

# BROOKINGS

## RESEARCH

# Protecting privacy in an AI-driven world

February 10, 2020

**Editor's note:** This report from The Brookings Institution's Artificial Intelligence and Emerging Technology (AIET) Initiative is part of "[AI Governance](https://www.brookings.edu/series/ai-governance/) (<https://www.brookings.edu/series/ai-governance/>)," a series that identifies key governance and norm issues related to AI and proposes policy remedies to address the complex challenges associated with emerging technologies.

Our world is undergoing an information Big Bang, in which the universe of data doubles every two years and quintillions of bytes of data are generated every day.<sup>1</sup> For decades, Moore's Law on the doubling of computing power every 18-24 months has driven the growth of information technology. Now—as billions of smartphones and other devices collect and transmit data over high-speed global networks, store data in ever-larger data centers, and analyze it using increasingly powerful and sophisticated software—Metcalf's Law comes into play. It treats the value of networks as a function of the square of the number of nodes, meaning that network effects exponentially compound this historical growth in information. As 5G networks and eventually quantum computing deploy, this data explosion will grow even faster and bigger.

The impact of big data is commonly described in terms of three "Vs": volume, variety, and velocity.<sup>2</sup> More data makes analysis more powerful and more granular. Variety adds to this power and enables new and unanticipated inferences and predictions. And velocity facilitates analysis as well as sharing in real time. Streams of data from mobile phones and other online devices expand the volume, variety, and velocity of

information about every facet of our lives and puts privacy into the spotlight as a global public policy issue.

Artificial intelligence likely will accelerate this trend. Much of the most privacy-sensitive data analysis today—such as search algorithms, recommendation engines, and adtech networks—are driven by machine learning and decisions by algorithms. As artificial intelligence evolves, it magnifies the ability to use personal information in ways that can intrude on privacy interests by raising analysis of personal information to new levels of power and speed.

“As artificial intelligence evolves, it magnifies the ability to use personal information in ways that can intrude on privacy interests by raising analysis of personal information to new levels of power and speed.”

Facial recognition systems offer a preview of the privacy issues that emerge. With the benefit of rich databases of digital photographs available via social media, websites, driver's license registries, surveillance cameras, and many other sources, machine recognition of faces has progressed rapidly from fuzzy images of cats<sup>3</sup> to rapid (though still imperfect) recognition of individual humans. Facial recognition systems are being deployed in cities and airports around America. However, China's use of facial recognition as a tool of authoritarian control in Xinjiang<sup>4</sup> and elsewhere has awakened opposition to this expansion and calls for a ban on the use of facial recognition. Owing to concerns over facial recognition, the cities of Oakland, Berkeley, and San Francisco in California, as well as Brookline, Cambridge, Northampton, and Somerville in Massachusetts, have adopted bans on the technology.<sup>5</sup> California, New Hampshire, and Oregon all have enacted legislation banning use of facial recognition with police body cameras.<sup>6</sup>

This policy brief explores the intersection between AI and the current privacy debate. As Congress considers comprehensive privacy legislation to fill growing gaps in the current checkerboard of federal and state privacy, it will need to consider if or how to address use personal information in artificial intelligence systems. In this brief, I discuss some potential concerns regarding artificial intelligence and privacy, including discrimination, ethical use, and human control, as well as the policy options under discussion.

# Privacy issues in AI

The challenge for Congress is to pass privacy legislation that protects individuals against any adverse effects from the use of personal information in AI, but without unduly restricting AI development or ensnaring privacy legislation in complex social and political thickets. The discussion of AI in the context of the privacy debate often brings up the limitations and failures of AI systems, such as predictive policing that could disproportionately affect minorities<sup>7</sup> or Amazon's failed experiment with a hiring algorithm that replicated the company's existing disproportionately male workforce.<sup>8</sup> These both raise significant issues, but privacy legislation is complicated enough even without packing in all the social and political issues that can arise from uses of information. To evaluate the effect of AI on privacy, it is necessary to distinguish between data issues that are endemic to all AI, like the incidence of false positives and negatives or overfitting to patterns, and those that are specific to use of personal information.

