


RESEARCH ARTICLE

Expanding the paradigm: Generative artificial intelligence and U.S. privacy norms

Elana Zeide 

College of Law, University of Nebraska College of Law, Lincoln, NE, USA
Email: zeide@unl.edu

Affiliated Fellow at: The Information Society Project, Yale Law School, Silicon Flatirons, University of Colorado–Boulder, The Cordell Institute for Policy in Medicine & Law, Washington University in St. Louis.

(Received 5 August 2024; revised 21 November 2024; accepted 2 December 2024)

Abstract

Generative artificial intelligence (AI) systems, such as large language models, image synthesis tools, and audio generation engines, present remarkable possibilities for creative expression and scientific discovery but also pose pressing challenges for privacy governance. By identifying patterns in vast troves of digital data, these systems can generate hyper-realistic yet fabricated content, surface sensitive inferences about individuals and groups, and shape public discourse at an unprecedented scale. These innovations amplify privacy concerns about nonconsensual data extraction, re-identification, inferential profiling, synthetic media manipulation, algorithmic bias, and quantification. This article argues that the current U.S. legal framework, rooted in a narrowly targeted sectoral approach and overreliance on individual notice and consent, is fundamentally mismatched to address the emergent and systemic privacy harms of generative AI. It examines how the unprecedented scale, speed, and sophistication of these systems strain core assumptions of data protection law, highlighting the misalignment between AI's societal impacts and individualistic, reactive approaches to privacy governance. The article explores distinctive privacy challenges posed by generative AI, surveys gaps in existing U.S. regulations, and outlines key elements of a new paradigm to protect individual and collective privacy rights that (1) shifts from individual to collective conceptions of privacy; (2) moves from reactive to proactive governance; and (3) reorients the goals and values of AI governance. Despite significant obstacles, it identifies potential policy levers, technical safeguards, and conceptual tools to inform a more proactive and equitable approach to governing generative AI.

Keywords: generative artificial intelligence; privacy; law; algorithmic bias; synthetic media; U.S. regulation

1. Introduction

The emergence of generative artificial intelligence (AI), a suite of technologies capable of creating novel and realistic content, represents a transformative development in AI (Surden, 2024, pp. 1944–1948). Generative AI refers to a class of machine learning models and techniques that can create new content – such as text, images, audio, video, and code – based on patterns learned from vast troves of training data (Surden, 2024). These systems, including natural language models like ChatGPT, image generation tools like DALL-E and Stable Diffusion, and audio synthesis engines like WaveNet, have already begun to revolutionize how we create, interact with, and perceive digital content (Davenport & Mittal, 2022; Kreaic et al., 2024).

© The Author(s), 2025. Published by Cambridge University Press. This is an Open Access article, distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives licence (<http://creativecommons.org/licenses/by-nc-nd/4.0>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided that no alterations are made and the original article is properly cited. The written permission of Cambridge University Press must be obtained prior to any commercial use and/or adaptation of the article.

While earlier AI excelled at tasks like image classification, speech recognition, and fraud detection, these systems were primarily constrained to identifying and exploiting correlations within narrowly defined problem domains (Surden, 2024). In contrast, generative AI systems produce output that is often indistinguishable from human-created content (Cooper et al., 2023, p. 7; Heikkilä, 2024). While these tools offer exciting new possibilities for creative expression, scientific discovery, and social innovation, they also raise profound privacy concerns (Bender et al., 2021; Cooper et al., 2023, p. 10; Helmus & Chandra, 2024; King & Meinhardt, 2024; Luccioni & Viviano, 2021; Matsumi & Solove, 2024; Song & Shmatikov, 2019; Weidinger et al., 2021; Weidinger et al., 2022).

This article explores the privacy challenges posed by generative AI and argues for a fundamental rethinking of privacy governance frameworks in response. Section 2 examines the technical characteristics and capabilities of generative AI systems that amplify existing privacy risks and introduce new challenges, including nonconsensual data extraction, data leakage and re-identification, inferential profiling, synthetic media generation, algorithmic bias, and quantification. Section 3 surveys the current landscape of U.S. privacy law and its shortcomings in addressing these emergent issues, highlighting the limitations of the sectoral approach, the FTC's constrained authority, the promise and pitfalls of state laws, and the inadequacy of individualistic privacy paradigms. Section 4 outlines critical elements of an alternative paradigm for generative AI privacy governance that: (1) shifts from individual to collective conceptions of privacy; (2) moves from reactive to proactive governance; and (3) reorients the goals and values of AI governance. The analysis concludes by discussing the political, legal, and cultural obstacles to regulatory reform in the United States while emphasizing the urgent need for action given the high stakes for individual autonomy and democratic values.

2. How generative AI challenges privacy

The rapid advancement of generative AI systems, with their enhanced ability to create highly realistic and persuasive content, magnifies existing privacy risks while also introducing new challenges that test the foundational assumptions of current privacy frameworks. The technical characteristics of generative AI systems exacerbate existing privacy threats. These risks include large-scale extraction of public data without individual consent and control; data leakage and re-identification; inferential privacy harms; generation of fake but convincing synthetic media; exacerbation of algorithmic bias and discrimination; and decontextualized quantification (Bommasani et al., 2022; Cooper et al., 2023, p. 10; Helmus & Chandra, 2024; King & Meinhardt, 2024; Lee et al., 2024; Shelby et al., 2023; Solove, 2024, 2025; Song & Shmatikov, 2019; Weidinger et al., 2021, 2022; Zeide, 2017). While society has long grappled with privacy concerns around big data and machine learning, the power, sophistication, and inscrutability of generative AI exceed the scope of current data governance paradigms, revealing weaknesses in established legal and ethical frameworks for protecting privacy and autonomy.

2.1 Nonconsensual data extraction and the failure of notice and consent

Generative AI systems are trained on vast amounts of data, often comprising billions of individual data points spanning a wide range of formats, domains, and sources (Baio, 2022; Schuhmann et al., 2022). For example, OpenAI trained its GPT-4 language model on a corpus of over 45 terabytes of text data, including books, articles, and websites (Achiam et al., 2024). While some of these data are clearly public material like stock photos, much of it includes individuals' names, addresses, and images published online under varying expectations of privacy (Cooper et al., 2023; King & Meinhardt, 2024).

This harvesting and processing of personal information and sensitive content occurs without notice, consent, or constraint (King & Meinhardt, 2024, pp. 17–19; Leffer, n.d.; Morrison, 2023; Solove, 2025, pp. 23–29). Current law explicitly allows (*HiQ Labs v. LinkedIn Corp.*, 2019) or implicitly

sanctions the collection and use of publicly available information (California Consumer Privacy Act (CCPA), 2023); Va. Code § 59.1-571, 2021; Utah Code Ann. § 13-61-101(29)(b), 2024). However, the unprecedented scope and granularity of data extraction by generative AI systems erode the assumptions of individual autonomy that underlie existing privacy frameworks (King & Meinhardt, 2024, pp. 17–19; Solove, 2025, pp. 23–29). In many cases, individuals are unaware that their data are being used to train these systems and cannot manage the downstream uses of their data. For example, Clearview AI, a facial recognition company, scraped billions of images from social media platforms to train its algorithms without the knowledge or consent of the individuals depicted (Hill, 2020; Tangalakis-Lippert, 2023).

2.2 Data leakage and re-identification

The training data used by generative AI systems can also be vulnerable to data leakage and re-identification attacks (Carlini et al., 2019; King & Meinhardt, 2024; Leffer, n.d.; Morrison, 2023; Solove, 2025; Staab et al., 2023; Weidinger et al., 2022; Winograd, 2023). Because these models capture patterns at a high level of granularity, they can inadvertently “memorize” and reproduce sensitive snippets of input data in synthetic outputs (Carlini et al., 2019; King & Meinhardt, 2024; Leffer, n.d.; Staab et al., 2023; Winograd, 2023). For example, a language model trained on a corpus of emails might reveal real names, addresses, or phone numbers (Carlini et al., 2021). Image synthesis models trained on photos from social media can produce pictures that depict recognizable individuals or locations (Fernandez et al., 2023). Moreover, malicious actors can craft adversarial prompts to extract specific sensitive information (Edwards, 2022).

