15TH ANNUAL PRIVACY PAPERS FOR POLICYMAKERS

2024

March 12, 2025

We are pleased to introduce FPF's 15th annual Privacy Papers for Policymakers. Each year we invite privacy scholars and authors to submit scholarship for consideration by a committee of reviewers and judges from the FPF Advisory Board. The selected papers are those judged to contain practical analyses of emerging issues that policymakers in Congress, in federal agencies, at the state level, and internationally will find useful.

This year's winning papers examine a variety of topical privacy issues:

- One paper investigates the motivations and purposes behind China's data protection regime, providing insights not only through an evaluation of China's data privacy law but also the historical context within which this data protection framework operates.
- Another paper analyzes the practice of data scraping for training generative AI systems within the context of privacy law. It argues that although scraping enables web searching, archival, and meaningful scientific research, scraping for AI can also be objectionable or even harmful.
- The third winning paper analyzes how the FTC's enforcement authority can address data-driven harms and identifies four influential forces that determine the "window" of FTC privacy enforcement possibility.
- Another paper evaluates AI ethics as a reflection of human and societal fairness and bias, offering a new perspective for analyzing AI technologies.
- Another winning paper provides methodologies, guidance, and case studies for practitioners tasked with undertaking fairness and equity assessments related to the development or use of AI systems.
- The sixth winning paper analyzes the value of "personhood credentials" — digital credentials that empower users to demonstrate that they are real people — to address the challenges of fraudulent identities online, especially in a world of increasingly capable AI tools.

For the ninth year in a row, we are proud to continue highlighting student work by honoring an excellent student paper: *Data Subjects' Reactions to Exercising Their Right of Access*.

We thank the scholars, advocates, and Advisory Board members who are engaged with us to explore the future of privacy.

Alan Raul
Board President,
FPF Board of Directors

Jules Polonetsky
CEO

Table of Contents

Awarded Papers

Authoritarian Privacy 4

The Great Scrape: The Clash Between Scraping and Privacy 6

Mirror, Mirror, on the Wall, Who's the Fairest of Them All? 8

The Overton Window and Privacy Enforcement 10

Navigating Demographic Measurement for Fairness and Equity 12

Personhood Credentials: Artificial Intelligence and the Value Of Privacy-Preserving Tools To Distinguish Who Is Real Online 14

Honorable Mentions

Aligning Algorithmic Risk Assessments with Criminal Justice Values 16

The Law of AI for Good 18

Awarded Student Paper

Data Subjects' Reactions to Exercising their Right of Access 20

Student Paper Honorable Mention

Artificial Intelligence is like a Perpetual Stew 22

Out of respect for copyright law and for ease of reference, this compilation is a digest of the papers selected by the Future of Privacy Forum Advisory Board and does not contain full text. The selected papers in full text are available through the referenced links.

Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Future of Privacy Forum.

Authoritarian Privacy

Mark Jia

The University of Chicago Law Review

Available at: https://lawreview.uchicago.edu/print-archive/authoritarian-privacy

Executive Summary

Privacy laws are traditionally associated with democracy. Yet autocracies increasingly have them. Why do governments that repress their citizens also protect their privacy? This Article answers this question through a study of China. Chinais a leading autocracy and the architect of a massive surveillance state. But China is also a major player in data protection, having enacted and enforced a number of laws on information privacy. To explain how this came to be, the Article first discusses several top-down objectives often said to motivate China's privacy laws: advancing its digital economy, expanding its global influence, and protecting its national security. Although each has been a factor in China's turn to privacy law, even together, they tell only a partial story.

Central to China's privacy turn is the party-state's use of privacy law to shore up its legitimacy amid rampant digital abuse. China's whiplashed transition into the digital age has given rise to significant vulnerabilities and dependencies for ordinary citizens. Through privacy law, China's leaders have sought to interpose themselves as benevolent guardians of privacy rights against other intrusive actors — individuals, firms, and even state agencies and local governments. So framed, privacy law can enhance perceptions of state performance and potentially soften criticism of the center's own intrusions. The party-state did not enact privacy law despite its surveillance state; it embraced privacy law to maintain it. This Article adds to our understanding of privacy law, complicates the relationship between privacy and democracy, and points toward a general theory of authoritarian privacy.

