanian, and Janet A. Vertesi. Fairness and Abstraction in Sociotechnical Systems. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 2019. (Acceptance rate: 24%)
- [22] Philip Adler, Casey Falk*, Sorelle A. Friedler, Tionney Nix*, Gabriel Rybeck*, Carlos Scheidegger, Brandon Smith*, and Suresh Venkatasubramanian. Auditing Black-box Models for Indirect Influence. *Knowledge and Information Systems*, 54.1 (2018): 95-122.

★ 2021 Knowledge and Information Systems Best Paper Award

- [23] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger, Suresh Venkatasubramanian. Decision Making with Limited Feedback: Error bounds for Recidivism Prediction and Predictive Policing. In *Algorithmic Learning Theory (ALT)* 2018. (Acceptance rate: 35%)
- [24] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger and Suresh Venkatasubramanian. Runaway Feedback Loops in Predictive Policing. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 2018. (Acceptance rate: 24%)
- [25] Richard L. Phillips*, Kyu Hyun Chang*, and Sorelle A. Friedler. Interpretable Active Learning. In *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 2018. (Acceptance rate: 24%)
- [26] Paul Raccuglia*, Katherine C. Elbert*, Philip D. F. Adler, Casey Falk*, Malia B. Wenny*, Aurelio

- Mollo*, Matthias Zeller, Sorelle A. Friedler, Joshua Schrier, and Alexander J. Norquist. Machine-learning-assisted materials discovery using failed experiments. *Nature*, 533: 73 - 76, May 5, 2016.
- [27] Philip Adler, Casey Falk*, Sorelle A. Friedler, Gabriel Rybeck*, Carlos Scheidegger, Brandon Smith*, and Suresh Venkatasubramanian. Auditing Black-box Models for Indirect Influence. In *Proceedings of the IEEE International Conference on Data Mining (ICDM)*, 2016. (Acceptance rate: 20%, acceptance as regular paper: 8.5%.)
- Expanded version available as part of: Auditing Black-box Models for Indirect Influence, *Knowledge and Information Systems*.
- [28] F. Betul Atalay, Sorelle A. Friedler, and Dianna Xu. Convex hull for probabilistic points. In Technical Papers of the 29th Conference on Graphics, Patterns and Images (*SIBGRAPI '16*), 2016. (Acceptance rate: 43%.)
- [29] Michael Feldman*, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and Removing Disparate Impact. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining (KDD)*, pages 259–268, 2015. (Acceptance rate: 19%.)
- [30] Sorelle A. Friedler and David M. Mount. A Sensor-Based Framework for Kinetic Data Compression. *Computational Geometry: Theory and Applications*, 48(3): 147 - 168, March 2015
- [31] Sorelle A. Friedler and David M. Mount. Approximation algorithm for the kinetic robust k-center problem. *Computational Geometry: Theory and Applications*, 43(6-7):572 - 586, 2010.
- [32] Sorelle A. Friedler and David M. Mount. Spatio-temporal Range Searching over Compressed Kinetic Sensor Data. In *Proc. of the European Symposium on Algorithms (ESA)*, pages 386 - 397, 2010. (Acceptance rate: 27%.)
- [33] Sorelle A. Friedler and David M. Mount. Compressing kinetic data from sensor networks. In *Proc. of the Fifth International Workshop on Algorithmic Aspects of Wireless Sensor Networks (AlgoSensors)*, pages 191 - 202, 2009. (Acceptance rate: 51%.)
- Expanded version available as part of: A Sensor-Based Framework for Kinetic Data Compression, *Computational Geometry: Theory and Applications*.
- [34] Sorelle A. Friedler, Yee Lin Tan, Nir J. Peer, and Ben Shneiderman. Enabling teachers to explore grade patterns to identify individual needs and promote fairer student assessment. *Computers & Education*, 51(4): 1467 - 1485, December 2008.

Policy-related Publications and Reports

- [P1] Sorelle A. Friedler and Marc Aidinoff. Analogies for AI Policymaking. *Washington Center for Equitable Growth*, Mar. 27, 2025.
- [P2] Aaron Klein, Cameron F. Kerry, Courtney C. Radsch, Mark MacCarthy, Sorelle Friedler, and Nicol Turner Lee. One year later, how has the White House AI Executive Order delivered on its promises?. *Brookings*, Nov. 4, 2024.
- [P3] Sorelle Friedler and Brian J. Chen. Response to OMB’s Artificial Intelligence Draft Memorandum. *Data & Society*, Dec. 5, 2023.

- [P4] Sorelle Friedler, Janet Haven, and Brian J. Chen. How the AI Executive Order and OMB memo introduce accountability for artificial intelligence. *Brookings*, Nov. 16, 2023.
- [P5] Sorelle Friedler, Ranjit Singh, Borhane Blili-Hamelin, Jacob Metcalf, and Brian J. Chen. AI Red-Teaming Is Not a One-Stop Solution to AI Harms: Recommendations for Using Red-Teaming for AI Accountability. *Data & Society Policy Brief*, Oct. 25, 2023.
- [P6] Brian J. Chen, Sorelle Friedler, and Serena Oduro. Response to “Request for Information: National Priorities for Artificial Intelligence.” *Data & Society*, July 7, 2023.
- [P7] Janet Haven and Sorelle Friedler. The Senate doesn’t need to start from scratch on AI legislation. *The Hill*, July 1, 2023.
- [P8] AI Policy and Governance Working Group. Response to the NTIA AI Accountability Policy Request for Comment. June 12, 2023.
- [P9] Sorelle Friedler, Suresh Venkatasubramanian, and Alex Engler. How California and other states are tackling AI legislation. *Brookings*, Mar. 22, 2023.
- [P10] Nicholas Diakopoulos and Sorelle Friedler. How to Hold Algorithms Accountable. *MIT Technology Review*, Nov. 17, 2016.