2.3 Inferential profiling and privacy harms

Generative AI systems exploit subtle patterns and correlations in large datasets to make probabilistic inferences about a person’s demographics, preferences, behaviors, and beliefs, even when such information is not explicitly disclosed (Cooper et al., 2023; Gillis, 2022; King & Meinhardt, 2024; Solove, 2025; Weidinger et al., 2022, p. 218). Language models trained on social media posts might learn to associate certain linguistic styles, topics, or sentiments with particular demographic groups, allowing them to make inferences about a user’s age, ethnicity, or socioeconomic status based on their writing patterns (Solow-Niederman, 2022; Zeide, 2015, 2022). Similarly, a computer vision model trained on user-uploaded images might be able – or at least claim to be able – to infer sensitive attributes like health conditions, political affiliations, or sexual orientation based on visual cues and contextual signals in the images (Wang & Kosinski, 2018).

2.4 Synthetic media, deepfakes, and disinformation

Generative AI’s ability to create highly realistic content opens the door to pervasive deception and manipulation (King & Meinhardt, 2024; Solove, 2025; Weidinger et al., 2022). Malicious actors can exploit readily available tools to produce “deepfakes” and other types of synthetic media that convincingly impersonate real people and mislead audiences on a massive scale (Salam, 2023). Deep learning models, for instance, can clone an individual’s voice from just a few seconds of audio and generate fake audio clips (Leffer, 2024). Natural language models can mass-produce fake news articles, product reviews, and social media posts that are nearly impossible to distinguish from authentic content. Image synthesis systems can create realistic faces of nonexistent individuals or seamlessly insert real people’s faces into fabricated scenarios (Westerlund, 2019).

These technologies are now accessible even to those with limited technical expertise or resources, who can leverage them for deceptive or harmful purposes (Heikkilä, 2024). Middle school children across the country, for example, alter images of their classmates to create “deepfake porn” (Rubin, 2023; Verma, 2023). As synthetic media capabilities grow increasingly sophisticated and accessible,

it is becoming increasingly difficult to distinguish between real and generated content, eroding the epistemic foundations of trust and truth in online interactions (Baker & Chadwick, 2021; Chesney & Citron, 2019; Kalpokas, 2019; MacKenzie & Bhatt, 2020; Shin, 2024; Tokaji, 2019; Whyte, 2020).

2.5 Algorithmic bias and discrimination

Generative AI systems also risk perpetuating and amplifying historical patterns of bias and discrimination reflected in their training data. For example, facial recognition algorithms exhibit higher error rates for women and people of color (Buolamwini & Gebru, 2018; Grother, 2022; Grother et al., 2019). Researchers have shown that language models like GPT-3 can exhibit gender and racial biases in their generated text, such as associating men with career-oriented terms and women with family-oriented terms or perpetuating harmful stereotypes about minority groups (Bolukbasi et al., 2016; Caliskan-Islam et al., 2016).

These biases can translate to real-world harms when entities use generative systems to allocate benefits and opportunities (Ajunwa, 2021, 2020, p. 1405; Geddes, 2023, p. 31; Kim, 2016; O'Neil, 2016; Solove, 2025, pp. 45–46; Solove & Matsumi, 2024; Zeide, 2022). A company that uses a biased language model to screen resumes or generate job descriptions may end up excluding qualified candidates from underrepresented backgrounds. A government agency that employs a skewed facial recognition tool to identify suspects or predict recidivism risk may disproportionately target and surveil communities of color (Barrett, 2017; Meijer & Wessels, 2019). Over time, such discriminatory outcomes can compound disadvantage and erode economic mobility for marginalized communities (Zeide, 2022).

2.6 Quantification and decontextualization

Automated profiling and decision-making by generative AI systems can also lead to the decontextualization and abstraction of individuals, reducing them to a set of quantifiable data points and statistical inferences (Citron & Pasquale, 2014; Cohen, 2000, p. 1405; Zeide, 2017, p. 169). This reductive approach to human identity and agency fails to capture the complexity and nuance of individual circumstances, leading to decisions that may be inaccurate, unfair, or devoid of situational understanding (Cohen, 2000, p. 1405; Geddes, 2023, p. 31; O'Neil, 2016; Solove & Matsumi, 2024, pp. 45–46; Zeide, 2017). A generative AI system used to predict someone's creditworthiness or risk of recidivism may rely on aggregate patterns and correlations learned from historical data that do not reflect the full context of that person's circumstances and capacity for behavioral change (Eaglin, 2017). This risks creating a system of self-fulfilling prophecies that undermine individuals' autonomy and agency (Cohen, 2013; Harcourt, 2008; Kerr & Earle, 2013; Lazaro, 2018; Solove, 2024; Véliz, 2021; Zeide, 2017, 2022). Moreover, using generative AI systems to automate high-stakes decisions shifts discretion away from domain experts to unaccountable actors (Engstrom & Haim, 2023, pp. 291–292; Zeide, 2017, pp. 168–169). These systems optimize based on what they measure, thereby shaping not only individual assessments but also determining the broader goals and values of a given context (Engstrom & Haim, 2023, pp. 291–292; Zeide, 2017, pp. 168–169). By displacing situated judgment with opaque and unaccountable systems, generative AI risks enabling private entities to shape public values and societal norms without adequate transparency or oversight (Engstrom & Haim, 2023, pp. 291–292; Zeide, 2017, pp. 168–169). In summary, the advanced capabilities of generative AI systems pose significant threats to privacy at both the individual and societal levels.

3. The inadequacy of U.S. privacy law in addressing generative AI challenges

The formidable capabilities and evolving risks of generative AI systems present substantial challenges to current privacy and data protection frameworks in the United States. This section highlights three key limitations of the current legal framework: (1) the fragmented and incomplete patchwork of

federal and state laws; (2) the mismatch between generative AI's collective harms and a framework premised on notice and choice; and (3) the inadequacy of individualistic privacy models in capturing AI's systemic impacts. These shortcomings necessitate a fundamental rethinking of privacy governance in the age of AI.

3.1. A fragmented and sectoral approach to privacy regulation

Unlike the European Union, which has comprehensive privacy and data protection regimes like the General Data Protection Regulation (GDPR) (Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation, 2016), Arts 12–22 (Rights of the data subject) and the Artificial Intelligence Act (Proposal for a Regulation of the European Parliament and of the Council Laying Down Harmonised Rules on Artificial Intelligence (Artificial Intelligence Act)), the United States employs a sectoral approach to privacy and data regulation. This fragmented approach protects specific data categories rather than establishing a general right to privacy across all contexts.

The current framework comprises a patchwork of federal and state laws imposing different obligations on collecting, using, and sharing personal data based on industry, data type, and jurisdiction. Key examples include the Health Insurance Portability and Accountability Act (42 U.S.C. § 1420d et seq.), the Gramm–Leach–Bliley Act (15 U.S.C. § 6821 et seq.), and the Children's Online Privacy Protection Act (15 U.S.C. 6501–6506), each imposing unique obligations on the handling of personal healthcare, financial, and children's data, respectively. The fragmented U.S. framework, which focuses narrowly on protecting specific categories of data based on context, struggles to regulate AI systems that can repurpose and recombine information in ways that transcend traditional sectoral boundaries (King & Meinhardt, 2024; Solove, 2024; Weidinger et al., 2022).