Author

Mark Jia is a scholar of comparative and transnational law, with particular focus on the United States and China. His research broadly seeks to understand the relationship between law and authoritarianism and between law and geopolitics. Recent works have addressed questions of constitutional law, international law, privacy law, legal interpretation, and legal theory. Professor Jia's scholarship has been or will be published in the University of Chicago Law Review, the New York University Law Review, the University of Pennsylvania Law Review, the Texas Law Review, and other journals. Before joining the academy, Professor Jia was an appellate lawyer and law clerk to Justice David Souter and Justice Ruth Bader Ginsburg of the U.S. Supreme Court and Judge William Fletcher of the U.S. Court of Appeals for the Ninth Circuit. He is a graduate of Princeton University, Oxford University, where he studied as a Rhodes Scholar, and Harvard Law School, where he was an articles co-chair of the Harvard Law Review.

The Great Scrape: The Clash Between Scraping and Privacy

Daniel J. Solove and Woodrow Hartzog

California Law Review, Vol. 113, (forthcoming 2025)

Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4884485

Executive Summary

Artificial intelligence (AI) systems depend on massive quantities of data, often gathered by "scraping" — the automated extraction of large amounts of data from the internet. A great deal of scraped data is about people. This personal data provides the grist for AI tools such as facial recognition, deep fakes, and generative AI. Although scraping enables web searching, archival, and meaningful scientific research, scraping for AI can also be objectionable or even harmful to individuals and society. Organizations are scraping at an escalating pace and scale, even though many privacy laws are seemingly incongruous with the practice. In this Article, we contend that scraping must undergo a serious reckoning with privacy law. Scraping violates nearly all of the key principles in privacy laws, including fairness; individual rights and control; transparency; consent; purpose specification and secondary use restrictions; data minimization; onward transfer; and data security. With scraping, data protection laws built around these requirements are ignored. Scraping has evaded a reckoning with privacy law largely because scrapers act as if all publicly available data were free for the taking. But the public availability of scraped data shouldn't give scrapers a free pass. Privacy law regularly protects publicly available data, and privacy principles are implicated even when personal data is accessible to others. This Article explores the fundamental tension between scraping and privacy law. With the zealous pursuit and astronomical growth of AI, we are in the midst of what we call the "great scrape." There must now be a great reconciliation.

Authors

Daniel J. Solove is the Eugene L. and Barbara A. Bernard Professor of Intellectual Property and Technology Law at the George Washington University Law School. He is the co-director of the GW Center for Law & Technology and is the director of the Privacy and Technology Law Program.

One of the world's leading experts in privacy law, Solove is the author of 10+ books and 100+ articles. He has published books with Oxford, Harvard, and Yale University Presses, and articles in the Harvard, Stanford, and Columbia Law Reviews, among others. His works have been translated into many languages.

Solove founded two companies, one that provides privacy training to organizations and another, TeachPrivacy, that involves education, events, and certification to privacy professionals. He founded the Privacy Law Scholars Conference, the largest academic conference in privacy law. He served as co-reporter for the ALI's Principles of Law, Data Privacy.

A graduate of Yale Law School, Solove clerked for Judge Stanley Sporkin, U.S. District Court for the District of Columbia and Judge Pamela Ann Rymer, U.S. Court of Appeals for the 9th Circuit. He also was an associate at Arnold & Porter LLP and a senior policy advisor at Hogan Lovells LLP.

Solove has been interviewed and quoted in hundreds of media articles and broadcasts. He has more than 1 million LinkedIn followers. He has written a children's fiction book about privacy. He has been a consultant for many Fortune 500 companies and celebrities. He is the most cited law professor born after 1970 and the most cited law professor in the law and technology field.

Woodrow Hartzog is a Professor of Law and Class of 1960 Scholar at Boston University School of Law. He is also a Faculty Associate at the Berkman Klein Center for Internet & Society at Harvard University, a Non-resident Fellow at The Cordell Institute for Policy in Medicine & Law at Washington University, and an Affiliate Scholar at the Center for Internet and Society at Stanford Law School. He is the author of Privacy's Blueprint: The Battle to Control the Design of New Technologies, published in 2018 by Harvard University Press, and the co-author of Breached! Why Data Security Law Fails and How to Improve It, published in 2022 by Oxford University Press.

