TRANSPARENCY
COALITION

PRIVACY HARMS IN THE AI AGE

Time for a System Update

BY LEIGH WICKELL

CONTENTS

AUTHOR

LEIGH WICKELL

Leigh Wickell is a data privacy attorney and a legal advisor to the Transparency Coalition. She has worked in legal and ethical compliance for eight years, including four years specializing in data privacy compliance at Mastercard and SAP Concur.

The Transparency Coalition is a non-profit organization dedicated to supporting the creation of AI laws, regulations, and policies that allow innovation to flourish while protecting the data rights and privacy of individuals and society.

Visit www.transparencycoalition.ai to learn more.

EXECUTIVE SUMMARY

Many of today's most powerful artificial intelligence (AI) systems process personal data in ways that do not comply with U.S. privacy laws, resulting in widespread privacy harms. In the United States, current regulatory and judicial systems are failing to deter these violations. Federal and state regulators do not have the capacity to properly prevent privacy violations. Supreme Court decisions over the past twenty years have eroded the ability of individuals to bring private rights of action against tech companies for privacy harms.

In light of this, policymakers must adjust their approach to privacy violations in the AI age. This should begin with an expansion of the legal understanding of privacy harms, and continue with a more serious consideration of statutory violations during the adjudication of private rights of action.

In order to ensure the proper handling of personal data within the burgeoning AI industry, three main actions are proposed:

1) **Federal and state policymakers, regulators, and enforcement agencies must step up enforcement of existing privacy laws.** This should include an expanded understanding of "actual harm" suffered by those whose personal data have been used without consent.

2) **AI developers must be required to fully describe their training data, and make that description available to consumers and regulators.** The description must include information regarding any personal information included in the datasets.

3) **Policymakers must create an appropriate framework for direct state and federal oversight of the AI industry.** This should include direct regulatory access to AI training data via API, as well as regulatory review of new AI models.

The global AI industry is growing at a frenetic pace, with many of the largest and most well-funded models emerging from companies based in the United States. American lawmakers have the duty and responsibility to establish a paradigm of transparency and accountability now, before these systems and companies become too large to regulate.

SECTION I
INTRODUCTION

Artificial intelligence (AI) has become a powerful force in our individual lives. AI systems now influence how employers hire job candidates, how marketers reach specific consumers, and how healthcare providers interact with patients. Public-facing chatbots such as OpenAI's ChatGPT and Microsoft's Copilot are changing the ways individuals search for information. Generative AI platforms like Midjourney allow users to conjure sophisticated visual creations using only a few text prompts.

The excitement around AI has been tempered by instances in which AI systems have produced unexpected — and alarming — results. For example, a mayoral candidate in Australia learned he was the subject of an AI "hallucination," in which the AI system falsely claimed he was part of a criminal scheme.[1] He is now suing the company for defamation. Another person was incorrectly listed as deceased in an obituary written by a faulty AI, startling his friends and family.[2] In Georgia, a radio host filed suit against OpenAI after the company's ChatGPT system falsely stated that he had been accused of fraud and embezzlement.[3]

These problems run deeper than false accusations. Many AI models have trained on datasets that contain personal information, and every AI model is capable of regurgitating any data it has previously processed. In 2023, researchers demonstrated that leading AI image diffusion models such as DALL-E, Imagen, and Stable Diffusion "do memorize and regenerate individual training examples."[4]

Nicholas Carlini, a research scientist at Google DeepMind, has observed that "neural networks often leak details of their training sets."[5] Carlini offers further examples of attack strategies that may completely recover subsets of training data. Several well-known measures are available to guard against these leaks, but are not commonly implemented by AI model developers.

For purposes of this paper, "personal information" describes any information relating to an identified or identifiable person.[6] This can include contact information, but may also include information like a person's preferences for internet settings or purchasing history, even if the person's contact information is not directly tied to those preferences or purchases. This conception of personal information originated in the European Union's General Data Protection Regulation[7], and has since been adopted by several U.S. state privacy laws, such as the California Consumer Privacy Act[8] and the Virginia Privacy and Data Protection Act.[9]

When an AI model discloses personal information or “hallucinates” about a person, it indicates the AI could have been trained on datasets containing personal information, and the model may invent additional information in the model output. Large language models like ChatGPT, for example, are only “trained to ‘produce a plausible sounding answer’ to user prompts,” regardless of the accuracy of the information in the answer.[10]

Using personal data to train AI models is highly controversial for several reasons, some of which this paper will discuss. Though AI developers have faced criticism in many areas, such as the lack of accountability for intellectual property violations and discrimination against protected classes, this paper will focus on addressing the privacy harms that occur when AI models improperly process personal information.

Hallucinations

Large language models (LLMs) like ChatGPT are not trained to seek factual accuracy. They are trained to produce a plausible sounding answer to a prompt. As a result, AI chatbots are known to occasionally produce information that is completely fabricated. These made-up answers are known as hallucinations.

Some experts believe hallucinations occur because AI systems compress the massive amounts of data used to train them, and details are lost in that compression. An AI system will seek to fill that gap in detail with plausible — but fictional — details.

SECTION II
PRIVACY HARMS

The right to privacy in the United States has not always been clearly understood. In 1960, legal scholar William Prosser clarified the right by describing four types of personal harms resulting from violations of a person's right to privacy.[11] Those types, elucidated in the chart below, are commonly referred to as intrusion, public disclosure, false light, and appropriation.