3.2. FTC authority and limitations

In the absence of comprehensive federal AI legislation, the Federal Trade Commission (FTC) has emerged as the de facto federal authority responsible for privacy protection in the context of AI systems, including generative models (DiResta & Sherman, 2023). The FTC's authority stems from its general mandate to protect consumers from unfair or deceptive practices under section 5 of the FTC Act. This includes taking enforcement actions against companies that violate their own privacy policies or engage in "unfair" practices that cause or are likely to cause substantial injury to consumers, that cannot be reasonably avoided by consumers, and that are not outweighed by countervailing benefits to consumers or competition (Federal Trade Commission Act § 5, 15 U.S.C. § 45(a), 2018).

In recent years, the Commission has taken a number of actions to address the privacy and fairness implications of AI systems, including issuing guidance (Federal Trade Commission, 2021, 2023a), conducting workshops (Federal Trade Commission, 2023b), and penalizing companies deploying AI without adequate safeguards against discriminatory impact on protected classes (*Federal Trade Commission v. Rite Aid Corp.*, 2024; Hanley & Goldfarb, 2021). For example, in 2016, the FTC reached a settlement with InMobi over its deceptive use of location tracking in its mobile ad targeting system (*FTC v. InMobi*, 2016). In 2019, the agency issued warnings to companies that purport to use AI for automated hiring decisions about the risks of perpetuating or exacerbating discriminatory biases (FTC, 2020). In 2021, the FTC issued guidance explicitly cautioning that bias in AI could lead to enforcement actions under laws prohibiting unfair or deceptive practices (FTC, 2021). That same year, the FTC required EverAlbum to delete biometric and facial recognition data that had been collected without adequate notice and consent (*FTC v. EverAlbum*, 2021).

However, several factors limit the FTC's ability to effectively oversee generative AI systems under its current section 5 authority. First, the Commission lacks the substantive rulemaking power to

issue binding regulations interpreting what constitutes an unfair or deceptive AI practice (Hartzog & Solove, 2014), unlike the supervisory powers of the European Data Protection Board or the UK Information Commissioner's Office. There is also no statutory authority to conduct general audits or inspections of AI developers' practices (Hartzog & Solove, 2014). This results in largely reactive, fact-specific enforcement actions focused on cases of procedural violations, such as deceptive marketing or inadequate disclosures, rather than addressing the broader societal risks and harms posed by AI systems (Hartzog & Solove, 2014; Hirsch, 2020; Waldman, 2019).

Second, the FTC's jurisdiction only extends to commercial practices that cause or are likely to cause substantial injury to consumers, which may not cover some of the more intangible and externalized impacts of generative AI, such as the erosion of public trust or the amplification of disinformation (Calo, 2021; Lamo & Calo, 2019).

Third, recent judicial decisions have further curtailed the FTC's enforcement tools. In *AMG Capital Management v. FTC* (2021), the Supreme Court held that section 13(b) of the FTC Act does not authorize the agency to seek monetary relief like restitution or disgorgement, removing a key deterrent against privacy abuses (pp. 1344–1347). This significantly narrow the agency's enforcement powers Slaughter (Federal Trade Commission, 2021).

3.3. *The promise and pitfalls of state privacy laws*

Given these limitations at the federal level, a growing number of states have taken up the mantle of regulating AI systems (2024 AI State Law Tracker, 2024). As of January 2025, 20 states have enacted comprehensive consumer data protection laws that impose heightened requirements for processing sensitive personal information and establish rights of access, correction, and deletion (2024 AI State Law Tracker, 2024). These state privacy laws often incorporate data protection mechanisms from the GDPR (Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation) (Text with EEA relevance, 2016, Arts 12–22; Chander et al., 2020).

In addition to these comprehensive privacy laws, states have also enacted or proposed more targeted laws governing specific AI applications. For example, Maryland's Algorithmic Decision Systems Risk Assessment Act mandates bias and privacy impact assessments for government contractors using AI tools, while California's proposed Automated Decision Systems Accountability Act would require businesses to evaluate high-risk AI systems and report on their data practices. Some states have specifically targeted generative AI and deepfakes, with California's A.B. 1280 (2023) prohibiting the dissemination of synthetic media in political campaigns with the intent to deceive voters, and Texas' S.B. 2382 (2023) creating a civil cause of action for individuals whose likeness is used in sexually explicit deepfakes without their consent. Other states, such as Tennessee and Utah, have gone further by enacting laws that prohibit or criminalize the production of deepfakes more broadly (Tenn. Code Ann. § 47-25-1101 et seq., 2023); Utah Code Ann. § 76-5b-206 et seq., 2024).

While these state efforts represent an important source of policy innovation and experimentation in AI governance, they also suffer from several key limitations in addressing the unique challenges posed by generative AI. They remain fragmented and reactive, often addressing narrow harms rather than the underlying structural conditions that enable them (King & Meinhardt, 2024; Solove, 2024). Most still rely heavily on a notice-and-choice model of privacy protection that is ill-equipped to govern the complex data ecosystems of machine learning pipelines (Hartzog & Richards, 2020, p. 1704; King & Meinhardt, 2024; Solove, 2022, pp. 983–984, 2024). Moreover, their enforcement mechanisms are often limited to attorney general actions or narrow private rights of action, rather than more comprehensive administrative oversight and auditing (Fitzgerald et al., 2024). Ultimately, while state AI laws play a vital role in experimenting with different

regulatory approaches and filling gaps in federal policy, they are not sufficient on their own to govern the far-reaching impacts of generative AI, which demands a more comprehensive, proactive, and cohesive governance framework that can address both individual and collective privacy harms, align innovation with public values, and hold AI developers and deployers accountable across contexts.

3.4. *The limitations of individualistic privacy paradigms*

The biggest conceptual limitation of U.S. privacy law in the context of generative AI is its over-reliance on individual notice and consent as the primary mechanism for protecting personal autonomy (King & Meinhardt, 2024; Solove, 2022, 2025). This approach is ill-suited to address the scale, complexity, and opacity of data flows powering generative AI systems, which rely on web crawling, data brokers, and other indirect sources of data collection that are not visible and thus not amenable to granular individual control (Burrell, 2016; Hartzog & Richards, 2020; King & Meinhardt, 2024; Pasquale, 2015; Solove, 2013, 2024, 2025; Veale & Zuiderveen Borgesius, 2021).

First, privacy and data protection laws, which rely heavily on individual rights and procedural safeguards, struggle to contend with the scale, complexity, and opacity of data flows in generative AI systems (Edwards & Veale, 2017; Hartzog & Richards, 2020, p. 1704; Kaminski, 2023; Solove, 2022, p. 993). These characteristics render many core provisions, such as access rights, correction mechanisms, and deletion requirements, technically and practically unfeasible (Carlini et al., 2021; Katell et al., 2020; Shokri, Stronati, Song & Shmatikov, 2017; Villaronga et al., 2018; Waldman, 2019). Even when individuals can exercise their rights, doing so may not be technically feasible, as generative AI models involve a compressed representation of their training data, making it difficult to erase or remove personal information (Carlini et al., 2021). Similarly, if a generative model is used to create and disseminate harmful synthetic media, it will be difficult to contain or reverse the viral spread of this content (Bloch-Wehba, 2020; Chesney & Citron, 2019; Lamo & Calo, 2019; Van der Sloot & Wagenveld, 2022).

Furthermore, the individualistic focus of U.S. privacy law extends beyond notice and choice to its emphasis on procedural rights and ex post remedies through private lawsuits (Edwards & Veale, 2017; Hartzog & Richards, 2020, p. 1704; Hirsch, 2020, p. 462; Solove, 2022, p. 993; 2025). However, these mechanisms are often inadequate or ineffective in the face of generative AI's structural risks and harms, as the opacity and inscrutability of these systems pose significant challenges for existing legal frameworks designed to ensure transparency and accountability (Cohen, 2019b; Selbst & Barocas, 2018; Wachter & Mittelstadt, 2019).