The privacy legislative proposals that involve these issues do not address artificial intelligence in name. Rather, they refer to “automated decisions” (borrowed from EU data protection law) or “algorithmic decisions” (used in this discussion). This language shifts people's focus from the use of AI as such to the use of *personal data* in AI and to the impact this use may have on individuals. This debate centers in particular on algorithmic bias and the potential for algorithms to produce unlawful or undesired discrimination in the decisions to which the algorithms relate. These are major concerns for civil rights and consumer organizations that represent populations that suffer undue discrimination.

Addressing algorithmic discrimination presents basic questions about the scope of privacy legislation. First, to what extent can or should legislation address issues of algorithmic bias? Discrimination is not self-evidently a privacy issue, since it presents broad social issues that persist even without the collection and use of personal information, and fall under the domain of various civil rights laws. Moreover, making these laws available for debate could effectively open a Pandora's Box because of the charged political issues they touch on and the multiple congressional committees with jurisdiction over various such issues. Even so, discrimination is based on personal attributes such as skin color, sexual identity, and national origin. Use of personal information about these attributes, either explicitly or—more likely and less obviously

—via proxies, for automated decision-making that is against the interests of the individual involved thus implicates privacy interests in controlling how information is used.

“This charade of consent has made it obvious that notice-and-choice has become meaningless. For many AI applications ... it will become utterly impossible.”

Second, protecting such privacy interests in the context of AI will require a change in the paradigm of privacy regulation. Most existing privacy laws, as well as current Federal Trade Commission enforcement against unfair and deceptive practices, are rooted in a model of consumer choice based on “notice-and-choice” (also referred to as “notice-and-consent”). Consumers encounter this approach in the barrage of notifications and banners linked to lengthy and uninformative privacy policies and terms and conditions that we ostensibly consent to but seldom read. This charade of consent has made it obvious that notice-and-choice has become meaningless. For many AI applications—smart traffic signals and other sensors needed to support self-driving cars as one prominent example—it will become utterly impossible.

Although almost all bills on Capitol Hill still rely on the notice-and-choice model in some degree, key congressional leaders as well as privacy stakeholders have expressed desire to change this model by shifting the burden of protecting individual privacy from consumers over to the businesses that collect data.<sup>9</sup> In place of consumer choice, their model focuses on business conduct by regulating companies’ processing of data—what they collect and how they can use it and share it. Addressing data processing that results in any algorithmic discrimination can fit within this model.

A model focused on data collection and processing may affect AI and algorithmic discrimination in several ways:

- Data stewardship requirements, such as duties of fairness or loyalty, could militate against uses of personal information that are adverse or unfair to the individuals the data relates to.
- Data transparency or disclosure rules, as well as rights of individuals to access information relating to them, could illuminate uses of algorithmic decision-making.

- Data governance rules that prescribe the appointment of privacy officers, conduct of privacy impact assessments, or product planning through “privacy by design” may surface issues concerning use of algorithms.
- Rules on data collection and sharing could reduce the aggregation of data that enables inferences and predictions, but may involve some trade-offs with the benefits of large and diverse datasets.

In addition to these provisions of general applicability that may affect algorithmic decisions indirectly, a number of proposals specifically address the subject.<sup>10</sup>

## AI policy options for privacy protection

The responses to AI that are currently under discussion in privacy legislation take two main forms. The first targets discrimination directly. A group of 26 civil rights and consumer organizations wrote a joint letter advocating to prohibit or monitor use of personal information with discriminatory impacts on “people of color, women, religious minorities, members of the LGBTQ+ community, persons with disabilities, persons living on I winsome, immigrants, and other vulnerable populations.”<sup>11</sup> The Lawyers’ Committee for Civil Rights Under Law and Free Press Action have incorporated this principle into model legislation aimed at data discrimination affecting economic opportunity, public accommodations, or voter suppression.<sup>12</sup> This model is substantially reflected in the Consumer Online Privacy Rights Act, which was introduced in the waning days of the 2019 congressional session by Senate Commerce Committee ranking member Maria Cantwell (D-Wash.). It also includes a similar provision restricting the processing of personal information that discriminates against or classifies individuals on the basis of protected attributes such race, gender, or sexual orientation.<sup>13</sup> The Republican draft counterproposal addresses the potential for discriminatory use of personal information by calling on the Federal Trade Commission to cooperate with agencies that enforce discrimination laws and to conduct a study.<sup>14</sup>