Prosser's privacy harms framework has been a foundation of privacy jurisprudence for more than a half-century. But it is not equipped to deal with today's information landscape, and often fails to address harms resulting from today's AI models altogether.[12] While a company's use of an individual's personal information to train its AI models without the individual's consent could be construed by a reasonable person as an intrusion or appropriation harm, it is unlikely to meet harm requirements imposed by courts on privacy cases in the Internet era.[13] Even where a company directly profits from the use of personal information that the data subject did not consent to, courts have been reluctant to find this situation meets harm thresholds under Prosser's appropriation theory. This trend is exemplified by the Appellate Court of Illinois decision in *Dwyer v. American Express.*[14] In that case, a group of cardholders sued American Express for disclosing their spending data to participating merchants as part of a joint

PROSSER'S PRIVACY HARMS (1960)	EXAMPLES
INTRUSION	The right to be free from unreasonable intrusion upon a person's seclusion, solitude, or private affairs.
PUBLIC DISCLOSURE	The right to be protected from unreasonable publicity given to a person's private life, such as the publication of an embarrassing private fact or photograph.
FALSE LIGHT	This is the harm to one's reputation due to the publication of a false statement that would be offensive to a reasonable person.
APPROPRIATION	Also called misappropriation or the right of publicity, appropriation encompasses the right of an individual to exclusively control the use of one's name and image in advertising, merchandise, and other forms of commerce.

marketing program. The plaintiffs claimed the practice violated their privacy and consumer rights under the Illinois Consumer Fraud and Deceptive Business Practices Act. The court rejected the claim, finding that the practice did “not deprive any of the cardholders of any value their individual names may possess,” and thus did not cause the plaintiffs any provable harm[15].

Indeed, many courts have found the risk of actual harm from privacy violations to be too tenuous to establish standing on which to bring a case. When a company merely processes a person’s information without their knowledge and that information is not 1) embarrassing, 2) misleading, nor 3) disclosed to the public at large, Prosser’s traditional harm framework may not recognize that the individual has suffered any harm at all.[16] In order to properly compensate individuals whose personal data has been processed inappropriately and to deter unfair practices by technology companies, an updated legal conception of privacy harms is needed.

A good place to start is with the work of Danielle Keats Citron and Daniel J. Solove.[17] In their 2022 paper, "Privacy Harms," Citron and Solove build on Prosser's four types of privacy harms and include five new categories: physical, economic, reputational, psychological, and autonomy harms. The chart on page 9 outlines their updated privacy framework and offers examples of each type of harm.

> Many privacy laws in the United States are enforced through a combination of regulatory oversight and private rights of action.

In addition to their expansion of privacy harm categories, Citron and Solove propose a shift of focus in privacy law: balancing the overriding importance of harm with a weightier consideration of statutory violations. Many privacy laws in the United States are enforced through a combination of regulatory oversight and private rights of action.[18] Due to the overburdened state of regulatory bodies, the right of private citizens to bring suit under these laws fills a yawning enforcement gap and ensures companies violating privacy laws are held accountable.[19] Citron and Solove argue that the decades-long trend of state and federal courts downplaying privacy harms — especially at the highest level (see sidebar, page 10) — has effectively eliminated a critically important mechanism for the enforcement of privacy laws.[20] In fact, the U.S. Supreme

CITRON AND SOLOVE'S PRIVACY HARMS (2022)	EXAMPLES
PHYSICAL HARMS	A stalker is able to find personal information about the victim online, even though the victim did not consent to disclosure.
ECONOMIC HARMS	A consumer suffers identity theft after an organization processing the consumer's personal data suffers a data breach.
REPUTATIONAL HARMS	An error in a consumer's credit report or background check leads to the consumer being denied a line of credit or a job.
PSYCHOLOGICAL HARMS	A person is worried that a security breach at their bank will result in identity theft and/or financial loss. A person receives intrusive text messages or phone calls in violation of the Telephone Communications Privacy Act.
AUTONOMY HARMS	A person does not want to share their location data with a mobile messaging application, but is required to share this data in order to use the app. A person submits their phone number to an organization for a specific purpose, but the organization has not informed the person that their phone number will be shared with marketing companies. A person submitted their personal information to a social media site several years ago, and the site continues using the personal information for additional purposes indefinitely.

Court's focus on demonstrable harm has become so ardent that even where privacy laws do not require a plaintiff to demonstrate harm resulting from the defendant's violation in order to recover, courts have imposed harm requirements anyway, rendering many privacy laws toothless.[21]

This rising harm bar has blocked the ability of many plaintiffs to establish the standing needed for their case to be heard. Under current standing doctrine, plaintiffs must allege an "actual or imminent, not conjectural or hypothetical" injury in fact, one that is "concrete and particularized." In one recent case, *TransUnion LLC v. Ramirez* (2021)[22], the Supreme Court spelled this out clearly. In *TransUnion*, the consumer credit reporting agency erroneously labeled the plaintiffs as potential terrorists in their credit reports. The Court found that wronged plaintiffs whose credit reports had not yet been circulated had

Eroding Privacy by Raising the 'Harm Bar'

Over the past 20 years the U.S. Supreme Court has undermined the effectiveness of privacy laws by demanding evidence not merely of lawbreaking, but of significant individual harm to each plaintiff in a private right of action. Significant SCOTUS cases are listed below.

1992, *Lujan v. Defenders of Wildlife*	Justice Scalia creates three-part test to establish standing in federal court; requires an actual or imminent “injury in fact.”
2004, *Doe v. Chao*	Court holds that statutory damages under the federal Privacy Act of 1974 are allowed only if plaintiffs establish “actual” damages.
2012, *FAA v. Cooper*	Court holds that “mental or emotional distress” does not constitute actual damage under the Privacy Act.
2016, *Spokeo Inc. v. Robins*	Court holds that even if a legislature grants the right to recover without proving harm, a court may impose its own requirement to prove harm in order to establish standing.

not suffered an actual injury — only those whose reports had been disseminated. Therefore, consumers whose reports were not circulated could not sue TransUnion. “No concrete harm, no standing,” wrote Justice Kavanaugh for the majority.