Workshop Papers and Technical Reports

- [TR1] Ashkan Bashardoust, Hannah C. Beilinson*, Sorelle A. Friedler, Jiajie Ma*, Jade Rousseau*, Carlos E. Scheidegger, Blair D. Sullivan, Nasanbayar Ulzii-Orshikh*, Suresh Venkatasubramanian. Information access representations and social capital in networks. arXiv:2010.12611, Oct. 16, 2023.
- [TR2] Venkateswaran Shekar, Vincent Yu, Benjamin J Garcia, David Benjamin Gordon, Gemma E Moran, David M Blei, Loïc M Roch, Alberto García-Durán, Mansoor Ani Najeeb, Margaret Zeile, Philip W Nega, Zhi Li, Mina A Kim, Emory M Chan, Alexander J Norquist, Sorelle Friedler, and Joshua Schrier. Serendipity based recommender system for perovskites material discovery: balancing exploration and exploitation across multiple models. ChemRxiv. July 19, 2022.
- [TR3] Kathi Fisler, Sorelle Friedler, Kevin Lin, Suresh Venkatasubramanian. Approaches for Weaving Responsible Computing into Data Structures and Algorithms Courses. In *Proceedings of the 53rd ACM Technical Symposium on Computer Science Education (SigCSE)*, pages 1049-1050, Mar. 3, 2022. <https://dl.acm.org/doi/abs/10.1145/3478432.3499222>
- [TR4] I. Elizabeth Kumar, Carlos Scheidegger, Suresh Venkatasubramanian, and Sorelle Friedler. Shapley Residuals: Quantifying the limits of the Shapley value for explanations. *ICML Workshop on Workshop on Human Interpretability in Machine Learning (WHI)*, 2020.

Expanded version available as part of: Problems with Shapley-value-based explanations as feature importance measures. In *International Conference on Machine Learning (ICML)*, 2020.
- [TR5] Dylan Slack*, Sorelle Friedler and Emile Givental*. Fairness Warnings. *NeurIPS Workshop on Human-Centric Machine Learning (HCML)*, 2019.
- [TR6] Dylan Slack*, Sorelle Friedler and Emile Givental*. Fair Meta-Learning: Learning How to Learn Fairly. *NeurIPS Workshop on Human-Centric Machine Learning (HCML)*, 2019.

Expanded version of above two papers available as: Fairness Warnings and Fair-MAML: Learning Fairly with Minimal Data, *Conference on Fairness, Accountability, and Transparency (FAccT)*, 2020.

- [TR7] Dylan Slack*, Sorelle A. Friedler, Chitradeep Dutta Roy, and Carlos Scheidegger. Assessing the Local Interpretability of Machine Learning Models. *NeurIPS Workshop on Human-Centric Machine Learning (HCML)*, 2019. <https://arxiv.org/abs/1902.03501>

- [TR8] Kadan Lottick*, Silvia Susai*, Sorelle Friedler, and Jonathan Wilson. Energy Usage Reports: Environmental awareness as part of algorithmic accountability. *NeurIPS Workshop on Tackling Climate Change with Machine Learning*, 2019.

- [TR9] Danielle Ensign, Sorelle A. Friedler, Scott Neville, Carlos Scheidegger and Suresh Venkatasubramanian. Runaway Feedback Loops in Predictive Policing. Presented as a talk at the *Fairness, Accountability, and Transparency in Machine Learning Workshop (FAT/ML)*, Aug. 14, 2017.

Expanded version available as part of: Runaway Feedback Loops in Predictive Policing, *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*, 2018.

- [TR10] Danielle Ensign, Sorelle Friedler, Scott Neville, Carlos Scheidegger and Suresh Venkatasubramanian. Decision Making with Limited Feedback: Error bounds for Recidivism Prediction and Predictive Policing. Presented as a poster at the *Fairness, Accountability, and Transparency in Machine Learning Workshop (FAT/ML)*, Aug. 14, 2017.

Expanded version available as part of: Error bounds for Recidivism Prediction and Predictive Policing, *Proceedings of Algorithmic Learning Theory (ALT)*, 2018.

- [TR11] Richard L. Phillips*, Kyu Hyun Chang*, and Sorelle A. Friedler. Interpretable Active Learning. Presented at the *ICML Workshop on Human Interpretability in Machine Learning (WHI)*, Aug. 10, 2017.

Expanded version available as part of: Interpretable Active Learning, *Proceedings of the Conference on Fairness, Accountability, and Transparency (FAccT)*.

- [TR12] Sorelle A. Friedler, Carlos Scheidegger, and Suresh Venkatasubramanian. On the (im)possibility of fairness. arXiv:1609.07236, Sept. 23, 2016. <http://arxiv.org/abs/1609.07236>

Modified version available as: The (im)possibility of fairness: different value systems require different mechanisms for fair decision making. *Communications of the ACM*, April, 2021.

- [TR13] Nicholas Diakopoulos, Sorelle Friedler, Marcelo Arenas, Solon Barocas, Michael Hay, Bill Howe, HV Jagadish, Kris Unsworth, Arnaud Sahuguet, Suresh Venkatasubramanian, Christo Wilson, Cong Yu, and Bendert Zevenbergen. Principles for accountable algorithms and a social impact statement for algorithms. *Dagstuhl working group write-up*. July, 2016. Available at: <https://www.fatml.org/resources/principles-for-accountable-algorithms>

How to Hold Algorithms Accountable is a description of this work for the public.

- [TR14] Ifeoma Ajunwa, Sorelle Friedler, Carlos E. Scheidegger, and Suresh Venkatasubramanian. Hiring by Algorithm: Predicting and Preventing Disparate Impact. Presented at the Yale Law School Information Society Project conference *Unlocking the Black Box: The Promise and Limits of Algorithmic Accountability in the Professions*, Apr. 2, 2016.

This is a translation of the paper *Certifying and Removing Disparate Impact* for a legal audience.