Finally, intellectual property laws, particularly trade secret protections, enable companies to assert proprietary control over training data and algorithms (Tschider, 2021, p. 711; Wexler, 2018, p. 1402), limiting external visibility and oversight. This legal barrier compounds the technical inscrutability of generative AI systems, further undermining accountability (Burrell, 2016; Citron, 2008; Kroll et al., 2017; Selbst & Barocas, 2018). As a result, there is often little public disclosure of information about data sources, model architectures, training procedures, and output generation processes, making it difficult for individuals to understand how their data are being used and to identify potential harms or abuses.

3.5. *The collective harms and societal risks of generative AI*

Beyond these individual challenges, generative AI systems pose a range of diffuse societal risks that are mismatched with the individualistic, reactive focus of existing U.S. privacy laws (Bhargava & Velasquez, 2020; Cinnamon, 2017; Cooper et al., 2023; Zeide, 2022). Anti-discrimination laws, which target discrete instances of intentional or disparate impact discrimination, struggle to address the

structural and emergent harms of generative AI, such as compounded disadvantage, intersectional bias, and the preemptive shaping of opportunities (Kerr & Earle, 2013b; Solove, 2024; Zeide, 2022). These laws fail to capture the systemic and diffuse impacts of generative AI on historically disadvantaged populations (Mayson, 2018). Many of the most troubling impacts of generative AI are invisible, embedded in automated systems, and occurring before formal decision points (Zeide, 2022). Instead of explicit rejection, biased or inaccurate assessments and predictions often preempt access to opportunity (Kerr & Earle, 2013b; Solove, 2024; Zeide, 2022). For example, as I have discussed in prior work, predictive hiring algorithms can filter out qualified job applicants based on biased assessments, creating a “silicon ceiling” that imperceptibly impedes economic mobility for marginalized communities (Zeide, 2022).

In light of these limitations, reactive, individual-centric regulation is inadequate to mitigate the structural risks of generative AI. Addressing these challenges requires a proactive, systemic, and collaborative approach that moves beyond individual rights and remedies.

4. Towards a paradigm for generative AI governance

The limitations of existing regulatory frameworks in addressing the privacy risks posed by generative AI systems underscore the need for a new paradigm of privacy governance. This section argues for a fundamental reorientation of privacy protection from a narrow focus on individual control and procedural safeguards to a more systemic approach. It outlines three key elements of this new paradigm: (1) shifting from individual to collective conceptions of privacy; (2) moving from reactive to proactive governance; and (3) reorienting the goals and values of AI governance. The section concludes by acknowledging the significant obstacles to implementing such a paradigm shift in the United States, including the lack of a comprehensive federal privacy law, the limitations of sectoral and state-level regulations, and the entrenched ideological resistance to precautionary governance of emerging technologies.

4.1. *Shifting from individual to collective conceptions of privacy*

Given the limitations of individual control and the societal impact of generative AI, privacy governance should shift toward a more holistic understanding of privacy as a social value and public good. A robust approach to generative AI privacy governance will require reorienting the foundations of privacy protection from individual control and procedural rights to recognizing privacy as a social foundation and collective good (Cohen, 2019b; Tisné, 2020). As demonstrated by the limitations of existing regulations, policymakers cannot mitigate the privacy risks of generative AI solely by empowering individuals to control how their personal data are collected and used by particular entities. Instead, achieving meaningful privacy protection in the generative AI era requires recognizing the collective and relational dimensions of privacy harms (Milner & Traub, 2021; Viljoen, 2021).

4.2. *Moving from reactive to proactive governance*

A second key element in a more robust privacy governance framework is a shift from reactive and retroactive enforcement actions to proactive and preventative oversight regimes (Kaminski, 2023). Rather than relying primarily on ex post remedies triggered by specific legal violations or consumer complaints, policymakers should institutionalize continuous monitoring, auditing, and impact assessment requirements to surface and mitigate potential risks before companies and organizations deploy generative AI systems at scale (Kaminski, 2023; Katell et al., 2020; Metcalf et al., 2021). This means prioritizing proactive, preventative, and participatory approaches to AI governance that can anticipate and mitigate risks before they cause harm. It requires investing in new institutional capacities and governance frameworks that can enable ongoing monitoring, assessment, and public engagement throughout the AI lifecycle.

The European Union's AI Act offers a potential model for systematic AI governance (European Commission, 2023). One key strength of the AI Act is its focus on the broader societal impacts of AI systems, rather than just individual privacy harms (European Commission, 2023). The Act creates a comprehensive regulatory framework for AI systems, imposing graduated requirements based on a technology's level of risk (European Commission, 2023). Notably, the Act would require "high-risk" AI systems to undergo mandatory conformity assessments to ensure compliance with essential requirements related to data quality, transparency, human oversight, and robustness before entering the EU market. It also creates ongoing monitoring obligations for high-risk systems and establishes a centralized database for registering stand-alone AI systems. While the Act has drawn criticism for the compliance burdens it would impose (Corbett, 2024), it represents a meaningful effort to extend regulatory scrutiny to the entire AI lifecycle and to create an institutional infrastructure for proactive and adaptive governance (Veale & Zuiderveen Borgesius, 2021).

4.3. Reorienting the goals and values of AI governance

Ultimately, the limitations of the current U.S. privacy framework in addressing the challenges of generative AI point to the need for a more fundamental reorientation of the goals and values underlying AI governance. Addressing these challenges requires not just new regulatory tools and oversight mechanisms, but a deeper shift in how we conceptualize the purposes and priorities of AI governance itself (Milner & Traub, 2021; Powles & Nissenbaum, 2018; Viljoen, 2021). This reorientation involves moving beyond a narrow focus on protecting individual privacy rights and towards a broader vision of promoting collective well-being, social justice, and democratic values in the development and deployment of AI systems (Crawford, 2021; West et al., 2019; Whittaker et al., 2019). It means reconceptualizing privacy not just as a matter of individual control over personal data, but as a collective good that is essential for human autonomy, dignity, and self-determination in the face of increasingly powerful and pervasive AI systems (Cohen, 2019b; Tisné, 2020).

Finally, it means grappling with the inherently political and value-laden nature of AI development and governance, and creating mechanisms for democratic deliberation and contestation over the goals, values, and trade-offs embedded in these systems (Benthall & Haynes, 2019). This requires moving beyond technocratic and instrumental approaches to AI ethics and governance, and towards more inclusive and participatory processes that empower affected communities to shape the trajectories of AI innovation in line with their values and interests (Crawford, 2021). While the specific regulatory tools and accountability mechanisms needed to operationalize these principles will likely vary across different contexts and jurisdictions, reorienting the underlying goals and values of AI governance is an essential first step toward a more proactive, equitable, and democratically legitimate approach to managing the risks and benefits of generative AI systems (Milner & Traub, 2021; Viljoen, 2021).

4.4. Obstacles to comprehensive AI privacy governance in the USA

Implementing a new paradigm of AI privacy governance in the USA will not be easy, as it must contend with the country's deeply rooted legal traditions, political economy, and ideological commitments. One major hurdle is the political challenge of passing a comprehensive federal privacy law that could provide a coherent and consistent framework for governing AI systems across different sectors and jurisdictions. While there have been several proposals for such a law in recent years, none have yet been enacted, partly due to disagreements over preemption, private rights of action, and the scope of covered data and entities (Kerry et al., 2020). The highly polarized and industry-captured

policymaking process can block or dilute even incremental reforms, making it difficult to achieve the kind of systemic change needed to address generative AI's privacy risks (Kaminski, 2023).