Mirror, Mirror, on the Wall, Who's the Fairest of Them All?

Alice Xiang

Daedalus, the Journal of the American Academy of Arts & Sciences, Vol. 153, 2024

Available at: https://direct.mit.edu/daed/article/153/1/250/119940/Mirror-Mirror-on-the-Wall-Who-s-the-Fairest-of

Executive Summary

Debates in AI ethics often hinge on comparisons between AI and humans: which is more beneficial, which is more harmful, which is more biased, the human or the machine? These questions, however, are a red herring. They ignore what is most interesting and important about AI ethics: AI is a mirror. If a person standing in front of a mirror asked you, "Who is more beautiful, me or the person in the mirror?" the question would seem ridiculous. Sure, depending on the angle, lighting, and personal preferences of the beholder, the person or their reflection might appear more beautiful, but the question is moot. AI reflects patterns in our society, just and unjust, and the worldviews of its human creators, fair or biased. The question then is not which is fairer, the human or the machine, but what can we learn from this reflection of our society and how can we make AI fairer? This essay discusses the challenges to developing fairer AI, and how they stem from this reflective property.

Author

Alice Xiang is the Global Head of AI Ethics at Sony. As the Vice President responsible for AI governance across Sony Group, she leads the team that guides the establishment of AI policies and governance frameworks across Sony's business units. Sony is one of the world's largest manufacturers of consumer and professional electronics products, the largest video game console company and publisher, and one of the largest music companies and film studios. In addition, as the Lead Research Scientist for AI ethics at Sony AI, Alice leads a lab of AI researchers working on cutting-edge research to enable the development of more responsible AI solutions.

Alice previously served as a General Chair for the ACM Conference on Fairness, Accountability, and Transparency (FAccT), the premier multidisciplinary research conference on these topics, and is currently a Steering Committee member. Alice also previously was a member of the leadership team of the Partnership on AI. As the Head of Fairness, Transparency, and Accountability Research, she led a team of interdisciplinary researchers and a portfolio of multi-stakeholder research initiatives. She also served as a Visiting Scholar at Tsinghua University's Yau Mathematical Sciences Center, where she taught a course on Algorithmic Fairness, Causal Inference, and the Law.

She has been quoted in the Wall Street Journal, MIT Tech Review, Fortune, Yahoo Finance, and VentureBeat, among others. She has given guest lectures at the Simons Institute at Berkeley, USC, Harvard, SNU Law School, among other universities. Her research has been published in top machine learning conferences, journals, and law reviews.

Alice is both a lawyer and statistician, with experience developing machine learning models and serving as legal counsel for technology companies. Alice holds a Juris Doctor from Yale Law School, a Master's in Development Economics from Oxford, a Master's in Statistics from Harvard, and a Bachelor's in Economics from Harvard.

The Overton Window and Privacy Enforcement

Alicia Solow-Niederman

Harvard Journal of Law & Technology, Vol. 37, 2024

Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4627376

Executive Summary

On paper, the Federal Trade Commission's consumer protection authority seems straightforward: the agency is empowered to investigate and prevent unfair or deceptive acts or practices. This flexible and capacious authority, coupled with the agency's jurisdiction over the entire economy, has allowed the FTC to respond to privacy challenges both online and offline. The contemporary question is whether the FTC can draw on this same authority to curtail the data-driven harms of commercial surveillance or emerging technologies like artificial intelligence.

This Essay contends that the legal answer is yes and argues that the key determinants of whether an agency like the Federal Trade Commission will be able to confront emerging digital technologies are social, institutional, and political. Specifically, it proposes that the FTC's privacy enforcement occurs within an "Overton Window of Enforcement Possibility." Picture the FTC Act's legal standards as setting forth a range of lawful enforcement behavior for the agency — a range within which further choices must be made. Within this lawful space, just as a politician's "Overton Window of Political Possibility" will not include every possible policy option, the agency's Window will not include every possible enforcement option. Rather, the Window for privacy enforcement — the space within which the agency might operate — will be sharply informed by four critical forces: social norms; institutional norms within the agency; the courts; and Congress.