- [TR15] Michael Feldman*, Sorelle A. Friedler, John Moeller, Carlos Scheidegger, and Suresh Venkatasubramanian. Certifying and Removing Disparate Impact. Presented at the *Fairness, Accountability, and Transparency in Machine Learning Workshop (FAT/ML)*, Dec. 12, 2014. <http://arxiv.org/abs/1412.3756>

Expanded version available as part of: Certifying and Removing Disparate Impact, in *Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.

- [TR16] F. Betul Atalay, Sorelle A. Friedler, and Dianna Xu. Probabilistic Kinetic Data Structures. Presented at the *Fall Workshop on Computational Geometry*, Oct. 25, 2013. <http://arxiv.org/abs/1412.1039>

Expanded version available as part of: Convex Hull for Probabilistic Points, in *Proceedings of the Brazilian Symposium on Computer Graphics and Image Processing (SIBGRAPI)*.

- [TR17] Sorelle A. Friedler and David M. Mount. Spatio-temporal Range Searching Over Compressed Kinetic Sensor Data. *Second Workshop on Massive Data Algorithms (MASSIVE 2010)*, June 17, 2010.

Expanded version available as part of: Spatio-temporal Range Searching over Compressed Kinetic Sensor Data, in *Proc. of the European Symposium on Algorithms (ESA)*.

- [TR18] Sorelle A. Friedler and David M. Mount. Realistic Compression of Kinetic Sensor Data. *University of Maryland Computer Science Department*, Technical Report CS-TR-4959, June 6, 2010. <http://hdl.handle.net/1903/10114>

Expanded version available as part of: A Sensor-Based Framework for Kinetic Data Compression, *Computational Geometry: Theory and Applications*.

Thesis

Sorelle A. Friedler. Geometric Algorithms for Objects in Motion. Dissertation committee: Prof. David Mount (chair), Prof. William Gasarch, Prof. Samir Khuller, Prof. Steven Selden, Prof. Amitabh Varshney. Defense date: July 30, 2010.

Book Reviews

Sorelle A. Friedler. Review of *Pioneering Women in American Mathematics: the Pre-1940 PhD's* by Judy Green and Jeanne LaDuke. *SIGACT News* 42(2): 37-41, 2011.

Sorelle A. Friedler. Review of *Change is Possible: Stories of Women and Minorities in Mathematics* by Patricia Clark Kenschaft. *SIGACT News* 41(2): 47-50, 2010.

PATENTS

Sorelle Alaina Friedler, Mohammed Waleed Kadous, Andrew Lookingbill. *Position indication controls for device locations*. US 20130131973 A1 (also WO 2013078125 A1). Publication date: May 23, 2013.

Mohammed Waleed Kadous, Isaac Richard Taylor, Cedric Dupont, Brian Patrick Williams, Sorelle Alaina Friedler. *Permissions based on wireless network data*. US 20130244684 A1 (also WO2013138304 A1). Publication date: Sep 19, 2013.

KEYNOTE TALKS

Fordham Data Science Symposium: Doing good with data,
Fordham University, NY
Blueprint for an AI Bill of Rights

Apr. 11, 2023

Symp. on Data Science, AI, and a Sustainable, Resilient, and Equitable Future, Pennsylvania State University, State College, PA <i>Blueprint for an AI Bill of Rights: Making automated systems work for the American people</i>	Oct. 13, 2022
Atlantic Council AI Connect Network Webinar, remote <i>Human-centered Implementation of AI Technologies</i>	May 26, 2022
Teaching Machine Learning Workshop at ECML, remote <i>Embedding Ethics in Data Structures Classes</i>	Sept. 8, 2021
Machines, Minds, and Morality: Ethics in a Changing Technological World Rosemont College, Rosemont, PA <i>Fairness and Abstraction: algorithmic discrimination and attempts to address it</i>	April 5, 2019
Scandinavian Symposium and Workshops on Algorithm Theory (SWAT) Malmö, Sweden <i>Optimizing Society? Ensuring Fairness in Automated Decision-Making</i>	June 18, 2018
NSF Workshop on Data Science for Secure and Privacy-aware (DSSP) Large Data Management and Mining, Snowbird, UT <i>Algorithmic Fairness: Guaranteeing fairness and non-discrimination in machine-learned decision making</i>	Sept. 26, 2016
CONGRESSIONAL BRIEFINGS	
U.S. House AI Task Force, Washington, DC Closed door briefing on AI Transparency	Jun. 6, 2024
California State Assembly Committee on Privacy and Consumer Protection, remote Informational hearing on <i>Understanding AI: Myths, Magic, and Machine Learning</i>	Feb. 27, 2024
Maryland State Joint Committee on Cybersecurity, Information Technology, and Biotechnology, remote <i>State Governance Considerations on AI</i>	Oct. 25, 2023
INVITED TALKS (selected)	
Haverford College, <i>Fairness, Accountability, and Transparency in Machine Learning</i>	2025
Georgetown, <i>Making AI Work for the People: from values to (socio)technical solutions</i>	2025
Georgia Tech, <i>From Principles to Practice: federal AI policy</i>	2024
Bryn Mawr College, <i>How AI Will Be Governed</i>	2024
Swarthmore College, <i>How AI Will Be Governed</i>	2024
Lafayette College, <i>Understanding AI</i>	2024
Math Sciences Research Institute, <i>Values and Fairness Definitions</i>	2023
National Council of State Legislatures, <i>Anatomy of an AI System</i>	2023
Colby College, Distinguished Speaker Series, <i>Blueprint for an AI Bill of Rights</i>	2023
University of Washington, <i>Blueprint for an AI Bill of Rights</i>	2023
Santa Fe Institute, <i>Fairness in Networks</i>	2021
The Ohio State University, <i>Fairness in Networks</i>	2020
Saint Louis University, <i>Algorithms in Criminal Justice</i>	2020
University of Pennsylvania, <i>Algorithms in Criminal Justice</i>	2020
Cornell Tech, <i>Fairness in Networks</i>	2019