Unfortunately, the intangible nature of most types of privacy harms means many of these suits cannot succeed under this requirement of "concrete harm."[23]

Citron’s and Solove’s taxonomy of privacy harms especially in its recognition of psychological and autonomy harm, succeeds in addressing privacy harms that occur frequently in the online economy.

Given the tendency of high courts to dismiss the legitimacy of psychological harms, however, there is a further need to quantify these harms in terms more likely to be recognized by the bench. In cases of data breach, emotional reactions of fear or anxiety often lead to a loss of real time and money (in the effort required to cancel credit cards and change passwords, for instance, or install stronger home security systems). The monetary value of data privacy itself remains largely speculative, although some scholars are beginning to quantify it. Angela W. Winegar and Cass R. Sunstein did preliminary work along these lines in a 2019 paper, “How Much Is Data Privacy Worth? A Preliminary Investigation.”[24] They found that surveyed consumers would pay $5 per month to maintain data privacy, but cautioned against the reliance on that figure due to a lack of information regarding the full value of data and the risk of privacy loss. More recent work by economists Jeffrey Prince and Scott Wallsten[25] found that U.S. consumers would demand between $3.50 and $8 per month from Facebook to share their contact information with third parties. These studies demonstrate the public's understanding that their personal information is valuable to corporations like Meta, the parent company of Facebook. They also show, in stark terms, that consumers value the ability to make decisions about what companies can do with their personal information — so much that consumers are willing to pay for that control. It's time for courts to recognize this value when demonstrable harm is required to bring suit.

Given the fact that the U.S. Supreme Court has established a requirement that plaintiffs show that they have suffered a "concrete harm," state and federal legislators should include language in privacy legislation explicitly stating that private rights of action may be brought based on legal violations alone and do not require a showing of harm. In other words, protecting privacy in the AI era will require the legislative branch to push back on the judiciary's restrictions on lawsuits brought under existing privacy laws.

Some may balk, but SCOTUS overstepped and undermined legislative intent in *Spokeo Inc. v. Robins* by allowing a court to require proof of harm even when not required by law. If new privacy laws are to ensure protection of personal information in the AI age, legislators must explicitly unblock the private right of action as a key enforcement mechanism.

> The monetary value of data and privacy must be established through more research before it can be used to clear Justice Scalia's three-part test for actual harm.

SECTION III

PRIVACY HARMS AND ARTIFICIAL INTELLIGENCE

As companies scramble to develop ever larger and more sophisticated AI models, the scope of data used to train those models grows wider and wider. The general rule when developing AI models is that more training data results in a better model.[26] Many of the most well-known AI models today have been trained using huge volumes of data scraped directly from the web.[27] While some large companies have developed their own web scraping tools to build training data sets, others license training datasets from other companies.[28] Due to the way that web scrapers tend to vacuum up data indiscriminately and outsource the task of maintaining data hygiene, many companies may not have a full view of the data being used to train their AI models.[29] Datasets created using web scraping are frequently found to include personal information, even if the person to whom the personal information relates did not consent to the use of their personal information to train AI models.[30] Some anomalous AI outputs indicate that personal health information may be used to train some models.[31] These privacy issues have been traced to AI models developed by Google, OpenAI, and Meta,[32] indicating that this is an industry-wide problem.

Similarly concerning are the unpredictable ways AI models can disclose personal information. Lack of clarity around how ChatGPT processes the data sent to it in prompts — which could include sensitive personal information such as a specific health condition or a troubling family conflict — mean there is a risk that ChatGPT could disclose this sensitive data to others.[33] Examples of personal information revealed in an AI model's output can include personal health data,[34] personal information exactly as it appeared in the training data for the model,[35] and even deidentified data that was combined with additional data which allowed the model to successfully re-identify the data.[36]

There are additional concerns when AI systems produce incorrect personal information. Large language models like ChatGPT are only trained to produce a plausible-sounding answer to user prompts, regardless of the accuracy.[37] This has resulted in a number of well-documented errors.

One expert theorizes that these hallucinations occur because AI systems compress the incredible amounts of data used to train them, and through that compression, they lose the details of the data.[38] When a user requests those details, the AI model will simply "'mak[e] things up'" to fill in the gaps left by compression of the training data.[39]

The examples cited in Section I — the Australian mayoral candidate, the Georgia radio host, and the living person described as deceased — aren't the only individuals negatively affected by AI hallucinations about them. Indeed, the hallucination problem is so widespread that Max Schrems, the Austrian attorney who inspired the creation of the EU's General Data Protection Regulation, has sued OpenAI over it.[40] Schrems's complaint states that when queried about Schrems's birthdate, the company's ChatGPT system provided incorrect dates multiple times. When Schrems requested that ChatGPT delete the incorrect information and provide accurate information instead, OpenAI failed to fulfill the request. The GDPR requires that OpenAI honor requests from data subjects to correct and delete their personal information, as well as a general obligation on data controllers to ensure the data they process is accurate.[41] Schrems argues that OpenAI violated the GDPR by failing to honor his requests regarding his personal data, and by processing inaccurate personal data.[42]

Consumer reaction to revelations about their personal data being used to train AI models shows that individuals feel harmed by these practices.[43] Public outcry over the use of personal information in AI training datasets exemplifies many of the wrongs contained within Citron and Solove's rubric of autonomy harms: lack of control, failure to inform, thwarted expectations, coercion, and manipulation. In an article describing the public backlash against Zoom Inc.'s use of video chat data in AI training datasets, *Axios* writers noted that consumers "often have little option beyond clicking the agree button or to stop using a service entirely."[44] These options fail to give individuals any real choice regarding how their personal data will be treated — a manipulative and coercive scenario the FTC has enforced against in the past as an unfair business practice.[45]