This approach highlights how the agency's enforcement actions do not occur in a rigidly fixed domain; rather, they unfold within a dynamic space that can change over time, subject to both forces inside the agency and external to it. What's more, understanding enforcement as a process in this way surfaces an often-overlooked point for federal legislation that seeks to endow new or existing agencies with additional regulatory authority: without a sufficiently large Window within which the agency can operate, all the theoretical grants of power in the world will have little impact on the ground. That's a sobering lesson. But it's empowering, too. For one, it suggests strategies for administrative officials who seek to exercise their enforcement authority, such as attempting to ground more progressive or novel actions in topics with thick social consensus. For another, it pushes policymakers seeking to empower agencies to consider institutional design; to account for the practical realities that an agency must confront, over time; and to think creatively about where there might be play in the joints.

Author

Alicia Solow-Niederman is an associate professor of law at The George Washington University Law School. Professor Solow-Niederman's scholarship sits at the intersection of law and technology. Her research focuses on how to regulate emerging technologies, such as artificial intelligence, in a way that reckons with social, economic, and political power. With an emphasis on algorithmic accountability, data governance, and information privacy, Professor Solow-Niederman explores how digital technologies can both challenge longstanding regulatory approaches and expose underlying legal values.

Professor Solow-Niederman's work has been published or is forthcoming in the Harvard Journal on Law & Technology, the Northwestern University Law Review, the Southern California Law Review, and the Berkeley Technology Law Journal, among other law reviews and peer-reviewed journals. Her piece on data breaches was selected as a winner of the 2017 Yale Law Journal Student Essay Competition. Professor Solow-Niederman is a member of the Electronic Privacy Information Center (EPIC) Advisory Board. She is also a faculty affiliate at Harvard University's Berkman Klein Center for Internet & Society and a visiting fellow at the Yale Law School Information Society Project, where she has worked with the Media Freedom and Information Access Clinic on a series of FOIA requests concerning state government use of AI.

Professor Solow-Niederman teaches or has taught courses in information privacy, technology and law, legislation and regulation, and torts. She clerked in U.S. District Court in the District of Columbia and served as a Climenko Fellow at Harvard Law School and a fellow on AI, law, and policy at UCLA School of Law. Professor Solow-Niederman received her J.D., cum laude, from Harvard Law School and a B.A. with distinction in communication and political science from Stanford University. In her free time, she enjoys distance running, crossword puzzles, and ice cream.

Navigating Demographic Measurement for Fairness and Equity

Miranda Bogen

Available at: https://cdt.org/insights/report-navigating-demographic-measurement-for-fairness-and-equity/

Executive Summary

Governments and policymakers increasingly expect practitioners developing and using AI systems in both consumer and public sector settings to proactively identify and address bias or discrimination that those AI systems may reflect or amplify. Central to this effort is the complex and sensitive task of obtaining demographic data to measure fairness and bias within and surrounding these systems. This report provides methodologies, guidance, and case studies for those undertaking fairness and equity assessments — from approaches that involve more direct access to data to ones that don't expand data collection. Practitioners are guided through the first phases of demographic measurement efforts, including determining the relevant lens of analysis, selecting what demographic characteristics to consider, and navigating how to hone in on relevant sub-communities. The report then explores a variety of approaches to uncover demographic patterns and responsibly handle demographic data. While there is no one-size-fits-all solution, the report makes clear that the lack of obvious access to raw demographic data should not be considered an insurmountable barrier to assessing AI systems for fairness, nor should it provide a blanket justification for widespread or incautious data collection efforts.

Author

Miranda Bogen is the founding director of the AI Governance Lab at the Center for Democracy & Technology. Building on CDT's decades of leadership fighting to advance civil rights and civil liberties in the digital age, the Lab provides public interest expertise in rapidly developing policy and technical conversations around artificial intelligence, advancing the interests of individuals whose lives and rights are impacted by AI.