Haverford College, <i>Algorithmic Fairness</i>	2018
Columbia University, <i>Auditing, Explaining, and Ensuring Fairness</i>	2018
University of Pennsylvania, <i>Feedback Loops in Predictive Policing</i>	2018
University of Minnesota, <i>Auditing, Explaining, and Ensuring Fairness</i>	2018
Rutgers University, <i>Auditing Black-box Models</i>	2017
Brown University, <i>Algorithmic Fairness</i>	2017
Rutgers / Bryn Mawr Undergraduate Workshop, <i>Data Structures for Kinetic Points</i>	2016
Consumer Financial Protection Bureau, <i>Biased Data, Biased Algorithms</i>	2016
Williams College, <i>Biased Data, Biased Algorithms</i>	2016
Microsoft Research, <i>Detecting and Preventing Discrimination in ML</i>	2016

CONTRIBUTED TALKS (conferences and workshops)

UC Berkeley Law, <i>Research and Policy Leading to the OMB AI Guidance</i>	2025
USGS, <i>How to successfully use AI for science</i>	2025
FACCT, <i>Measuring and mitigating voting access disparities</i>	2023
SIGCSE, <i>Weaving responsible computing into data structures and algorithms</i>	2022
Networked Justice Symposium, <i>A Case for Community-Based Notions of Fairness</i>	2021
Northwestern University Law + Computation, <i>Accountability for risk assessments</i>	2021
National Academy of Sciences, <i>Interpretable ML for Scientific Understanding</i>	2019
Telluride ML for Materials Workshop, <i>Interpretable ML for Scientific Discovery</i>	2018
Google Workshop, <i>Fairness definitions: axioms and representations</i>	2018
CCC Workshop, <i>Fair representations</i>	2018
FAT/ML, <i>Runaway Feedback Loops in Predictive Policing</i>	2017
Obfuscation Workshop, <i>Obfuscating Data to Prevent Discrimination</i>	2017
AALAC Workshop on Data Ethics, <i>Algorithmic Fairness</i>	2017
Fairness for Digital Infrastructure Workshop, <i>On the (Im)possibility of Fairness</i>	2017
ICDM, <i>Auditing Black-box Models for Indirect Influence</i>	2016
SIBGRAPI, <i>Convex Hull for Probabilistic Points</i>	2016
Dagstuhl Data Responsibly, <i>Auditing Black-box Models</i>	2016
Data & Society Workshop, <i>Hiring by Algorithm</i>	2016
National Council on Measurement in Education, <i>Fairness and ML for Education</i>	2016
Yale Law, <i>Unlocking the Black Box, Hiring by Algorithm</i>	2016
FAT/ML, <i>Certifying and Removing Disparate Impact</i>	2014
Fall Workshop on Computational Geometry, <i>Probabilistic Kinetic Data Structures</i>	2013
AALAC/Mellon 23 Working Group on Information, <i>Information Content in Motion</i>	2012
ESA, <i>Spatio-temporal Range Searching Over Compressed Kinetic Sensor Data</i>	2010
MASSIVE, <i>Spatio-temporal Range Searching Over Compressed Kinetic Sensor Data</i>	2010
AlgoSensors, <i>Compressing Kinetic Data From Sensor Networks</i>	2009

POLICY-RELATED TALKS

The Spencer Foundation, <i>Can AI be a tool for equity and justice in education?</i>	2024
The Leadership Conference on Civil and Human Rights, <i>AI and Civil Rights</i>	2023
Data & Society, <i>Decoding the AI Executive Order</i>	2023
Data & Society, CDT, and William & Mary Law, <i>AI Bill of Rights: One Year Later</i>	2023
Bipartisan Policy Center, <i>The Future of AI Governance</i>	2023
FACCT, <i>Conversation with US EEOC Chair Burrows</i>	2023
FACCT, <i>AI Governance and Policy in the US - AI Bill of Rights</i>	2023
FACCT, <i>From Research Insight to Policy Impact - How to Engage in AI Policy</i>	2023

Institute for Advanced Study, <i>Steering AI for the Public Good</i>	2023
Princeton University, <i>Confused by All the Chatter? Chatbots and Other LLMs</i>	2023
UNESCO Commission on the Status of Women, <i>The Gender Digital Revolution</i>	2023
Brookings Institute, <i>Unpacking the White House Blueprint for an AI Bill of Rights</i>	2022
OECD Working Group on AI, <i>Blueprint for an AI Bill of Rights</i>	2022
NSF Fairness in AI PI Meeting, <i>Improving the Policy and Practical Impact of Fair-AI</i>	2022
Mozilla Meetups, <i>The Building Blocks of a Trusted Internet</i>	2022
Carnegie Mellon University Responsible AI Initiative, <i>Launch Event</i>	2022
SIIA, <i>Building U.S. Leadership in Responsible AI Use, fireside chat</i>	2022
Global Partnership on AI, <i>AI for Drug Discovery</i>	2021
Philadelphia Fulbright Enrichment Seminar, <i>Big Data for the Public Good</i>	2018
University of Pennsylvania, <i>Optimizing Government: Policy Challenges in the ML Age</i>	2017