It's important to remember that when a company trains their AI models on personal data, the company derives economic benefit from that data. Though courts have largely failed to recognize that plaintiffs are harmed when personal information is used in AI training data, the FTC has not.[46] In a 2023 complaint, the FTC alleged that Amazon violated the Children's Online Privacy Protection Act (COPPA) by retaining voice recordings of children to train its Alexa AI systems, even when parents had requested deletion.[47] The FTC explicitly described how Amazon realized economic benefits from this data: "Children's

COPPA: Protecting Children's Privacy

The Children's Online Privacy Protection Act of 1998 (COPPA) protects the privacy of children online by requiring parental consent for the collection or use of any personal information of users under 13 years of age. The Act was passed due to concern over Internet marketing techniques targeting children without parental consent.

unique speech patterns and accents differ from adults, and their voice recordings provide Amazon with a valuable data bank for training the Alexa algorithm to understand children," and that this data bank was used "for purposes such as refining Alexa's voice recognition and natural language processing capabilities." [48] In a press release about the proposed court order, the FTC stated that this practice "benefit[ed] its bottom line at the expense of children's privacy." [49] In other words, a company using personal data to train an AI model derives economic benefit from that data through the improvement of the AI model's capabilities. Even where the harm to an individual person may be minor, the cumulative effect of using personal data to train AI models results in societal harm where public trust in the right to privacy suffers death by a thousand cuts, all to increase the company's own profits.[50]

Even where privacy laws offer consumers the ability to remove their personal information from AI training datasets, the companies in charge of those AI models can undermine and circumvent those choices, leading to further harm. For example, some U.S. privacy laws grant consumers the right to demand companies delete any personal information about an individual in the company's data stores.[51] However, the lack of proper record-keeping and annotation in datasets can make it impossible for companies to find and delete all information about a person.[52] The FTC complaint against Amazon alleges that Amazon failed to delete written transcripts and backups of children's conversations with the Alexa system, despite telling parents requesting deletion that it had done so.[53] On top of this, simply deleting the information from the training data does not erase it from the AI model that was trained on that data. Short of retraining the entire AI model — at a cost of tens of millions of dollars — accomplishing that task may be impossible. Upon being trained on a dataset, an AI model cannot be programmed to "forget" a specific portion of data. Once seen, it cannot be unseen.[54] The difficulty of deleting the data can be so great that some companies may simply decline deletion requests from individuals, even though this approach may not comply with applicable law.[55]

Upon being trained on a dataset, an AI model cannot be programmed to "forget" a specific portion of data. Once seen, it cannot be unseen.

SECTION IV
POTENTIAL SOLUTIONS

The U.S. has several options to reduce privacy harms caused by AI model development and deployment.

The first and most obvious action is to step up enforcement of existing privacy laws. The FTC has brought enforcement actions against several companies for privacy harms to autonomy resulting from manipulation and thwarted expectations of consumer choices, among others.[56] In one of the most striking recent cases, the FTC alleged that the online mental health service BetterHelp shared private customer health information with Facebook, Instagram, and other third parties for marketing purposes, which the FTC described as a deceptive business practice and a likely violation of HIPAA requirements.[57] We can see that the FTC has successfully enforced existing privacy laws when personal information is processed in ways that consumers would not expect. The FTC now needs to enforce these laws in the context of AI. Even a federal agency as powerful as the FTC, however, works within limited resources, which is why private rights of action can function as an important secondary enforcement mechanism.

Section II described the difficulties faced by plaintiffs in a private right of action to show harm in cases that involve personal data and privacy. Given the case law of the past thirty years, and *Spokeo* in particular, the "concrete harm" requirement imposed by SCOTUS must be countered by explicit legislative repudiation establishing private rights of action as privacy enforcement tools. Legislatures should also encourage the growth of research in the valuation of data and privacy. The question, ‘What is a person’s personal data truly worth?’ remains largely unanswered — but when research provides an answer, plaintiffs will have a much stronger foundation for standing in private rights of action. Those legal actions should serve as a strong deterrent against the practice of developing AI models with little regard for the unauthorized use of personal data.

A second way to improve the protection of personal information within AI training datasets would be to require AI developers to describe each of these datasets in detail. Transparency is an important principle of data privacy. As such, state and federal laws should require AI systems to come with a data card describing the dataset used to train the AI system.[58] This data card could include information such as: each type of data element included in the data set; the source or owner of the data set and how the data set was obtained; the time period during which the data was collected; and the dates the dataset was used to train and re-train the AI model. If a dataset used to train AI models includes personal information, the AI owner must establish a clear legal basis for their use of the data. The advantage of this approach is the balance between a relatively low burden on the AI model owner and the at-a-glance transparency regarding the personal information being processed. A number of different data card formats have been proposed in the past few years, including the Dataset Nutrition Label developed by Consumer Reports, the Datasheets for Datasets concept from Timnit Gebru et al, and the Google Research team's Data Card.[59] Though an industry standard has yet to emerge, this is a promising concept that the Transparency Coalition (TCAI) views as a necessary tool for the ethical development of AI models. TCAI co-founders Rob Eleveld and Jai Jaisimha have put forward a further refinement of the data card concept, a Data Declaration that would include specific documentation required for transparency in AI model development. The elements included in that Data Declaration are listed in the chart below.