This approach to algorithmic discrimination implicates debates over private rights of action in privacy legislation. The possibility of such individual litigation is a key point of divergence between Democrats aligned with consumer and privacy advocates on one hand, and Republicans aligned with business interests on the other. The former argue

that private lawsuits are a needed force multiplier for federal and state enforcement, while the latter express concern that class action lawsuits, in particular, burden business with litigation over trivial issues. In the case of many of the kinds of discrimination enumerated in algorithmic discrimination proposals, existing federal, state, and local civil rights laws enable individuals to bring claims for discrimination. Any federal preemption or limitation on private rights of action in federal privacy legislation should not impair these laws.

The second approach addresses risk more obliquely, with accountability measures designed to identify discrimination in the processing of personal data. Numerous organizations and companies as well as several legislators propose such accountability. Their proposals take various forms:

- **Transparency:** This refers to disclosures relating to uses of algorithmic decision-making. While lengthy, detailed privacy policies are not helpful to most consumers, they do provide regulators and other privacy watchdogs with a benchmark by which to examine a company's data handling and hold that company accountable. Replacing current privacy policies with "privacy disclosures" that require a complete description of what and how data is collected, used, and protected would enhance this benchmark function. In turn, requiring that these disclosures identify significant uses of personal information for algorithmic decisions would help watchdogs and consumers know where to look out for untoward outcomes.<sup>15</sup>
- **Explainability:** While transparency provides advance notice of algorithmic decision-making, explainability involves retroactive information about the use of algorithms in specific decisions. This is the main approach taken in the European Union's General Data Protection Regulation (GDPR). The GDPR requires that, for any automated decision with "legal effects or similarly significant effects" such as employment, credit, or insurance coverage, the person affected has recourse to a human who can review the decision and explain its logic.<sup>16</sup> This incorporates a "human-in-the-loop" component and an element of due process that provide a check on anomalous or unfair outcomes.

A sense of fairness suggests such a safety valve should be available for algorithmic decisions that have a material impact on individuals' lives. Explainability requires (1) identifying algorithmic decisions, (2) deconstructing specific decisions, and (3) establishing a channel by which an individual can seek an explanation. Reverse-engineering algorithms based on machine learning can be difficult, and even

impossible, a difficulty that increases as machine learning becomes more sophisticated. Explainability therefore entails a significant regulatory burden and constraint on use of algorithmic decision-making and, in this light, should be concentrated in its application, as the EU has done (at least in principle) with its “legal effects or similarly significant effects” threshold. As understanding increases about the comparative strengths of human and machine capabilities, having a “human in the loop” for decisions that affect people’s lives offers a way to combine the power of machines with human judgment and empathy.

- **Risk assessment:** In the 1974 Privacy Act, risk assessments were originally developed as “privacy impact assessments” within the federal government. They have since evolved as widely used privacy-management tools to evaluate and mitigate privacy risks in advance, and are required by the GDPR for novel technology or high-risk uses of data. Proposals for privacy legislation from Sen. Ron Wyden (D-Ore.)<sup>17</sup> and Intel Corporation<sup>18</sup> would require that *any* automated decision-making be preceded by an assessment of its risks to individuals. Wyden has also filed a separate, stand-alone bill on algorithmic decision-making, the Algorithmic Accountability Act.<sup>19</sup> Risk assessments for algorithmic decision-making provide an opportunity to anticipate potential biases in design and data as well as the potential impact on individuals. For the regulatory burden to be proportionate, the level of risk assessment should be appropriate to the significance of the decision-making in question, which depends on the consequences of the decisions, the number of people and volume of data potentially affected, and the novelty and complexity of algorithmic processing.
- **Audits:** Audits evaluate privacy practices retrospectively. Most legislative proposals contain some general accountability requirements to ensure companies comply with their privacy programs, and some include self-audits or third-party audits. Paired with proactive risk assessments, auditing outcomes of algorithmic decision-making can help match foresight with hindsight; although, like explainability, auditing machine-learning routines is difficult and still developing. One of the clear lessons from the AI debate, as summarized in a review of best practices by Brookings scholar Nicol Turner Lee with Paul Resnick and Genie Barton, is that “it’s important for algorithm operators and developers to always be asking themselves: *Will we leave some groups of people worse off as a result of the algorithm’s design or its unintended consequences?*” (emphasis in original).<sup>20</sup>