PEDAGOGICAL / GENERAL AUDIENCE TALKS

FACCT, <i>Plenary Panel: Pathways of Change and the Future of Responsible AI</i>	2025
Grace Hopper, <i>Practical Steps for Tackling Bias in Software Engineering in the Age of AI</i>	2024
NSF TIP/CISE/SBE Distinguished Innovation and Entrepreneurship, <i>Ethical AI</i>	2023
KDD, <i>Fairness in Networks: Social Capital, Information Access, and Interventions</i>	2021
AALAC, <i>Data Science in the Liberal Arts, Data Science Ethics in CS2 and Algorithms</i>	2021
MozFest, <i>Responsible Computing Curricula - How do we do it?</i>	2021
Beth Am Israel, <i>Jewish Perspectives on Ethics of AI</i>	2020
St. Joseph's University, McNulty Scholars Program, <i>Fairness and Abstraction</i>	2019
IEEE Intl. Conf. on Data Science and Advanced Analytics, <i>Data Ethics</i>	2019
Simons Institute, Summer Cluster: Fairness, <i>Algorithmic Governance</i>	2019
FACCT, <i>Hands-on Tutorial: pip install fairness</i>	2019
Princeton University CITP, <i>Principles for Accountable Algorithms</i>	2017
ICDM Workshop on Privacy and Discrimination, <i>Ethical Data Mining</i>	2016
FAT/ML, <i>Opening panel</i>	2016
Dagstuhl Seminar on Data, Responsibly, <i>Teaching Data Ethics</i>	2016
NYU Law, Bernstein Institute for Human Rights, <i>Data Hygiene and Algorithmic Oversight</i>	2016
SXSW, <i>Biased Algorithms and the Future of Prejudice</i>	2016
Data & Civil Rights Conference, <i>Discriminatory Machine Learning</i>	2015
Grace Hopper, <i>Diverse Paths to Teaching and Research at LACs</i>	2015

INVITED WORKSHOPS

The Spencer Foundation, <i>Towards New Horizons of AI, Learning, and Equity in Education</i>	2024
CCC and MacArthur, <i>Community-Driven Approaches to Research in Technology & Society</i>	2023
GAO Comptroller General Forum on AI Oversight	2020
Dagstuhl Seminar on ML Meets Visualization to Make AI Interpretable	2019
Simons Institute, Summer Cluster: Fairness	2019
Harvard Center for Research on Computation and Society, <i>Ethics into CS Curricula</i>	2018
Telluride Workshop, <i>Machine Learning and Informatics for Chemistry and Materials</i>	2018
Princeton University, <i>Limits of Artificial Intelligence in Public Policy</i>	2018
Google Workshop on Fairness in Machine Learning	2018
NYU, <i>Data Science Ethics Education Workshop</i>	2018
NSF BIGDATA PI Meeting	2017
University of Pennsylvania, <i>Fairness for Digital Infrastructure</i>	2017
MacArthur Foundation, <i>Algorithmic Decision-Making Tools in Criminal Justice</i>	2016
University of Pennsylvania, <i>USACM Algorithmic Transparency and Accountability</i>	2016
Dagstuhl Seminar, <i>Data, Responsibly</i>	2016
Data & Civil Rights Conference	2015
NSF Workshop on the Rise of Data in Materials Research	2015

SELECTED PRESS

Related to *Blueprint for an AI Bill of Rights*:

- Cristiano Lima. White House unveils 'AI bill of rights' as 'call to action' to rein in tool. *The Washington Post*, Oct. 4, 2022.
- Melissa Heikkilä. The White House just unveiled a new AI Bill of Rights: It's the first big step to hold AI to account. *MIT Technology Review*, Oct. 4, 2022.
- Angus Loten. White House Issues 'Blueprint for an AI Bill of Rights'. *The Wall Street Journal*, Oct. 4, 2022.
- Rachel Metz. The White House released an 'AI Bill of Rights'. *CNN Business*, Oct. 5, 2022.
- Garance Burke. White House unveils artificial intelligence 'Bill of Rights'. *AP News*, Oct. 4, 2022.

Related to *Anthropogenic biases in chemical reaction data hinder exploratory inorganic synthesis*:

- Kira Welter. Human biases cause problems for machines trying to learn chemistry. *Chemistry World*, Sept. 13, 2019.
- Bob Yirka. Chemists show how bias can crop up in machine learning algorithm results. *Phys.org*, Sept. 12, 2019.
- Sam Lemonick. Machine learning can have human bias: Algorithm performance suffers when humans choose how to train it. *Chemical & Engineering News*, Sept. 11, 2019.
- Editorial. Look out for potential bias in chemical data sets. *Nature*, Sept. 11, 2019.

Related to *Runaway feedback loops in predictive policing*:

- Ethan Baron. Bay Area police try out controversial AI software that tells them where to patrol. *The Mercury News*, Mar. 10, 2019.
- Caroline Haskins. Academics Confirm Major Predictive Policing Algorithm is Fundamentally Flawed. *Motherboard*, Feb. 14, 2019.
- Daniel Cossins. Discriminating algorithms: 5 times AI showed prejudice. *New Scientist*, Apr. 12, 2018.
- Matt Reynolds. Biased policing is made worse by errors in pre-crime algorithms. *New Scientist*, Oct. 4, 2018.

Related to *Fairness and abstraction in sociotechnical systems*:

- Tristan Greene. Why the criminal justice system should abandon algorithms. *The Next Web*, Feb 6, 2019.
- Karen Hao. This is how AI bias really happens — and why it's so hard to fix. *Technology Review*, Feb. 4, 2019.

Related to *Machine-learning-assisted materials discovery using failed experiments*:

- Adam Marcus and Ivan Oransky. What scientists could learn from startups. *The Week and STAT*, May 12, 2016.
- Daniela Hernandez. Why Machines Should Learn From Failures. *The Wall Street Journal*, May 6, 2016.
- Jordana Cepelewicz. Lab Failures Turn to Gold in Search for New Materials. *Scientific American*, May 6, 2016.
- Philip Ball. Computer gleans chemical insight from lab notebook failures. *Nature News*, May 4, 2016.

Related to *On the (im)possibility of fairness*:

- Jordan Pearson. To Make AI Less Biased, Give It a Worldview. *Motherboard*, Sept. 27, 2016.

Related to *Certifying and removing disparate impact*:

- Lauren J. Young. Computer Scientists Find Bias in Algorithms. *IEEE Spectrum*, August 21, 2015.
- Julianne Pepitone. Can Resume-Reviewing Software Be As Biased As Human Hiring Managers? *NBC News*, August 17, 2015.

Kiona Smith-Strickland. Computer Programs Can Be as Biased as Humans. *Gizmodo*, August 16, 2015.