Documentation required in a Data Declaration

FIELD NAME	POSSIBLE VALUES
Data Set Name	Text
Data Set Owner	Text
Data Set Description	Text
Data Set Size	Numerical
Data Set Category	Web text, images, music, video, books
Data Set License Type	Commercial license, Proprietary, Public Domain, Fair Use Claim
Data Set License Name	e.g GPL, Apache, Creative Commons
Data Set Collection Period	Start date, End date (or Present)
Data Set Usage Period	Start date, End date (or Present)
Data Set Contains Personal or Personally Identifiable Information	Yes or No
Personal Information Opt-in Obtained	Yes or No
Personal Information License Mechanism	EULA, Terms of Service, Privacy Policy, Click Through

Documentation required in a Data Declaration (cont'd)

FIELD NAME	POSSIBLE VALUES
Personal Information Anonymized Prior to Training	Yes or No
Data Set Contains Copyrighted Information	Yes or No
License Governing Copyrighted Information	Fair use, Commercial License
Synthetic Training Data Use	Yes or No

Source: Transparency Coalition.AI

Finally, in order to ensure that personal information is protected from unauthorized use by the developers of AI models, the U.S. must invest in its regulatory capacity. Because AI involves a groundbreaking, wide-sweeping set of technologies, an entirely new and robust regulatory system may be needed. The Federal Trade Commission cannot regulate AI alone. The Transparency Coalition has proposed direct government oversight over AI training data and model development.[60] This would involve direct regulatory access to AI training data via API to screen for sensitive personal or copyrighted data, as well as regulatory review of new AI models, which could screen for unauthorized disclosure of personal data by the model.[61] An additional advantage of this approach would be to ensure that training data sets are limited to only data that is necessary to improve the specific model being trained. Due to the widespread privacy issues across the AI industry, it seems clear that the industry cannot police itself. Direct oversight is urgently needed.

> To ensure that personal information is protected from unauthorized use by the developers of AI models, the U.S. must invest in its regulatory capacity. The FTC cannot regulate AI alone.

SECTION V
CONCLUSION

AI development is occurring at a frenetic pace and has many potential positive benefits to society. Unfortunately, privacy protections for personal information have not yet been successfully been built into the AI development process. Steps must be taken to ensure that data privacy becomes a foundational value of AI development moving forward. The current state of judicial and regulatory enforcement against privacy violations are insufficient to address problems arising in the current wave of AI development. To fix these systems and ensure U.S. individuals' personal information is protected, lawmakers, courts and regulators must update their approaches.