Because of the difficulties of foreseeing machine learning outcomes as well as reverse-engineering algorithmic decisions, no single measure can be completely effective in avoiding perverse effects. Thus, where algorithmic decisions are consequential, it makes sense to combine measures to work together. Advance measures such as transparency and risk assessment, combined with the retrospective checks of audits and human review of decisions, could help identify and address unfair results. A combination of these measures can complement each other and add up to more than the sum of the parts. Risk assessments, transparency, explainability, and audits also would strengthen existing remedies for actionable discrimination by providing documentary evidence that could be used in litigation. Not all algorithmic decision-making is consequential, however, so these requirements should vary according to the objective risk.

## Looking ahead

The window for this Congress to pass comprehensive privacy legislation is narrowing. While the Commerce Committee in each house of Congress has been working on a bipartisan basis throughout 2019 and have put out discussion drafts, they have yet to reach agreement on a bill. Meanwhile, the California Consumer Privacy Act went into effect on Jan. 1, 2020,<sup>21</sup> impeachment and war powers have crowded out other issues, and the presidential election is going into full swing.

“The window for this Congress to pass comprehensive privacy legislation is narrowing.”

In whatever window remains to pass privacy legislation before the 2020 election, the treatment of algorithmic decision-making is a substantively and politically challenging issue that will need a workable resolution. For a number of civil rights, consumer, and other civil society groups, establishing protections against discriminatory algorithmic decision-making is an essential part of legislation. In turn, it will be important to Democrats in Congress. At a minimum, some affirmation that algorithmic discrimination based on personal information is subject to existing civil rights and nondiscrimination laws, as well as some additional accountability measures, will be essential to the passage of privacy legislation.

*The Brookings Institution is a nonprofit organization devoted to independent research and policy solutions. Its mission is to conduct high-quality, independent research and, based on that research, to provide innovative, practical recommendations for policymakers and the public. The conclusions and recommendations of any Brookings publication are solely those of its author(s), and do not reflect the views of the Institution, its management, or its other scholars.*

*Microsoft provides support to The Brookings Institution's [Artificial Intelligence and Emerging Technology \(AIET\) Initiative \(https://www.brookings.edu/project/artificial-intelligence-and-emerging-technology-initiative/\)](https://www.brookings.edu/project/artificial-intelligence-and-emerging-technology-initiative/), and Amazon and Intel provide general, unrestricted support to the Institution. The findings, interpretations, and conclusions in this report are not influenced by any donation. Brookings recognizes that the value it provides is in its absolute commitment to quality, independence, and impact. Activities supported by its donors reflect this commitment.*

## AUTHORS



**Cameron F. Kerry** Ann R. and Andrew H. Tisch Distinguished Visiting Fellow - Governance Studies, Center for Technology Innovation  @Cam\_Kerry

---

## Footnotes

1. Bernard Marr, "How Much Data Do We Create Every Day? The Mind-Blowing Stats Everyone Should Read," *Forbes*, May 21, 2018. A quintillion is a 1 followed 30 zeroes.
2. "NIST Big Data Interoperability Framework: Volume 1, Definitions," September 2015. <http://dx.doi.org/10.6028/NIST.SP.1500-1>
3. <https://web.stanford.edu/class/aerchive/cs/cs106a.1188/lectures/lecture26.pdf>
4. Paul Mozur, "One Month, 500,000 Face Scans: How China Is Using A.I. to Profile a Minority," *New York Times*, April 14, 2019.