Background on Algorithmic Fairness:

Megan Rose Dickey. Algorithmic accountability: Algorithms are designed to make our lives easier. The problem is, they're designed by us. *TechCrunch*, Apr 30, 2017.

Stephanie Pappas. Bad News: Artificial Intelligence Is Racist, Too. *Live Science*, Apr 13, 2017.

Nidhi Subbaraman. Scientists Taught A Robot Language. It Immediately Turned Racist. *BuzzFeed News*, Apr 13, 2017.

Jeremy Hsu. AI Learns Gender and Racial Biases from Language. *IEEE Spectrum*, Apr 13, 2017.

Sam Levin. A beauty contest was judged by AI and the robots didn't like dark skin. *The Guardian*, September 8, 2016.

David Ingold and Spencer Soper. Amazon Doesn't Consider the Race of Its Customers. Should It? *Bloomberg*, April 21, 2016.

Rose Eveleth. The Inherent Bias of Facial Recognition. *Motherboard*, March 21, 2016.

Laura Sydel. Can Computer Programs be Racist and Sexist? *NPR*, March 15, 2016.

Lauren Kirchner. When big data becomes bad data. *ProPublica*, September 2, 2015.

Hal Hodson. No one in control: The algorithms that run our lives. *New Scientist*, February 4, 2015.

PUBLIC SERVICE

White House Office of Science and Technology Policy, July 2021 - December 2022

Advisory Board Member

Leadership Conference on Civil and Human Rights, Center for Civil Rights and Technology, 2024 - present

Electronic Privacy Information Center (EPIC), 2024 - present

Informal Expert Advice

Private conversations held with U.S. federal or state congressional offices to inform proposed legislation.

Offices of U.S. Senators Lujan, Markey, Padilla, Peters, Sanders, Schumer

Offices of U.S. Congressional Representatives Beyer and Jacobs

Committee staff for the Senate Committee on Homeland Security and Government Affairs, the Joint Economic Committee, the House Energy and Commerce Committee, and the House Science Committee

California State Assemblymember Rebecca Bauer-Kahan

California Assembly Privacy and Consumer Protection Committee

Maryland State Senator Katie Fry Hester

Research Advisory Council Member

Arnold Ventures' Advancing Pretrial Policy and Research, Fall 2020 - July 2021

First Judicial District of Pennsylvania (Philadelphia Courts), Fall 2019 - July 2021

Committee Member

ACM Tech Policy Briefs Committee, Fall 2024 - present

PROFESSIONAL SERVICE

Co-Founder

Conference on Fairness, Accountability, and Transparency (FAccT)

Executive Committee Member

2018 – 2020, 2025 – present, Conference on Fairness, Accountability, and Transparency (FAccT)

Program Committee Co-chair

2018 Conference on Fairness, Accountability, and Transparency (FAccT)
2015 and 2016 Workshops on Fairness, Accountability, and Transparency in Machine Learning

Guest Editor

Big Data, "Special Issue on Social and Technical Trade-Offs," 2017

Program Committee Area Chair

2023, 2024 Conference on Fairness, Accountability, and Transparency (FAccT)
2020 International Conference on Machine Learning (ICML)
2020 Black in AI Workshop

Program Committee Member

2021 International Conference on Machine Learning (ICML)
2020 Conference on Fairness, Accountability, and Transparency (FAccT)
2019, 2020 Conference on Neural Information Processing Systems (NeurIPS)
2017 – 2020 ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD), Applied Data Science Track
2019, 2018 Black in AI Workshop
2019 AAAI Conference on Artificial Intelligence
2018, 2017 Workshop on Ethics in Natural Language Processing
2018 SIAM Algorithm Engineering and Experiments (ALENEX)
2017 International Conference on Information and Knowledge Management (CIKM)
2017 Workshop on Fairness, Accountability, and Transparency in Machine Learning (FAT/ML)
2017 Workshop on Technology and Consumer Protection (ConPro)
2017, 2016, 2015 SIAM International Conference on Data Mining (SDM)

Workshop Co-organizer

2021 AALAC Workshop on Data Science in the Liberal Arts
2017, 2016, 2015 Workshops on Fairness, Accountability, and Transparency in Machine Learning
2017 AALAC Workshop on Data Ethics

Committee Member

2017, 2016 CRA Outstanding Undergraduate Researchers Selection Committee
2015 ACM Student Research Competition poster judge at Grace Hopper

External Reviewer

2023 External Reviewer, University of Richmond Computer Science Department

External Thesis Committee Member

Andrew Bell, Ph.D., New York University (advisor: Julia Stoyanovich)
Yaaseen Mahomed, M.S., University of Pennsylvania (advisor: Danaë Metaxa)
Ke Yang, Ph.D., Drexel University (advisor: Julia Stoyanovich)

HAVERFORD COLLEGE TEACHING AND SERVICE

Haverford College is a leading liberal arts college with a student body of 1,470 composed entirely of undergraduate students. The teaching load is 5 course units per year, with lab / discussion sections and senior thesis advising counting towards the total teaching load. The college operates under a philosophy of shared governance among the administration, faculty, students, and staff.