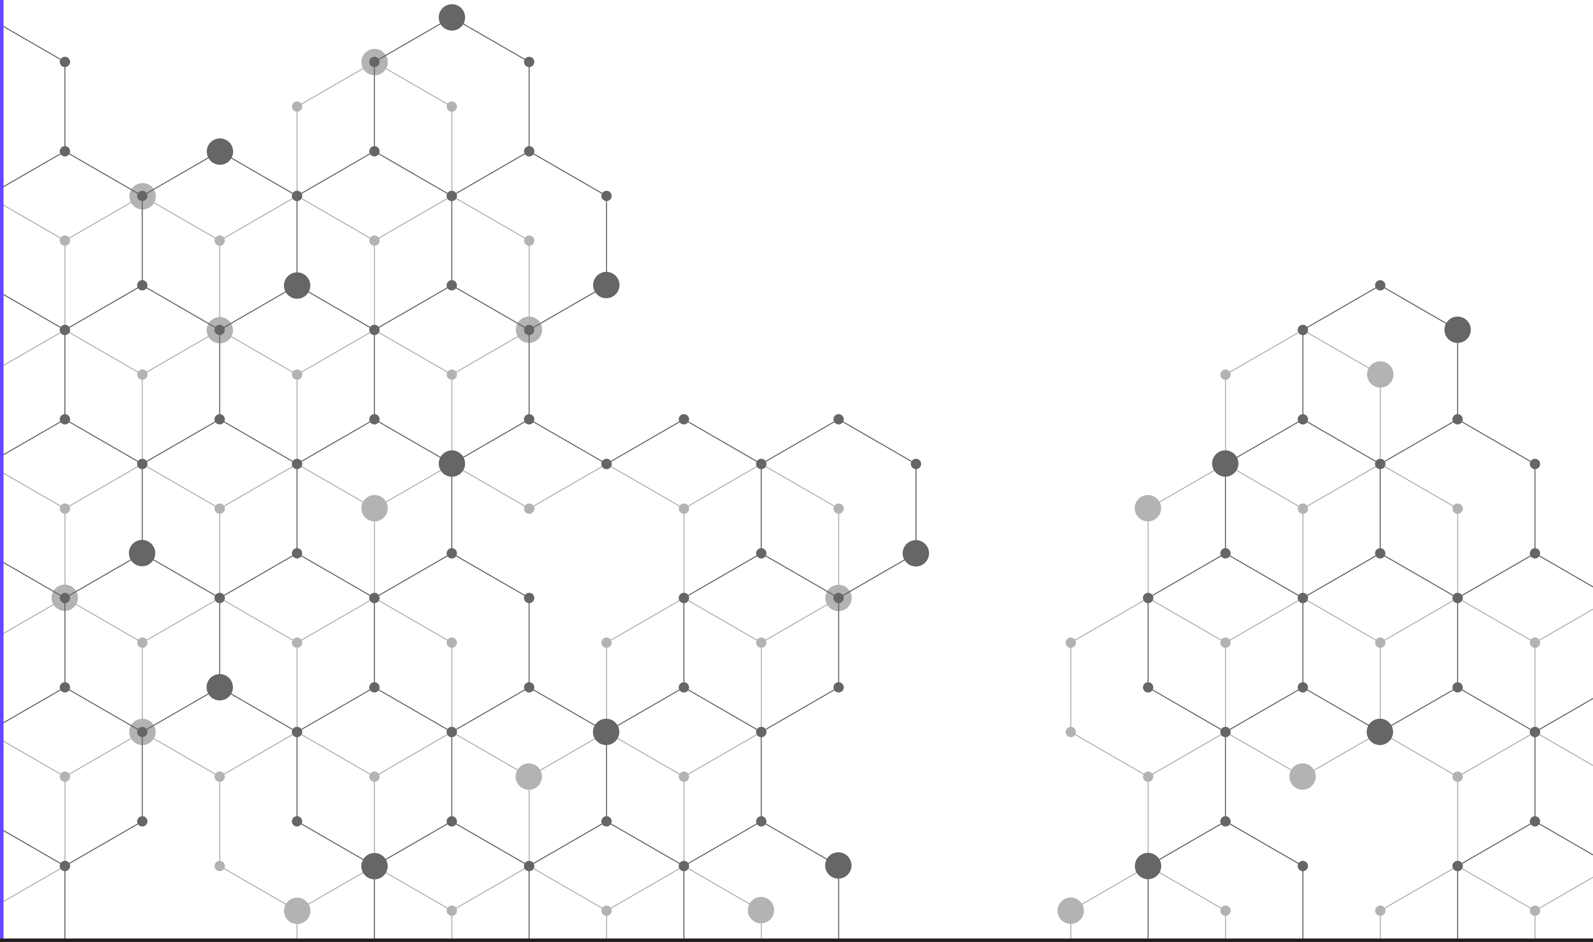

ENDNOTES

[1] Kaye, Byron. “Australian Mayor Readies World’s First Defamation Lawsuit Over ChatGPT Content.” *Reuters*, 5 Apr. 2023, https://www.reuters.com/technology/australian-mayor-readies-worlds-first-defamation-lawsuit-over-chatgpt-content-2023-04-05/. Accessed Feb. 20 2024.

[2] Sato, Mia. “The Unsettling Scourge of Obituary Spam.” *The Verge*, Feb. 12 2024, https://www.theverge.com/24065145/ai-obituary-spam-generative-clickbait. Accessed Feb. 20 2024.

[3] Poritz, Isaiah. "OpenAI Fails to Escape First Defamation Suit From Radio Host." *Bloomberg Law*, Jan. 16, 2024. https://news.bloomberglaw.com/ip-law/openai-fails-to-escape-first-defamation-suit-from-radio-host. Accessed May 4, 2024.

[4] Carlini, Nicholas, et. al. "Extracting Training Data from Diffusion Models." *Proceedings of the 32nd USENIX Security Symposium,* Aug. 9-11, 2023. https://www.usenix.org/system/files/usenixsecurity23-carlini.pdf. Accessed April 24, 2024.

[5] Carlini et. al., "Extracting Training Data from Diffusion Models," at 5254.

[6] Carlini et. al., "Extracting Training Data from Diffusion Models," at 5264-5265.

[7] General Data Protection Regulation, Art. 4(1): “’Personal data’ means any information relating to an identified or identifiable natural person…an identifiable natural person is one who can be identified, directly or indirectly, in particular by reference to an identifier such as a name, an identification number, location data, an online identifier or to one of more factors specific to the physical, physiological, genetic, mental, economic, cultural or social identity of that natural person.” European Parliament and the Council of the European Union. General Data Protection Regulation Art. 4(1). Official Journal of the European Union, 27 Apr. 2016. https://eur-lex.europa.eu/legal-content/EN/ TXT/PDF/?uri=CELEX:32016R0679.

[8] “’Personal information’ means information that identifies, relates to, describes, is reasonably capable of being associated with, or could reasonably be linked, directly or indirectly, with a particular consumer or household.” State of California. California Consumer Privacy Act of 2018. *California Civil Code Sec. 1798.140(v)(1)*, 2018. *California Legislative Information*, https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5.

[9] “’Personal data’ means any information that is linked or reasonably linkable to an identified or identifiable natural person.” Commonwealth of Virginia, Virginia Consumer Data Protection Act. *Code of Virginia Sec. 59.1-575*, 2021. *Virginia’s Legislative Information System*, https://law.lis.virginia.gov/vacodefull/title59.1/chapter53/.

[10] “The reality, [Brown Univerisity professor Suresh] Venkatasubramanian said, is that large language models – the technology underpinning AI tools like ChatGPT – are simply trained to ‘produce a plausible-sounding answer’ to user prompts. ‘So, in that sense, any plausible-sounding answer, whether it’s accurate or factual or made up or not, is a reasonable answer, and that’s what it produces…There is no knowledge of truth there.’” Thorbecke, Catherine. “AI Tools Make Things Up a Lot, and That’s a Huge Problem.” *CNN Business*, 29 Aug. 2023, https://www.cnn.com/2023/08/29/tech/ai-chatbot-hallucinations/index.html. Accessed 15 Mar. 2024.

[11] Prosser, William L. “Privacy.” *California Law Review*, vol. 48, no. 3, Aug. 1960, pp. 383-423, https://lawcat.berkeley.edu/record/1109651?ln=en. Accessed Feb. 20 2024.

[12] Citron, Danielle Keats and Solove, Daniel J. “Privacy Harms,” at 810. *Boston University Law Review*, vol. 102, no. 3, 13 Apr. 2022, pp. 793-863, https://www.bu.edu/bulawreview/files/2022/04/CITRON-SOLOVE.pdf. Accessed 29 Feb. 2024.

[13] “The overall effect of [recent privacy opinions from the United States Supreme Court] has been to drastically limit the enforcement of the Privacy Act through private rights of action. Plaintiffs now have to prove willful conduct as well as establish harm, and they are forbidden from using emotional distress, which is a common type of harm in privacy cases.” Citron and Solove at 812.

[14] *Patrick E. Dwyer v. American Express Co.*, 273 Ill. App. 3d 742 (Ill. App. Ct. 1995).

[15] *Dwyer v. American Express* at 749.