An AI policy expert and responsible AI practitioner, Miranda has led work at the intersection of policy and AI fairness and governance in senior roles in industry and civil society. She served as co-chair of the Fairness, Transparency, and Accountability Working Group at the Partnership on AI, conducted foundational research at the intersection of machine learning and civil rights at Upturn, and most recently guided strategy and implementation of responsible AI practices at Meta. Bogen co-authored widely cited research on the potential for discrimination in personalized advertising and the role of artificial intelligence in the hiring process, and her work has informed international policy discussions on the civil and human rights implications of artificial intelligence, including citations in the White House Blueprint for an AI Bill of Rights. Bogen's writing and analysis has appeared in publications including the Harvard Business Review, NPR, Slate, and Newsweek, and her work has been featured in The Wall Street Journal, The Atlantic, The Economist, Reuters, Wired, MIT Technology Review, Last Week Tonight, and more.

Bogen holds a Masters from The Fletcher School of Law and Diplomacy at Tufts University with a focus on international technology policy, and graduated summa cum laude and Phi Beta Kappa from UCLA with degrees in Political Science and Middle Eastern & North African Studies.

Personhood Credentials: Artificial Intelligence and the Value of Privacy-Preserving Tools to Distinguish Who Is Real Online

Steven Adler, Zoë Hitzig, Shrey Jain, Catherine Brewer, Wayne Chang, Renée DiResta, Eddy Lazzarin, Sean McGregor, Wendy Seltzer, Divya Siddarth, Nouran Soliman, Tobin South, Connor Spelliscy, Manu Sporny, Varya Srivastava, John Bailey, Brian Christian, Andrew Critch, Ronnie Falcon, Heather Flanagan, Kim Hamilton Duffy, Eric Ho, Claire R. Leibowicz, Srikanth Nadhamuni, Alan Z. Rozenshtein, David Schnurr, Evan Shapiro, Lacey Strahm, Andrew Trask, Zoe Weinberg, Cedric Whitney, Tom Zick

Available at: https://arxiv.org/pdf/2408.07892

Executive Summary

Anonymity is an important principle online. However, malicious actors have long used misleading identities to conduct fraud, spread disinformation, and carry out other deceptive schemes. With the advent of increasingly capable AI, bad actors can amplify the potential scale and effectiveness of their operations, intensifying the challenge of balancing anonymity and trustworthiness online. In this paper, we analyze the value of a new tool to address this challenge: "personhood credentials" (PHCs), digital credentials that empower users to demonstrate that they are real people — not AIs — to online services, without disclosing any personal information. Such credentials can be issued by a range of trusted institutions — governments or otherwise. A PHC system, according to our definition, could be local or global, and does not need to be biometrics-based. Two trends in AI contribute to the urgency of the challenge: AI's increasing indistinguishability from people online (i.e., lifelike content and avatars, agentic activity), and AI's increasing scalability (i.e., cost-effectiveness, accessibility). Drawing on a long history of research into anonymous credentials and "proof-of-personhood" systems, personhood credentials give people a way to signal their trustworthiness on online platforms, and offer service providers new tools for reducing misuse by bad actors. By contrast, existing countermeasures to automated deception — such as CAPTCHAs — are inadequate against sophisticated AI, while stringent identity verification solutions are insufficiently private for many use-cases. After surveying the benefits of personhood credentials, we also examine deployment risks and design challenges. We conclude with actionable next steps for policymakers, technologists, and standards bodies to consider in consultation with the public.

Authors

Zoë Hitzig is a Junior Fellow at the Harvard Society of Fellows and a Research Scientist at OpenAI. Her research at the intersection of economics and computer science centers on privacy and transparency in markets and algorithms. She is the author of two books of poetry, and occasionally writes about economics and technology in venues like WIRED, Artforum, and The Drift.

Shrey Jain is a Product Manager at Microsoft, advancing AI for healthcare with a focus on multimodal imaging models. Previously an Applied Scientist at Microsoft Research, he tackled privacy-preserving AI, built cryptographic tools for disinformation resilience, and co-founded the Plural Technology Collaboratory within Microsoft Research Special Projects. Shrey studied Engineering Science at the University of Toronto, researched healthcare AI at MIT CSAIL, and founded Flatten — a public health nonprofit advised by Geoffrey Hinton — that supported the Canadian and Somali governments with COVID-19 rapid response efforts.