Another significant obstacle is the strong protection afforded to freedom of expression under the First Amendment, which courts interpret to cover a wide range of data-driven activities, from collecting and disseminating publicly available information to creating and sharing synthetic media (Franks, 2019; Wu, 2012). These free speech protections, along with intellectual property rights, can hinder the ability of regulators to impose substantive restrictions on AI-generated content or mandate disclosure of proprietary AI systems (Franks, 2019; Massaro et al., 2017, pp. 2481–2525). This constitutional constraint, coupled with the U.S. policy landscape's historical favor for a *laissez-faire* and innovation-friendly approach to technological development, creates a challenging environment for proactive AI governance (Cohen, 2019a; Thierer, 2016, pp. 33–38).

The permissive stance of U.S. law, favoring market-driven solutions and self-regulation over precautionary regulation, is exemplified by numerous safe harbors and immunities for online platforms and technology providers, most notably section 230 of the Communications Decency Act, which shields platforms from liability for user-generated content (Kaminski, 2023; Citron & Wittes, 2017; Kosseff, 2019). This preference for innovation over precaution (Thierer, 2016, pp. 33–38), reflected in permissive liability regimes and judicial doctrines, places the burden of proof on regulators to demonstrate clear and concrete harm before intervening in the development and deployment of AI systems (Calo & Citron, 2020; Thierer, 2016). As a result, proactive regulation and public participation in AI governance can be chilled, making it difficult to address generative AI's privacy risks in a comprehensive and timely manner (Buchanan et al., 2021)

Despite these obstacles, there is a growing recognition among policymakers, experts, and the public of the need for a more proactive, equitable, and democratically accountable approach to AI governance (Milner & Traub, 2021; Viljoen, 2021; Zhang & Dafoe, 2020). Overcoming the current barriers will require a sustained effort to build political will, public awareness, and institutional capacity for a new paradigm of AI privacy governance. This may involve innovative legal strategies, multi-stakeholder partnerships, and public education campaigns to shift the discourse and create the conditions for meaningful reform. While the path forward is challenging, the stakes are too high to maintain the status quo in the face of generative AI's transformative impact on privacy, autonomy, and democracy.

5. Conclusion

The meteoric rise and widespread adoption of generative AI systems present significant threats to privacy and the regimes that seek to protect it. The ability of these systems to generate novel and realistic content based on patterns learned from vast troves of personal data raises profound risks, from nonconsensual data extraction and inferential profiling to the spread of synthetic media and the amplification of algorithmic biases. Existing regulatory approaches, which rely heavily on individual notice and consent, *ex post* enforcement, and a narrow conception of privacy harms, are ill-equipped to address generative AI's systemic and diffuse impacts. Addressing these challenges will require a fundamental reorientation of privacy governance from reliance on individual control and procedural safeguards to a more collective, proactive, and precautionary approach that recognizes privacy as a public good and collective responsibility.

Acknowledgements. Many thanks to Harry Surden, Ignacio Cofone, Neil M. Richards, and Paul Weitzel for their valuable input.

Funding statement. None.

Competing interests. None declared.

References

- 2024 *AI State Law Tracker*, April 2024. Retrieved May 3, 2024, from <https://www.huschblackwell.com/2024-ai-state-law-tracker>
- Ajunwa, I. (2020). The paradox of automation as anti-bias intervention. *Cardozo Law Review*, 41(5), 1671–1742.
- Ajunwa, I. (2021). An Auditing Imperative for Automated Hiring Systems. *Harvard Journal of Law & Technology*, 34(2), 621–699.
- AMG Capital Management v. FTC**, 141 S. Ct. 1341 (2021).
- Baio, A. (2022, August). *Exploring 12 million of the 2.3 billion images used to train stable diffusion's image generator*. Retrieved August 1, 2024, from <https://waxy.org/2022/08/exploring-12-million-of-the-images-used-to-train-stable-diffusions-image-generator/>
- Baker, C. R., & Chadwick, A. (2021). Corrupted infrastructures of meaning: Post-truth identities online. In H. Tumber & S. Waisbord (Eds.), *The Routledge companion to media misinformation and populism* (pp. 312–322). Routledge.
- Barrett, L. (2017). Reasonably suspicious algorithms: Predictive policing at the United States border. *New York University Review of Law and Social Change*, 41, 327.
- Bender, E. M., Gebru, T., McMillan-Major, A., & Shmitchell, S. (2021, March). On the dangers of stochastic parrots: Can language models be too big?. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, (pp. 610–623). <https://doi.org/10.1145/3442188.3445922>.
- Benthall, S., & Haynes, B. D. (2019). Racial categories in machine learning. In *Proceedings of the conference on fairness, accountability, and transparency (FAT* 19)* January 29–31, 2019, Atlanta, GA, USA (pp. 289–298). <https://doi.org/10.1145/3287560.3287575>.
- Bhargava, V. R., & Velasquez, M. (2020). Ethics of the attention economy: The problem of social media addiction. *Business Ethics Quarterly*, 31(3), 321–359. <https://doi.org/10.1017/beq.2020.32>
- Bloch-Wehba, H. (2020). Automation in moderation. *Cornell International Law Journal*, 53(1), 41–96.
- Bolukbasi, T., Chang, K.-W., Zou, J., Saligrama, V., & Kalai, A. (2016). Man is to computer programmer as woman is to homemaker? Debiasing word embeddings. *Advances in Neural Information Processing Systems*, 4349–4357. <https://doi.org/10.5555/3157382.3157584>.
- Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., Brynjolfsson, E., Buch, S., Card, D., Castellon, R., Chatterji, N., Chen, A., Creel, K., Davis, J. Q., Demszky, D., and Liang, P. (2022). On the opportunities and risks of foundation models, arXiv preprint arXiv:2108.07258.
- Buchanan, B., Lohn, A., Musser, M., & Sedova, K. (2021, May). *Truth, lies, and automation: How language models could change disinformation*. Center for Security and Emerging Technology. <https://doi.org/10.51593/2021CA003>
- Buolamwini, J., & Gebru, T. (2018). Gender shades: Intersectional accuracy disparities in commercial gender classification. In *Proceedings of the 1st Conference on Fairness, Accountability and Transparency 2018*, New York, NY, USA, 81, 77–91. <https://proceedings.mlr.press/v81/buolamwini18a/buolamwini18a.pdf>.
- Burrell, J. (2016). How the machine ‘thinks’: Understanding opacity in machine learning algorithms. *Big Data and Society*, 3(1), 1–12. <https://doi.org/10.1177/2053951715622512>
- California Consumer Privacy Act (CCPA)**, Cal. Civ. Code §§ 1798.100–1798.199 (West 2023).
- Caliskan-Islam, A., Bryson, J. J., & Narayanan, A. (2016). Semantics derived automatically from language corpora necessarily contain human biases. *Science*, 356(6334), 183–186. <https://doi.org/10.1126/science.aal4230>
- Calo, R. (2021). The boundaries of privacy harm. *Indiana Law Journal*, 95(1), 141–167.
- Calo, R., & Citron, D. K. (2020). The automated administrative state: A crisis of legitimacy. *Emory Law Journal*, 70, 797.
- Carlini, N., Liu, C., Erlingsson, Ú., Kos, J., & Song, D. (2019). The secret sharer: Evaluating and testing unintended memorization in neural networks. In *In 28th USENIX security symposium (USENIX security 19)* August 14–16, 2019, Santa Clara, CA, USA. (pp. 267–284) USENIX Association. <https://www.usenix.org/system/files/sec19-carlini.pdf>.
- Carlini, N., Tramer, F., Wallace, E., Jagielski, M., Herbert-Voss, A., Lee, K., Roberts, A., Brown, T., Song, D., Erlingsson, U., Oprea, A., & Raffel, C. (2021). Extracting training data from large language models. In *The 30th USENIX Security Symposium* August 11–13, 2021 Vancouver, B.C., Canada, (pp. 2633–2650). USENIX Association. <https://www.usenix.org/conference/usenixsecurity21/presentation/carlini-extracting>.
- Chander, A., Kaminski, M. E., & McGeveran, W. (2020). Catalyzing privacy law. *Minnesota Law Review*, 105(5), 1733–1802.
- Chesney, R., & Citron, D. K. (2019). Deep fakes: A looming challenge for privacy, democracy, and national security. *California Law Review*, 107(6), 1753–1820.
- Cinnamon, J. (2017). Social injustice in surveillance capitalism. *Surveillance & Society*, 15(5), 609–625.
- Citron, D. K. (2008). Technological due process. *Washington University Law Review*, 85, 1249–1313.
- Citron, D. K., & Pasquale, F. (2014). The scored society: Due process for automated predictions. *Washington Law Review*, 89(1), 1–33.
- Citron, D. K., & Wittes, B. (2017). The internet will not break: Denying bad Samaritans section 230 immunity. *Fordham Law Review*, 86(2), 401–423.
- Cohen, J. (2013). What Privacy Is For. *Harvard Law Review*, 126, 1904.
- Cohen, J. E. (2000). Examined lives: Informational privacy and the subject as object. *Stanford Law Review*, 52(5), 1373–1438.