[16] See generally the discussion in Section II, "The Challenges of Privacy Harms," pp. 816-819, in Citron and Solove's *Privacy Harms.*

[17] Citron and Solove, generally.

[18] “Litigation by private parties thus supplements enforcement by regulatory agencies and state attorneys general, and in a number of instances, private litigation serves as the primary enforcement mechanism of a law.” Citron and Solove at 815.

[19] Citron and Solove at 815.

[20] Citron and Solove at 818.

[21] See Citron and Solove’s discussion of deterrence as a goal of enforcement at 819-826.

[22] *TransUnion LLC v. Ramirez*, 594 US_(2021).

[23] See Justice Thomas's dissent. *TransUnion LLC v. Ramirez.*

[24] Winegar, Angela, and Sunstein, Cass R., "How Much Is Data

Worth? A Preliminary Investigation," *Journal of Consumer Policy*, v. 42 (July 1, 2019), p. 425-440.

[25] Prince, Jeffrey, and Wallsten, Scott. "How Much is Privacy Worth Around the World and Across Platforms?" Technology Policy Institute, Jan. 2020. https://techpolicyinstitute.org/wp-content/uploads/2020/02/Prince_Wallsten_How-Much-is-Privacy-Worth-Around-the-World-and-Across-Platforms.pdf

[26] "In AI development, the dominant paradigm is that the more training data, the better." Heikkilä, Melissa. "OpenAI's Hunger for Data Is Coming Back to Bite It." *MIT Technology Review*, MIT Technology Review, 20 Apr. 2023, https://www.technologyreview.com/2023/04/19/1071789/openais-hunger-for-data-is-coming-back-to-bite-it/

[27] "To build large generative AI models, developers turn to the public-facing Internet. …[D]evelopers amass their training sets through automated tools that catalog and extract data from the Internet. Web "crawlers" travel from link to link indexing the location of information in a database, while Web "scrapers" download and extract that same information." Leffer, Lauren. "Your Personal Information Is Probably Being Used to Train Generative AI Models." Scientific American, *Scientific American*, 19 Oct. 2023, https://www.scientificamerican.com/article/your-personal-information-is-probably-being-used-to-train-generative-ai-models/. Accessed Feb. 20 2024.

[28] "A very well-resourced company, such as Google's owner, Alphabet, which already builds Web crawlers to power its search engine, can opt to employ its own tools for the task, says machine learning researcher Jesse Dodge of the nonprofit Allen Institute for AI. Other companies, however, turn to existing resources such as Common Crawl, which helped feed OpenAI's GPT-3, or databases such as the Large-Scale Artificial Intelligence Open Network…" Leffer, "Your Personal Information."

[29] "[I]t is common in the AI industry to build data sets for AI models by scraping the web indiscriminately and then outsourcing the work of removing duplicates or irrelevant data points, filtering unwanted things, and fixing typos. These methods, and the sheer size of the data set, mean tech companies tend to have a very limited understanding of what has gone into training their models. …Tech companies don't document how they collect or annotate AI training data and don't even tend to know what's in the data set, says Nithya Sambasivan, a former research scientist at Google…"Heikkilä, "OpenAI's Hunger for Data."

[30] "Importantly, companies' use of personal data scraped from the web, including for the training of AI technologies, can undermine consumers' trust and thereby have detrimental consequences for the digital economy." Fazlioglu, Müge. "Training AI on Personal Data Scraped from the Web." International Association of Privacy Professionals, Nov. 8 2023, https://iapp.org/news/a/training-ai-on-personal-data-scraped-from-the-web/. Accessed Feb. 20 2024.

[31] "For example, deidentified or anonymized data can be made identifiable when combined with other deidentified data and processed by a machine-learning algorithm. Output data unpredictably may contain identifiable health information." Bennet, Barbara and Matta, Neha M. "Beware Privacy Risks In Training AI Models with Health Data." Frost Brown Todd Attorneys, Nov. 9 2023, https://frostbrowntodd.com/beware-privacy-risks-in-training-ai-models-with-health-data-3/. Accessed Feb. 20 2024.

[32] See generally the section "Where do AI training data come from?" in Leffer, "Your Personal Information."

[33] "People tend to share intimate, private information with the chatbot, telling it about things like their mental state, their health, or their personal opinions. Leautier says it is problematic if there's a risk that ChatGPT regurgitates this sensitive data to others." Heikkilä, "OpenAI's Hunger for Data."

[34] See generally Heikkilä, "OpenAI's Hunger for Data;" and Bennet and Matta, "Beware Privacy Risks."

[35] "AI models can regurgitate the same material that was used to train them—including sensitive personal data and copyrighted work. Many widely used generative AI models have blocks meant to prevent them from sharing identifying information about individuals, but researchers have repeatedly demonstrated ways to get around these restrictions." Leffer, "Your Personal Information."

[36] "For example, deidentified or anonymized data can be made identifiable when combined with other deidentified data and processed by a machine-learning algorithm." Bennet and Matta, "Beware Privacy Risks."

[37] "The reality, [Brown Univerisity Professor Suresh] Venkatasubramanian said, is that large language models – the technology underpinning AI tools like ChatGPT – are simply trained to 'produce a plausible-sounding answer' to user prompts. 'So, in that sense, any plausible-sounding answer, whether it's accurate or factual or made up or not, is a reasonable answer, and that's what it produces…There is no knowledge of truth there.'" Thorbecke, Catherine. "AI Tools Make Things Up a Lot, and That's a Huge Problem." *CNN Business*, 29 Aug. 2023, https://www.cnn.com/2023/08/29/tech/ai-chatbot-hallucinations/index.html. Accessed 15 Mar. 2024.