UNDERGRADUATE THESIS ADVISEES by graduation year

- 2025 Amina Ahmed (Bryn Mawr), *Quantifying Impact: The Role of Scholarly Work in Shaping Government Policy*
 Jadyn Elliot, *Cultural Bias in Large Language Models: A Study of Character Dynamics in AI-Generated TV Scripts*
 Harry Genth, *Reassessing Impact: Citation Counts, Perceived Value, and the Role of Institutional Prestige*
 Joselyne Malan, *Understanding the Influential Factors of a Co-Authorship Network In the Faculty of the Computer Science Community*
 Mihir Patel, *Improving Network Fairness With Edge Addition*
 Grace Proebsting, *Investigating LLM Sycophancy with Linear Probing*
- 2024 Charlie Crawford, *Scene and Unseen: GPT Bias in Script Writing*
 Mia Ellis-Einhorn, *Influence and Equity in Faculty Co-Authorship Networks*
 Annie Wernerfelt, *Model feature importance scores should reflect recourse*
- 2023 Jade Rousseau, *Fairness in Information Access: Emphasizing the Network*
 Xiaorong (Sharon) Wang, *Pre-training and Fine-tuning BERT: Energy and Carbon Considerations*
- 2022 Jiajie (Jason) Ma, *Information Access Representation of Social Networks: A Theoretical Analysis*
 Nasanbayar Ulzii-Orshikh, *Community-based Fairness in Networks*
- 2021 Zachary Broadman, *Using Information Access to Characterize and Group Social Media Users*
 Monique Byars, *Inclusivity and Transparency in Machine Learning Model Auditing*
 Steve Lee, *Identity and Computer Science: A Mismatch?*
 Ruiming (Ray) Li, *Quantifying Uncertainty in Shapley-value-based Explanations for Machine Learning Models*
 Jason Ngo, *Understanding Machine Learning Models Through Shapley Values and the Associated Uncertainties*
 Vincent Yu, *Recommender Systems for Scientific Explorations*
- 2020 Hannah Beilinson, *Fairness and Information Access Clustering in Social Networks*
 Emile Givental, *Interpretable Meta Learning*
 Kadan Lottick, *Energy Consumption in Machine Learning*
 Charlie Marx, *Indirect Influence and Fairness in Machine Learning*
 Gareth Nicholas, *Active Meta-Learning*
 Silvia Susai, *Analyzing Energy Efficiency in Neural Networks*
- 2019 Yasmine Ayad, *Analyzing the COMPAS Algorithm in Criminal Defendant Risk Assessment*
 Arthur Chang, *Intersectionality and Fairness in Machine Learning*
 Yilin Li, *Adversarial Examples under Fairness Constraints*
 Jai Nimgaonkar, *Re-evaluation of the ProPublica Article on Machine Bias*
 Dylan Slack, *Expert-Assisted Transfer Reinforcement Learning*
 Chris Villalta, *State Influence Calculations for Deep Q-Networks*
- 2018 Richard Phillips, *Explaining Neural Network Predictions with Image Certainty Translation*
 Derek Roth, *A Comparison of Fairness-Aware Machine Learning Algorithms*
- 2017 Kyu Chang, *Explaining Active Learning Queries*
 Evan Hamilton, *Benchmarking Fairness Aware Machine Learning*
 Tionney Nix, *A Rule Learning Approach to Discovering Contexts of Discrimination*
- 2016 Casey Falk, *Auditing Deep Neural Networks and Other Black-box Models*
 Jason Feinberg, *k-Robust Nearest Neighbor Search and Classification*
 Brian Guggenheimer, *The Red Pen: Applying Computer Vision to Automate the Grading of Traditional Assignments*
 Geoffrey Martin-Noble, *Optimizing a Machine Learning System for Materials Discovery*
 Gabriel Rybeck, *Indirect Discrimination in the Age of Big Data*
 Brandon Smith, *Auditing Deep Neural Networks to Understand Recidivism Prediction*

- 2015 Michael Feldman, *Computational Fairness: Preventing Machine-Learned Discrimination*
Aaron Lowe, *Persistence in Learning: Persistent Homology and its Application to Machine Learning*
- 2014 Paulina Cueto, *Identifying the Relationship Between Evolutionary Distance and the Accuracy of Cis-Regulatory Module Predictions*
Harry Levin, *Computerized Redistricting: Examining the Weighted Points Version of the Capacitated k -Center Problem*
Karl Moll, *Community Detection in Multidimensional Social Networks*
Paul Raccuglia, *Dark Reactions: Recommender Guided Materials Discovery*
Yingying (Daisy) Sheng (Bryn Mawr College), *A Practical Evaluation of Kinetic Data Structure on Android Devices*

UNDERGRADUATE RESEARCH STUDENTS

Class of 2028: Sushmit Chakma. Class of 2027: Fejiro Anigboro, Daniel Elleston. Class of 2026: David Dai, Nayera Hasan, Class of 2025: Jady Elliot, Mohammad Fanous, Mihir Patel. Class of 2024: Charlie Crawford, Mia Ellis-Einhorn. Class of 2023: Jade Rousseau, Sharon Wang. Class of 2022: Calvin Barrett, Isaac Chang, William Lawrence, Joseph Kawamura, Iryna Khovryak, Jason Ma, Femi Obiwumi, Nasanbayar Ulzii-Orshikh, Ziyao Wang. Class of 2021: Eniola Ajao, Haosong Huang, Steve Lee, Ruiming (Ray) Li, Kaito Nakatani, Jason Ngo, Jan Estrada Pabón, Lizzie Spano, Vincent Yu, Ivy Zhang. Class of 2020: Hannah Beilinson, Emile Givental, Kadan Lottick, Charles Marx, Gareth Nicholas, Ben Rogers-Boehme, Matthew Scharf, Silvia Susai. Class of 2019: Monique Byars, Yutong Li, Jai Nimgaonkar, Dylan Slack, Christopher Villalta. Class of 2018: Tosin Alliyu, Skyler Ellenburg, Dylan Emery, Richard Phillips. Class of 2017: Tionney Nix, Derek Roth, Nora Tien, Daniel Washburn. Class of 2016: Casey Falk, Jason Feinberg, Arthur Emidio Teixeira Ferreira, Brian Guggenheimer, Geoffrey Martin-Noble, Joshua Serota. Class of 2014: Paul Raccuglia.

COURSES DEVELOPED AND TAUGHT

Haverford College, 2012 – present

*Yearly course load is 5 courses, including credit for senior thesis supervision and instruction of lab and discussion sections attached to each course. Courses added to the Haverford curriculum indicated with **, courses substantially redesigned indicated with *. Courses at the 300-level are equivalent to 400-level university courses elsewhere.*

CS 101 Fluency with Information Technology

A general introduction to computing, including: understanding news related to technology, understanding of the basics of how computer hardware and software work, internet literacy and website creation, and basic programming and database skills.