- Cohen, J. E. (2019a). Turning privacy inside out. *Theoretical Inquiries in Law*, 20(1), 1–30.
- Cohen, J. E. (2019b). *Between truth and power: The legal constructions of informational capitalism*. Oxford University Press.
- Cooper, A. F., Lee, K., Grimmelmann, J., Ippolito, D., Callison-Burch, C., Choquette-Choo, C. A., Mireshghallah, N., Brundage, M., Mimno, D., Zahrah Choksi, M., Balkin, J. M., Carlini, N., De Sa, C., Frankie, J., Ganguli, D., Gipson, B., Guadamuz, A., Leng Harris, S., Jacobs, A. Z., Joh, E. & Kamath, G. (2023). Report of the 1st workshop on generative AI and law. *arXiv preprint arXiv:2304.05351*. <https://doi.org/10.48550/arXiv.2311.06477>.
- Corbett, J. (2024). *World's most extensive AI rules approved in EU despite criticism*. Bloomberg. <https://www.bloomberg.com/news/articles/2024-03-13/eu-embraces-new-ai-rules-despite-doubts-it-got-the-right-balance?embedded-checkout=true>.
- Crawford, K. (2021). *Atlas of AI: Power, politics, and the planetary costs of artificial intelligence*. Yale University Press.
- Davenport, T. H., & Mittal, N. (2022). How generative AI is changing creative work. *Harvard Business Review*, 100(4), 22–30.
- DiResta, A. E., & Sherman, Z. E. (2023, July). The FTC is regulating ai: A comprehensive analysis | Insights | Holland & Knight. Retrieved May 12, 2024, from <https://www.hklaw.com/en/insights/publications/2023/07/the-ftc-is-regulating-ai-a-comprehensive-analysis>
- Eaglin, J. M. (2017). Constructing recidivism risk. *Emory Law Journal*, 67, 59–122.
- Edwards, B. (2022, February 10). *AI-powered Bing Chat spills its secrets via prompt injection attack* [updated]. Ars Technica. <https://arstechnica.com/information-technology/2023/02/ai-powered-bing-chat-spills-its-secrets-via-prompt-injection-attack/>
- Edwards, L., & Veale, M. (2017). Slave to the algorithm: Why a right to an explanation is probably not the remedy you are looking for. *Duke Law & Technology Review*, 16, 18.
- Engstrom, D. F., & Haim, A. (2023). Regulating government AI and the challenge of sociotechnical design. *Annual Review of Law and Social Science*, 19(1), 277–298.
- European Commission. (2023). Proposal for a regulation of the European Parliament and of the Council laying down harmonised rules on artificial intelligence (artificial intelligence act) and amending certain Union legislative acts, COM(2023) 206 final. <https://eur-lex.europa.eu/legal-content/EN/TXT/?uri=CELEX%3A52023PC0206>
- Federal Trade Commission. (2020, April 8). Using artificial intelligence and algorithms: Federal Trade Commission guidance on AI. https://privacysecurityacademy.com/wp-content/uploads/2021/01/Using-Artificial-Intelligence-and-Algorithms_-_Federal-Trade-Commission.pdf.
- Federal Trade Commission. (2021). Aiming for truth, fairness, and equity in your company's use of AI. <https://www.ftc.gov/business-guidance/blog/2021/04/aiming-truth-fairness-equity-your-companys-use-ai>
- Federal Trade Commission. (2023a, September). Creative economy and generative AI. Retrieved May 13, 2024, from <https://www.ftc.gov/news-events/events/2023/10/creative-economy-generative-ai>
- Federal Trade Commission. (2023b). AI risk mitigation guidance for financial institutions. https://www.ftc.gov/system/files/ftc_gov/pdf/AI_Risk_Mitigation_Guidance_for_Financial_Institutions.pdf
- Federal Trade Commission Act § 5, 15 U.S.C. § 45(a) (2018).
- Federal Trade Commission v. Rite Aid Corp. 2024. Stipulated order for permanent injunction and civil penalty judgment, No. 2:23-cv-5023, E.D. Pa. Mar. 8
- Fernandez, V., Sanchez, P., Pinaya, W. H. L., Behrmann, J., Schölkopf, B., & Botvinick, M. (2023). Privacy distillation: Reducing re-identification risk of multimodal diffusion models. *arXiv preprint arXiv:2309.03690*. <https://doi.org/10.48550/arXiv.2306.01322>
- Fitzgerald, C., Williams, K., & Cross, R. J. (2024). The state of privacy: How state “privacy” laws fail to protect privacy and what they can do better.
- Franks, M. A. (2019). *The cult of the constitution: Our deadly devotion to guns and free speech*. Stanford University Press.
- FTC v. EverAlbum, Inc. (2021) No. 20-cv-03172.
- FTC v InMobi Pte Ltd. (2016) No. 3:16-cv-02192 (2016).
- Geddes, K. (2023). The death of the legal subject. *Vanderbilt Journal of Entertainment & Technology Law*, 25(1), 1–52.
- Gillis, T. B. (2022). The input fallacy. *Minnesota Law Review*, 106(4), 1175–1248.
- Grother, P. (2022). Face Recognition Vendor Test (FRVT) Part 8: Summarizing demographic differentials. *National Institute of Standards and Technology*, NISTIR 8429. <https://doi.org/10.6028/NIST.IR.8271>
- Grother, P., Ngan, M., & Hanaoka, K. (2019). Face Recognition Vendor Test (FVRT): Part 3, demographic effects. *National Institute of Standards and Technology*, NISTIR 8280. <https://doi.org/10.6028/NIST.IR.8280>
- Hanley, S., & Goldfarb, Z. (2021). FTC approves settlement with AI firm over biased algorithms. *The Washington Post*, January 29. <https://www.washingtonpost.com/technology/2021/01/29/ftc-facial-recognition-ai/>
- Harcourt, B. E. (2008). *Against prediction: Profiling, policing, and punishing in an actuarial age*. University of Chicago Press.
- Hartzog, W., & Richards, N. M. (2020). Privacy's constitutional moment and the limits of data protection. *Boston College Law Review*, 61(5), 1687–1761.
- Hartzog, W., & Solove, D. J. (2014). The FTC and the new common law of privacy. *Columbia Law Review*, 114(3), 583–676.
- Heikkilä, M. (2024, April 25). An AI startup made a hyperrealistic deepfake of me that's so good it's scary. *MIT Technology Review*. <https://www.technologyreview.com/2024/04/25/1091772/new-generative-ai-avatar-deepfake-synthesis/>