Honorable Mention

Aligning Algorithmic Risk Assessments with Criminal Justice Values

Dennis D. Hirsch, Angie Westover-Muñoz, Christopher B. Yaluma, and Jared Ott

Available at: https://moritzlaw.osu.edu/sites/default/files/2024-12/ARA%20ADA.pdf

Executive Summary

The use of risk assessment (RA) tools has become a key component of the criminal justice system in the United States. Much of the existing scholarship concentrates on normative and technical aspects of RAs, or on recommendations for their improvement. However, there has been little empirical work on how courts and other criminal justice actors perceive and utilize these tools on the ground. In this study, we provide an in-depth picture of how the Courts of Common Pleas think about and use algorithmic risk assessments. Primarily, we focus on the use of risk assessment tools in Ohio Courts of Common Pleas and compare Ohio practices with best practices highlighted in the literature. To investigate, we surveyed Ohio Courts of Common Pleas judges, probation officers, and court administrators regarding their views on and use of algorithmic risk assessment tools. We further conducted interviews with judges and a diverse array of stakeholders that included victim's rights, civil liberties, and civil rights groups, as well as public defenders and county prosecutors.

The findings show that judges largely see risk assessment tools as essential to their decision-making, with most trusting the tools to improve risk-related judgments. Our findings on Ohio's use of risk assessment tools are mixed. Judges agree the tools should guide, not dictate, decisions, aligning with best practices. However, many lack sufficient training — a crucial recommendation. We conclude with broad recommendations for enhancing the use of risk assessment tools in the judicial system.

Authors

Dennis D. Hirsch is a Professor of Law and of Computer Science at The Ohio State University (OSU) where he also serves as Director of the Program on Data and Governance and as a core faculty member of the Translational Data Analytics Institute (TDAI). His research and teaching, and the program that he directs, focus on the law, policy, ethics, and management of advanced analytics and AI. The author of numerous articles and an award-winning book, Professor Hirsch is the co-editor of the SSRN eJournal on Artificial Intelligence — Law, Policy, and Ethics, the founding Chair of TDAI's Responsible Data Science Community of Practice, the founder of the Ohio Data Ethics Working Group, a member of the Advisory Board for the International Association of Privacy Professionals' (IAPP) AI Governance Center, and a member of the OECD's Expert Group on AI Risk and Accountability. He received his J.D. from Yale Law School.

Angie Westover-Muñoz is a manager and researcher specializing in the governance of technology and advanced analytics. She currently serves as the Program Manager for the Program on Data and Governance at the Moritz College of Law, Ohio State University. In this role, she collaborates with practitioners to develop and implement programs for the responsible use of artificial intelligence and advanced analytics. She is also a Ph.D. candidate in Public Policy and Management at the Glenn College of Public Affairs, Ohio State University, with her dissertation focusing on the governance of data integrated systems in smart city initiatives. Her research interests include the adoption, use and governance of technology and data by local and regional governments, as well as private organizations, to support responsible decision-making and policy implementation.

Christopher B. Yaluma is a Senior Research Associate at The Ohio State University's Program on Data and Governance at Moritz College of Law. Chris designs and manages research projects on the governance and management of AI and advanced analytics within private and public organizations. He employs a multidisciplinary approach in his research, often utilizing quasi-experimental research designs. Chris obtained his Ph.D. in Public Policy and Management from John Glenn College of Public Affairs at The Ohio State University.

Jared Ott is an assistant professor of communication studies at Indiana University East. He researches mediated communication message effects and processes, and teaches courses on new media technology and digital responsibilities and rights, among others. He previously worked as a researcher for the Program on Data and Governance at Ohio State's Moritz College of Law, where he contributed to projects examining business' responsible AI management and the use of algorithmic risk assessments in the Ohio judicial system.

Honorable Mention

The Law for AI for Good

Orly Lobel

Florida Law Review, Vol. 75, 2023

Available at: https://www.floridalawreview.com/article/91298-the-law-of-ai-for-good