[38] The CTO of AI startup Vectera, Amin Ahmad, "says that LLMs create a compressed representation of all the training data fed through its artificial neurons. 'The nature of compression is that the fine details can get lost,' he says. A model ends up primed with the most likely answers to queries from users but doesn't have the exact facts at its disposal. 'When it gets to the details it starts making things up,' he says." Levy, Steven. "In Defense of AI Hallucinations." *Wired*, 5 January 2024, *Wired.com*, https://www.wired.com/story/plaintext-in-defense-of-ai-hallucinations-chatgpt/. Accessed 15 Mar. 2024.

[39] Levy, "In Defense of AI Hallucinations."

[40] Complaint: *European Center for Digital Rights v. OpenAI OpCo, LLC*, https://noyb.eu/sites/default/files/2024-04/OpenAI%20Complaint_EN_redacted.pdf.

[41] General Data Protection Regulation, Art. 16 and 17.

[42] Complaint: *European Center for Digital Rights v. OpenAI.*

[43] "Video conferencing company Zoom encountered a massive backlash over concerns the contents of video chat might be used to train AI systems. The move prompted an apologetic post from Zoom's CEO, but the company is far from alone in seeking more consumer data in order to train AI models." Heath, Ryan and Fried, Ina. "Terms-of-Service Land Grab: Tech Firms Seek Private Data to Train AI." *Axios*, 18 Aug. 2023, https://www.axios.com/2023/08/18/ai-legal-user-data. Accessed 20 Feb. 2024.

See also Fazlioglu, "Training AI on Personal Data," describing class action lawsuit claiming "stolen" personal data for "unjust enrichment" via AI training.

[44] Heath and Fried, "Terms-of-Service Land Grab."

[45] "The FTC has recognized that trade practices that prevent consumers from 'effectively making their own decisions' are ones that cause substantial injury." Citron and Solove at 848.

[46] *FTC and DOJ Charge Amazon with Violating Children's Privacy Law by Keeping Kids' Alexa Voice Recordings Forever and Undermining Parents' Deletion Requests.* Federal Trade Commission, 31 May 2023, https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever. Accessed 11 Mar. 2024.

[47] According to the FTC complaint, "Amazon also failed for a significant period of time to honor parents' requests that it delete their children's voice recordings by continuing to retain the transcripts of those recordings and failing to disclose that it was doing so, also in violation of COPPA." These transcripts were retained to "improve Alexa's speech recognition and processing capabilities." United States, Complaint filed in District Court for the Western District of Washington. *United States of America v. Amazon.com, Inc. and Amazon.com Services LLC.* Filed 31 May 2023. *Federal Trade Commission*, https://www.ftc.gov/news-events/news/press-releases/2023/05/ftc-doj-charge-amazon-violating-childrens-privacy-law-keeping-kids-alexa-voice-recordings-forever, at 3 and 7. Accessed 11 Mar. 2024.

[48] *U.S. v. Amazon* complaint at 6-7.

[49] Federal Trade Commission, *FTC and DOJ Charge Amazon.*

[50] Citron and Solove at 816 and 818.

[51] CCPA, for example, includes the following right to delete: "A consumer shall have the right to request that a business delete any personal information about the consumer which the business has collected from the consumer." California Consumer Privacy Act Sec. 1798.105.

[52] "OpenAI is going to find it near-impossible to identify individuals' data and remove it from its models, says Margaret Mitchell, an AI researcher and chief ethics scientist at startup Hugging Face, who was formerly Google's AI ethics co-lead." Heikkilä, "OpenAI's Hunger for Data."

See also Federal Trade Commission, *"FTC and DOJ Charge Amazon."*

[53] United States, *U.S. v. Amazon* complaint at 8.

[54] Describing the difficulty of removing information a large language model has been trained on from that system. Burgess, "ChatGPT Has a Big Privacy Problem."

[55] "In California and a handful of other states, recently passed digital privacy laws give consumers the right to request that companies delete their data. In the European Union, too, people have the right to data deletion. So far, however, AI companies have pushed back on such requests by claiming the provenance of the data can't be proven—or by ignoring the requests altogether—says Jennifer King, a privacy and data researcher at Stanford University." Leffer, "Your Personal Information."

[56] Citron and Solove at 848 and 852.

[57] U.S. Federal Trade Commission case proceedings. "BetterHelp, Inc., In the Matter of," July 14, 2023. https://www.ftc.gov/legal-library/browse/cases-proceedings/2023169-betterhelp-inc-matter.

[58] Similar to the data cards described in this article: Pushkarna, Mahima, Zaldivar, Andrew, and Kjartansson, Oddur. "Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI." *ACM Digital Library*, June 2022, https://dl.acm.org/doi/fullHtml/10.1145/3531146.3533231. Accessed 14 Mar. 2024.

[59] Gebru, Timnit, et al. *Datasheets for Datasets.* December 1, 2021. https://arxiv.org/pdf/1803.09010. Accessed May 17, 2024.

Mahima Pushkarna, Andrew Zaldivar, and Oddur Kjartansson. 2022. *Data Cards: Purposeful and Transparent Dataset Documentation for Responsible AI.* In 2022 ACM Conference on Fairness, Accountability, and Transparency (FAccT '22), June 21–24, 2022, Seoul, Republic of Korea. Association for Computing Machinery. https://doi.org/10.1145/3531146.3533231. Accessed May 17, 2024.

[60] "Regulators must have API access to review training data inputs, where the API provides [access] to an outline of files incorporated and/or personal data processed in AI model training. … In order to review and approve generative AI projects, an AI model clearinghouse is needed…where large AI projects are reviewed for approval…" Transparency Coalition.ai Association. "Solutions." *Transparency Coalition.ai*, https://www.transparencycoalition.ai/solutions. Accessed 20 Feb. 2024.

[61] Transparency Coalition.ai Association, "Solutions."