- Helmus, T. C., & Chandra, B., *Generative artificial intelligence threats to information integrity and potential policy responses* (2024)
- Hill, K. (2020). The secretive company that might end privacy as we know it. <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- HiQ Labs, Inc. v. LinkedIn Corp., 938 F.3d 985 (9th Cir. 2019).
- Hirsch, D. D. (2020). From individual control to social protection: New paradigms for privacy law in the age of predictive analytics. *Maryland Law Review*, 79(3), 439–505.
- Kalpokas, I. (2019). *A political theory of post-truth*. Palgrave Pivot. <https://doi.org/10.1007/978-3-030-20345-9>
- Kaminski, M. E. (2023). Regulating the risks of AI. *Boston University Law Review*, 103(10), 1347–1411.
- Katell, M., Young, M., Dailey, D., Herman, B., Guetler, V., Tam, A., ... Krafft, P. M. (2020). Toward situated interventions for algorithmic equity: Lessons from the field. In *Proceedings of the 2020 conference on fairness, accountability, and transparency* (pp. 45–55) Association for Computing Machinery.
- Kerr, I., & Earle, J. (2013). Prediction, preemption, presumption: How big data threatens big picture privacy. *Stanford Law Review Online*, 66, 65–72. <https://www.stanfordlawreview.org/online/privacy-and-big-data-prediction-preemption-presumption/>
- Kerry, C. F., Chin, C., & Lee, N. T. (2020). Bridging the gaps: A path forward to federal privacy legislation. <https://www.brookings.edu/research/bridging-the-gaps-a-path-forward-to-federal-privacy-legislation/>
- Kim, P. T. (2016). Data-driven discrimination at work. *William & Mary Law Review*, 58, 857.
- King, J., & Meinhardt, C. (2024). *Rethinking privacy in the AI era: Policy provocations for a data-centric world*. Stanford Cyber Policy Center. <https://hai.stanford.edu/sites/default/files/2024-02/White-Paper-Rethinking-Privacy-AI-Era.pdf>
- Kosseff, J. (2019). *The twenty-six words that created the Internet*. Cornell University Press.
- Kreacik, A., Uribe, L., Lasater-Wille, A., & Romeo, J. (2024). *How generative AI is transforming business and society: The good, the bad, and everything in between*. MIT Press. <https://www.oliverwymanforum.com/content/dam/oliver-wyman/ow-forum/gcs/2023/AI-Report-2024-Davos.pdf>
- Kroll, J. A., Huey, J., Barocas, S., Felten, E. W., Reidenberg, J. R., Robinson, D. G., & Yu, H. (2017). Accountable algorithms. *University of Pennsylvania Law Review*, 165, 633.
- Lamo, M., & Calo, R. (2019). Regulating bot speech. *UCLA Law Review*, 66(4), 988–1028.
- Lazaro, C. (2018). Algorithmic divination: From prediction to preemption of the future. *Law, Technology and Culture*, 20(1), 1–28.
- Lee, K., Ippolito, D., & Cooper, A. F. (2024). The devil is in the training data. *Technology Policy Working Paper Series*.
- Leffer, L. (2024, January 15). *AI audio deepfakes are quickly outpacing detection*. Scientific American. <https://www.scientificamerican.com/article/ai-audio-deepfakes-are-quickly-outpacing-detection/>
- Leffer, L. (n.d.). *Your personal information is probably being used to train generative AI models*. Scientific American. <https://www.scientificamerican.com/article/your-personal-information-is-probably-being-used-to-train-generative-ai-models/>
- Luccioni, A. S., & Viviano, J. D. (2021). What's in the box? A preliminary analysis of undesirable content in the common crawl corpus. *arXiv preprint arXiv:2105.02732*.
- MacKenzie, A., & Bhatt, I. (2020). Lies, bullshit and fake news: Some epistemological concerns. *Postdigital Science and Education*, 2(1), 9–13. <https://doi.org/10.1007/s42438-018-0025-4>
- Massaro, T. M., Norton, H., & Kaminski, M. E. (2017). SIRI-OUSLY 2.0: What artificial intelligence reveals about the First Amendment. *Minnesota Law Review*, 101, 2481–2525. doi:<https://scholarship.law.umn.edu/mlr/179/>
- Matsumi, H., & Solove, D. J. (2024). The prediction society: AI and the problems of forecasting the future. *Stanford Law Review*, forthcoming. https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4075677
- Mayson, S. G. (2018). Bias in, bias out. *Yale Law Journal*, 128, 2218.
- Meijer, A., & Wessels, M. (2019). Predictive policing: Review of benefits and drawbacks. *International Journal of Public Administration*, 42(12), 1031–1039.
- Metcalfe, J., Moss, E., Watkins, E. A., Singh, R., & Elish, M. C. (2021). Algorithmic impact assessments and accountability: The co-construction of impacts. In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency (FAccT '21)*, Association for Computing Machinery, New York, NY, USA (pp. 735–746).
- Milner, Y., & Traub, A. (2021). *Data capitalism and algorithmic racism*. <https://dataspace.princeton.edu/handle/88435/dsp01z603r1587>.
- Morrison, S. (2023, July 12). *The tricky truth about how generative AI uses your data*. Vox. <https://www.vox.com/recode/2023/7/12/23786023/generative-ai-data-privacy-consent>
- O'Neil, C. (2016). *Weapons of math destruction: How big data increases inequality and threatens democracy*. Crown.
- OpenAI, Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F. L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., Avila, R., Babuschkin, I., Balaji, S., Balcom, V., Baltescu, P., Bao, H., Bavarian, M., Belgum, J., ... Zoph, B. (2024). GPT-4 Technical Report.
- Pasquale, F. (2015). *The black box society: The secret algorithms that control money and information*. Harvard University Press.
- Powles, J., & Nissenbaum, H. (2018). *The seductive diversion of 'solving' bias in artificial intelligence*. One Zero, 7. https://nissenbaum.tech.cornell.edu/papers/The_Seductive_Diversion.pdf

- Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 on the protection of natural persons with regard to the processing of personal data and on the free movement of such data, and repealing Directive 95/46/EC (General Data Protection Regulation), OJ 2016 L 119/1.
- Rubin, A. (2023, November 3). *Teens exploited by fake nudes illustrate threat of unregulated AI*. Axios. <https://www.axios.com/2023/11/03/ai-deepfake-nude-images-new-jersey-high-school>
- Salam, E. (2023, June 14). US mother gets call from “kidnapped daughter” – But it’s really an AI scam. *The Guardian*. <https://www.theguardian.com/technology/2023/jun/14/us-mother-call-kidnapped-daughter-ai-scam-deepfake>
- Schuhmann, C., Vencu, R., Beaumont, R., Kaczmarczyk, R., Mullis, C., Katta, A., Coombes, T., Jitsev, J., & Komatsuzaki, A. (2022). LAION-5B: An open large-scale dataset for training next generation image-text models. In *36th Conference on Neural Information Processing Systems (NeurIPS 2022), Track on Datasets and Benchmarks*. <https://doi.org/10.48550/arXiv.2210.08402>
- Selbst, A. D., & Barocas, S. (2018). The intuitive appeal of explainable machines. *Fordham Law Review*, 87, 1085.
- Shelby, R., Rismani, S., Henne, K., Venema, P., & Passi, S. (2023). Sociotechnical harms of algorithmic systems: Scoping a taxonomy for harm reduction. In *Proceedings of the 2023 AAAI/ACM Conference on AI, Ethics, and Society* (pp. 723–741). Association for Computing Machinery. <https://doi.org/10.1145/3544548.3581282>
- Shin, D. (2024). Conclusion: Misinformation and AI—How algorithms generate and manipulate misinformation. In *Artificial misinformation: Exploring human-algorithm interaction online* (pp. 259–277). Springer.
- Shokri, R., Stronati, M., Song, C., & Shmatikov, V. (2017). Membership inference attacks against machine learning models (pp. 3–18). IEEE.
- Solove, D. J. (2013). Privacy self-management and the consent dilemma. *Harvard Law Review*, 126(7), 1880–1903. <https://harvardlawreview.org/2013/05/introduction-privacy-self-management-and-the-consent-dilemma/>
- Solove, D. J. (2022). The limitations of privacy rights. *Notre Dame Law Review*, 98(3), 975–1036.
- Solove, D. J. (2024). Murky consent: An approach to the fictions of consent in privacy law. *Boston University Law Review*, 104, 593–639.
- Solove, D. J. (2025). Artificial intelligence and privacy. 77 *Florida Law Review* (forthcoming, January 2025), preprint draft at https://papers.ssrn.com/sol3/Papers.cfm?abstract_id=4713111
- Solove, D. J., & Matsumi, H. (2024). AI, algorithms, and awful humans. *Fordham Law Review*, 92(5), 1923–1956.
- Solow-Niederman, A. (2022). Information privacy and the inference economy. *Northwestern University Law Review*, 117(1), 357–424.
- Song, C., & Shmatikov, V. (2019). Overlearning reveals sensitive attributes. *arXiv preprint arXiv:1905.11742*. <https://doi.org/10.48550/arXiv.1905.11742>
- Staab, R., Vero, M., Balunović, M., & Vechev, M. (2023). Beyond memorization: Violating privacy via inference with large language models. *arXiv preprint arXiv:2310.07298*. <https://doi.org/10.48550/arXiv.2310.07298>
- Surden, H. (2024). ChatGPT, large language models, and law. *Fordham Law Review*, 92(6), 1941–1985.
- Tangalakis-Lippert, K. (2023, April 15). *Clearview AI scraped 30 billion images from Facebook and other social media sites*. Business Insider. <https://www.businessinsider.com/clearview-ai-30-billion-images-facebook-twitter-youtube-scraped-2023-4>
- Tenn. Code Ann. §§ 47-25-1101 to 47-25-1108 (2023).
- Thierer, A. (2016). *Permissionless innovation: The continuing case for comprehensive technological freedom*. Mercatus Center at George Mason University.
- Tisné, M. (2020). *The data delusion: Protecting individual data isn't enough when the harm is collective*. Stanford Cyber Policy Center White Paper. https://fsi-live.s3.us-west-1.amazonaws.com/s3fs-public/the_data_delusion_formatted-v3.pdf
- Tokaji, D. P. (2019). Truth, democracy, and the limits of law. *Saint Louis University Law Journal*, 64(1), 569–594.
- Tschider, C. A. (2021). AI’s legitimate interest: Towards a public benefit privacy model. *Houston Journal of Health Law & Policy*, 21, 125.
- Utah Consumer Privacy Act, Utah Code Ann. §§ 13-61-101 (2024).
- Van der Sloot, B., & Wagenveld, Y. (2022). Deepfakes: Regulatory challenges for the synthetic society. *Computer Law & Security Review*, 46, 105716. <https://doi.org/10.1016/j.clsr.2022.105716>
- Veale, M., & Zuiderveen Borgesius, F. (2021). Demystifying the draft EU artificial intelligence act—Analysing the good, the bad, and the unclear elements of the proposed approach. *Computer Law Review International*, 22(4), 97–112.
- Véliz, C. (2021). If AI is predicting your future, are you still free? <https://www.nytimes.com/2020/01/18/technology/clearview-privacy-facial-recognition.html>
- Verma, P. (2023, November 4). AI fake nudes are booming. It’s ruining real teens’ lives. *The Washington Post*. <https://www.washingtonpost.com/technology/2023/11/05/ai-deepfake-porn-teens-women-impact/>
- Viljoen, S. (2021). Democratic data: A relational theory for data governance. *Yale Law Journal*, 131(2), 573–654.
- Villaronga, E. F., Kieseberg, P., & Li, T. (2018). Humans forget, machines remember: Artificial intelligence and the right to be forgotten. *Computer Law & Security Review*, 34(2), 304–313.
- Virginia Consumer Data Protection Act (VCDPA), Va. (2021). Code §, 59.1–571.
- Wachter, S., & Mittelstadt, B. (2019). A right to reasonable inferences. *Columbia Business Law Review*, 2019, 494–620.

- Waldman, A. E. (2019). Privacy law's false promise. *Washington University Law Review*, 97, 773–834.
- Wang, Y., & Kosinski, M. (2018). Deep neural networks are more accurate than humans at detecting sexual orientation from facial images. *Journal of Personality and Social Psychology*, 114(2), 246–257. <https://doi.org/10.1037/pspa0000098>
- Weidinger, L., Mellor, J., Rauh, M. et al. (2021). Ethical and social risks of harm from language models. *arXiv preprint arXiv:2112.04359*.
- Weidinger, L., Uesato, J., Rauh, M., Griffin, C., Huang, P.-S., Mellor, J., Glaese, A., Cheng, M., Balle, B., Kasirzadeh, A., Gabriel, I., Muennighoff, N., Higgins, I., Song, X., Nichol, A., Diamanti, A., Coppin, T., Gao, Z., El-Showk, S., ... Irving, G. (2022). Taxonomy of risks posed by language models. In *Proceedings of the 2022 ACM conference on fairness, accountability, and transparency* (pp. 214–229). Association for Computing Machinery. <https://doi.org/10.1145/3531146.3533088>
- West, S. M., Whittaker, M., & Crawford, K. (2019). *Discriminating systems: Gender, race and power in AI*. AI Now Institute. <https://ainowinstitute.org/wp-content/uploads/2023/04/discriminatingystems.pdf>
- Westerlund, M. (2019). The emergence of deepfake technology: A review. *Technology Innovation Management Review*, 9(11), 39–52. <https://doi.org/10.22215/timreview/1282>
- Wexler, R. (2018). Life, liberty, and trade secrets: Intellectual property in the criminal justice system. *Stanford Law Review*, 70, 1343–1430.
- Whittaker, M., Alper, M., Bennett, C., Hendren, S., Kazianas, L., Mills, M., Morris, M., Rankin, J., Rogers, E., Salas, M., & Myers West, S. (2019). *Disability, bias, and AI*. AI Now Institute. <https://ainowinstitute.org/publication/disabilitybiasai-2019>
- Whyte, C. (2020). Deepfake news: AI-enabled disinformation as a multi-level public policy challenge. *Journal of Cyber Policy*, 5(2), 199–217. <https://doi.org/10.1080/23738871.2020.1797135>
- Winograd, A. (2023). Loose-lipped large language models spill your secrets: The privacy implications of large language models. *Harvard Journal of Law & Technology*, 36(2), 615–648. <https://jolt.law.harvard.edu/assets/articlePDFs/v36/36HarvJLTech615.pdf>
- Wu, T. (2012). Machine speech. *University of Pennsylvania Law Review*, 161, 1495.
- Zeide, E. (2015, September 15). Algorithms can be lousy fortunetellers. *Slate*. http://www.slate.com/articles/technology/future_tense/2015/09/what_happens_when_big_data_uses_bad_data_to_make_predictions.html
- Zeide, E. (2017). The structural consequences of big data-driven education. *Big Data*, 5(2), 164–172. <https://doi.org/10.1089/big.2016>
- Zeide, E. (2022). The silicon ceiling: How artificial intelligence constructs an invisible barrier to opportunity. *UMKC Law Review*, 91, 403–436.
- Zhang, B., & Dafoe, A. (2020, February). U.S. Public Opinion on the Governance of Artificial Intelligence. In *proceedings of the AAAI/ACM Conference on AI, Ethics, and Society (AIES '20)*. Association for Computing Machinery, New York, NY, USA (pp. 187–193). <https://doi.org/10.1145/3375627.3375827>.

Elana Zeide is an Assistant Professor at the University of Nebraska College of Law, where she examines how artificial intelligence and digital systems affect privacy, equity, and opportunity. Her scholarship focuses on the governance of technology in education and hiring contexts. She has published extensively on student privacy and algorithmic decision-making, exploring how law can promote responsible innovation. She advises schools, companies, and policymakers on the ethical implementation of automated systems.