Executive Summary

Legal policy and scholarship are increasingly focused on regulating technology to safeguard against risks and harms, neglecting the ways in which the law should direct the use of new technology, and in particular artificial intelligence (AI), for positive purposes. This article pivots the debates about automation, finding that the focus on AI wrongs is descriptively inaccurate, undermining a balanced analysis of the benefits, potential, and risks involved in digital technology. Further, the focus on AI wrongs is normatively and prescriptively flawed, narrowing and distorting the law reforms currently dominating tech policy debates. The law-of-AI-wrongs focuses on reactive and defensive solutions to potential problems while obscuring the need to proactively direct and govern increasingly automated and datafied markets and societies. Analyzing a new Federal Trade Commission (FTC) report, the Biden administration's 2022 AI Bill of Rights and American and European legislative reform efforts, including the Algorithmic Accountability Act of 2022, the Data Privacy and Protection Act of 2022, the European General Data Protection Regulation (GDPR) and the new draft EU AI Act, the article finds that governments are developing regulatory strategies that almost exclusively address the risks of AI while paying short shrift to its benefits. The policy focus on risks of digital technology is pervaded by logical fallacies and faulty assumptions, failing to evaluate AI in comparison to human decision-making and the status quo. The article presents a shift from the prevailing absolutist approach to one of comparative cost-benefit. The role of public policy should be to oversee digital advancements, verify capabilities, and scale and build public trust in the most promising technologies.

Author

Orly Lobel is the Warren Distinguished Professor of Law and founding director of the Center for Employment and Labor Policy (CELP) at University of San Diego. She is the award-winning author of best-selling books and numerous high-impact articles. A graduate of Tel-Aviv University and Harvard Law School, Lobel clerked on the Israeli Supreme Court and is a member of the American Law Institute. She has recently been named the most cited legal scholar in employment law and overall one of the most cited younger legal scholars in the United States. She has received several grants for her scholarship including most recently a grant from the AI and Humanities Project. She is the winner of the 2023 Vanguard award from the California Bar Association.

Lobel served on President Obama's policy team on innovation and labor market competition, advised the Biden Administration's Federal Trade Commission (FTC) and other federal and state agencies on tech policy. Lobel consults private tech leaders on competition, human capital, equality, innovation, labor markets, and tech policy. In 2023, she keynoted the United Nation's AI for Good Summit in Geneva and the EU State of the Union in Florence. She also served this past year as a G7 representative of the World Economic Forum to Japan's governmental taskforce on digital transformation.

Her books You Don't Own Me: How Mattel v. MGA Entertainment Exposed Barbie's Dark Side (in development by CBS into a series) (Norton) and Talent Wants to Be Free: Why We Should Learn to Love Leaks, Raids and Free Riding (Yale University Press) are the recipient of several prestigious awards and have been reviewed in top scholarly journals and media. Her new book The Equality Machine: Harnessing Tomorrow's Technologies for a Brighter, More Inclusive Future (PublicAffairs) has received raving reviews and was named by The Economist Best Book of 2022 ("brilliant").

Awarded Student Paper

Data Subjects' Reactions to Exercising Their Right of Access

Arthur Borem, Elleen Pan, Olufunmilola Obielodan, Aurelie Roubinowitz, Luca Dovichi, Michelle L. Mazurek, Blase Ur

Available at: https://www.usenix.org/conference/usenixsecurity24/presentation/borem

Executive Summary

Recent privacy laws have strengthened data subjects' right to access personal data collected by companies. Prior work has found that data exports companies provide consumers in response to Data Subject Access Requests (DSARs) can be overwhelming and hard to understand. To identify directions for improving the user experience of data exports, we conducted an online study in which 33 participants explored their own data from Amazon, Facebook, Google, Spotify, or Uber. Participants articulated questions they hoped to answer using the exports. They also annotated parts of the export they found confusing, creepy, interesting, or surprising. While participants hoped to learn either about their own usage of the platform or how the company collects and uses their personal data, these questions were often left unanswered. Participants' annotations documented their excitement at finding data records that triggered nostalgia, but also shock and anger about the privacy implications of other data they saw. Having examining their data, many participants hoped to request the company erase some, but not all, of the data. We discuss opportunities for future transparency-enhancing tools and enhanced laws.

Author

Arthur Borem is a PhD student in Computer Science at the University of Chicago advised by Blase Ur. His work focuses on the privacy and security implications arising from the large-scale collection and use of personal data by online platforms. In particular, he's exploring and designing strategies and tools for implementing data subject rights guaranteed by the GDPR and other privacy regulations that more effectively promote user autonomy and platform transparency. Arthur has a BS in Computer Science from Brown University and before starting his PhD he was a software engineer at Asana.