CS 104 Topics in Introductory Programming **

Topics in Introductory Programming is designed to give a general introduction to programming as related to data analysis across many fields. Students will be introduced to standard introductory programming imperative and object oriented techniques as well as data structures necessary to create efficient and understandable algorithmic solutions to problems. Data for analysis will be drawn from a single discipline that will vary per semester, forming a theme for topical study. Topical investigations will include the ethics of data use in that field, how data is commonly generated and used, and implementation of important discipline-specific algorithms.

CS 105 Introduction to Computer Science

Introduction to the intellectual and software tools used to create and study algorithms: formal and informal problem specification; problem solving and algorithm design techniques; reliability, formal verification, testing, and peer code review techniques; program clarity, complexity and

efficiency; functional and imperative paradigms; associated programming skills.

CS 106 Introduction to Data Structures *

An introduction to the fundamental data structures of computer science: strings, lists, stacks, queues, trees, BSTs, graphs, hashables, and their accompanying algorithms. Principles of algorithmic analysis and object-oriented reasoning and design will be introduced. More practical issues, such as memory management, will also be covered.

CS 207 Data Science and Visualization **

An introduction to techniques for the automated and human-assisted analysis of data sets. These “big data” techniques are applied to data sets from multiple disciplines and include cluster, network, and other analytical methods paired with appropriate visualizations.

CS 266 Artificial Intelligence and Society **

An introduction to and analysis of the impacts of artificial intelligence on society, including ethical, historical, policy, and technical perspectives on AI. Course topics will include: defining intelligence, algorithmic discrimination, AI harms, data, physical infrastructure and environmental impacts, workers, and the future of AI. Students will also do laboratory work related to technical and policy approaches to these issues including transparency mechanisms, AI audits, fair machine learning, and other mitigations for AI harms.

CS 340 Analysis of Algorithms *

Analysis of algorithms and their corresponding data structures from a precise mathematical point of view, including performance bounds, asymptotic and probabilistic analysis, and correctness. Algorithm types include greedy, dynamic programming, divide and conquer, network flow algorithms, and linear programming, as well as approximation algorithms. NP-completeness and associated reductions are also covered.

CS 360 Machine Learning

In this course, we will explore both classical and modern approaches, with an emphasis on theoretical understanding. There will be a significant math component (statistics and probability in particular), as well as a substantial implementation component (as opposed to using high-level libraries). However, during the last part of the course we will use a few modern libraries such as TensorFlow and Keras. By the end of this course, you should be able to form a hypothesis about a dataset of interest, use a variety of methods and approaches to test your hypothesis, and be able to interpret the results to form a meaningful conclusion. We will focus on real-world, publicly available datasets, not generating new data.

CS 395 Mobile Development for Social Change **

An advanced course focusing on standard software engineering principles, object oriented programming, event-driven and multi-threaded programming, Android-specific mobile development concepts, and designing a positive user experience in the context of a semester-long placement with a local non-profit.

CS 399 Senior Thesis Seminar *

Seminar required for seniors writing theses, dealing with the oral and written exposition of advanced material.

CS 480 Independent Research **

Independent study research course that does not include additional students doing research for senior thesis and does not count towards the total teaching load.

University of Maryland, College Park, 2005 – 2010**CMSC 330 Organization of Programming Languages**

A study of programming languages, including their syntax, semantics, and implementation. Several different models of languages are discussed, including dynamic, scripting (e.g., Ruby, Python) functional (e.g., OCaml, Haskell, Scheme), and memory safe systems programming (e.g., Rust). Explores language features such as formal syntax, scoping and binding of variables, higher-order programming, typing, and type polymorphism.

CMSC 451 Design and Analysis of Computer Algorithms

This course presents fundamental techniques for designing efficient computer algorithms, proving their correctness, and analyzing their performance. Topics to be covered include graph algorithms, greedy algorithms, divide-and-conquer algorithms, dynamic programming, network flow algorithms, computational intractability, approximation algorithms, and randomized algorithms.

CMSC 212 Introduction to Low-Level Programming Concepts (TA, discussion section leader)**CMSC 311 Computer Organization (TA, discussion section leader)****Springside School, Philadelphia, middle school mathematics, 2004 – 2005****COLLEGE SERVICE****Appointed College-wide Service**

Kim Institute for Ethical Inquiry and Leadership, co-director	Summer 2024 - present
Strategic Curriculum and Personnel Committee (elected)	Fall 2024 - Spring 2026
— Creation and piloting of enrollment data analysis	
— Routine allocation of tenure-track positions	
— Study and proposal of streamlined search procedures	
KINSC Steering Committee	Spring 2024
Faculty Affairs and Planning Committee (elected)	Fall 2020 - Spring 2021
— EPC restructuring into ECC and SCPC	
— Fall 2020 strike resolution support	
— Faculty salary study and systematized comparison structure	
— Additional faculty governance issues	
STEM Forum / HHMI CS departmental representative	Spring 2019 - present
Haverford Innovation Program Advisory Committee	Fall 2018 - Spring 2020
Hurford Center for the Arts and Humanities Steering Committee	Fall 2019 - Spring 2020
CRAFT Advisory Committee	Fall 2019 - Spring 2020
Computer Science Search Committee	Fall 2018 - Spring 2019
Haverford Innovation Program Staff Search Committee, chair	Fall 2016
Computational Studies Working Group, member	Fall 2016 - Spring 2017
Panel for Cases of Sexual and Racial Harassment, Faculty representative	Fall 2016 - Spring 2017
Visual Studies Search Committee, member	Fall 2015 - Spring 2016
Visual Studies Working Group, member	Fall 2014 - Spring 2015

Departmental Service

Departmental Diversity Coordinator	Fall 2020 - present
Computer Science Departmental Search Committee	Spring 2015, 2017, 2019

Advising Fall 2014 - present

Pre-major (freshman and sophomore) and declared computer science major (junior and senior) academic advisees each graduating year.