Student Paper Honorable Mention

Artificial Intelligence is Like a Perpetual Stew

Nathan Reitinger

American University Law Review, Vol. 73, 2024

Available at: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4685772

Executive Summary

Governments around the world — in response to the abusive, dangerous, and unchecked powers that AI wields (e.g., the ability to identify anyone on earth with a single photograph, unrecognizable misinformation, and hidden biases) — are beginning to consider the right ways to safeguard AI. Unfortunately, any safeguards that may be enacted must come face to face with this simple question: What is AI? This question turns out to be of paramount importance to policymakers on a variety of levels. How should we ensure that criminal or tortious AI decisions do not create liability chasms if AIs act without mens rea; do AI models memorize — and therefore violate copyright law — when they are trained; how should we flag and remove biased model output when those outputs have real consequences for real people; how should we protect ourselves from the spread of misinformation when it is not possible to tell the difference between reality and deepfakes? Each of these questions hinges on the way AI operates under the hood: Knowing how AI works is vital to regulating AI. With that in mind, this paper attempts to teach AI in a way that allows non-experts to grasp AI's fundamentals and apply that knowledge toward crafting proper guardrails. The piece does this with simple, yet accurate, examples, and provides an analogy of how to think about machine learning models. A machine learning model is like a perpetual stew. Like the perpetual stew, these models are built with recipes, the recipes are tweaked per the model's specific purposes (i.e., adding a pinch of salt), and then the models "live" (i.e., produce accurate or tasty outputs) for as long as they are maintained. This simple analogy allows readers to understand why, for instance, it is not generally possible to remove a piece of sensitive data from a model — just like it is difficult to remove a pinch of salt from a stew. In turn, the piece is crucial for any policymaker considering how to regulate the promises and perils of AI.

Author

Nathan Reitinger is a computer scientist and legal scholar with over 13 law and computer science publications. His publications land in top technology-law journals, like the Stanford Technology Law Review, Jurimetrics, and the American University Law Review, and top computer science conferences, like the IEEE Symposium on Security and Privacy, USENIX, and the Privacy Enhancing Technologies Symposium. He has deep expertise in the fields of privacy and artificial intelligence, using technical methodologies to analyze legal questions of theory and doctrine.

Thank you to our 2024 Reviewers and Finalist Judges

Submissions received numeric rankings from a diverse team of academics, consumer advocates, and industry privacy professionals from the FPF Advisory Board, with each submission being evaluated for originality; overall quality of writing; and applicability to policymaking. For more information, visit fpf.org/privacy-papers-for-policy-makers.

Advisory Board Reviewers

Lael Bellamy
Partner, DLA Piper

Debra Berlyn
Project GOAL

Sara Collins
Public Knowledge

Rachel Cummings
Columbia University

Jo Davaris
Booking Holdings

Michael Dolan
Best Buy

Stacy Feuer
Senior Vice President,
ESRB Privacy Certified

Elise Houlik
Chief Privacy Officer, Intuit

Paul Lekas
Software & Information
Industry Association

Cathy Mulrow-Peattie

Kenneth Propp
Georgetown University
Law Center & Atlantic Council
Europe Center

Claire Readhead
PwC

Lauren Smith
Cruise

Rachel Thompson
SVP Assistant General Counsel,
Privacy & Data Protection,
Mastercard

Alexander White
Privacy Commissioner
for Bermuda

Ron Whitworth
Truist

Cobun Zweifel-Keegan
IAPP

Finalist Judges

Daniel Hales
Policy Fellow for Youth Privacy, Future of Privacy Forum

Jules Polonetsky
CEO, Future of Privacy Forum

John Verdi
Senior Vice President for Policy, Future of Privacy Forum

Adonne Washington
Policy Counsel for Data, Mobility, Location, Future of Privacy Forum

PRIVACY PAPERS FOR POLICYMAKERS

2009–2024

Future of Privacy Forum (FPF) is a nonprofit organization that serves as a catalyst for privacy leadership and scholarship, advancing principled data practices in support of emerging technologies.

1350 EYE STREET NW | SUITE 350 | WASHINGTON, DC 20005 FPF.ORG | 202.220.8062