5. Sara Merken, "Berkeley Bans Government Face Recognition Use, Joining Other Cities," Bloomberg Law, October 16, 2019; Nikolas DeCosta-Klipa, "Brookline becomes 2nd Massachusetts community to ban facial recognition," Boston.com, December 12, 2019. Tori Bedford, "Cambridge Votes To Ban Face Surveillance Technology," WGBH, January 13, 2020.
6. <https://securitytoday.com/articles/2019/10/10/california-to-become-third-state-to-ban-facial-recognition-software-in-police-body-cameras.aspx>
7. Matthias Spielkamp, "Inspecting Algorithms for Bias," MIT Technology Review, June 12, 2017.
8. Jeffrey Dastin, "Amazon scraps secret AI recruiting tool that showed bias against women," Reuters, October 9, 2018
9. Pallone statement at "Protecting Consumer Privacy in the Era of Big Data" hearing, House Subcommittee on Consumer Protection and Commerce, February 26, 2019. Wicker statement at "Policy Principles for a Federal Data Privacy Framework in the United States" hearing, Senate Commerce, Science, & Transportation Committee, February 27, 2019. See: principles for privacy legislation joined by key Democratic leaders on privacy in the Senate called for putting the burden of protecting privacy on the companies that collection personal information:  
[https://www.democrats.senate.gov/imo/media/doc/Final\\_CMTE%20Privacy%20Principles\\_11.14.19.pdf](https://www.democrats.senate.gov/imo/media/doc/Final_CMTE%20Privacy%20Principles_11.14.19.pdf) ↗
10. See Intel's draft privacy legislation: <https://usprivacybill.intel.com/wp-content/uploads/IntelPrivacyBill-05-25-19.pdf> ↗; Center for Democracy and Technology's draft privacy legislation: <https://cdt.org/files/2018/12/2018-12-12-CDT-Privacy-Discussion-Draft-Final.pdf> ↗; and Senator Wyden's discussion draft: <https://www.wyden.senate.gov/news/press-releases/wyden-releases-discussion-draft-of-legislation-to-provide-real-protections-for-americans-privacy> ↗
11. <https://lawyerscommittee.org/civil-rights-civil-liberties-and-consumer-groups-urge-congress-to-protect-marginalized-communities-from-discriminatory-privacy-abuses/> ↗
12. Stanley Augustin, "Lawyers' Committee for Civil Rights Under Law and Free Press Action Release Proposed 'Online Civil Rights and Privacy Act' to Combat Data Discrimination," Lawyers' Committee for Civil Rights Under Law, March 11, 2019.
13. See S. 2968: <https://www.congress.gov/bill/116th-congress/senate-bill/2968?q=%7B%22search%22%3A%5B%22consumer+online+privacy+rights+act%22%5D%7D&s=1&r=1> ↗
14. See press release: <https://www.commerce.senate.gov/2019/12/chairman-wicker-s-discussion-draft-the-united-states-consumer-data-privacy-act> ↗
15. Cameron F. Kerry and Caitlin Chin, "Hitting refresh on privacy policies: Recommendations for notice and transparency," The Brookings Institution, January 6, 2020.

16. See the General Data Protection Regulation official text: <https://gdpr-info.eu/> ↗
17. “Wyden Releases Discussion Draft of Legislation to Provide Real Protections for Americans’ Privacy,” Ron Wyden Press Release, November 1, 2018.
18. See Intel’s draft privacy legislation: <https://usprivacybill.intel.com/wp-content/uploads/IntelPrivacyBill-05-25-19.pdf> ↗
19. See S. 1108 – Algorithmic Accountability Act of 2019: <https://www.congress.gov/bill/116th-congress/senate-bill/1108?q=%7B%22search%22%3A%5B%22algorithmic+accountability+act%22%5D%7D&s=1&r=2> ↗
20. Nicol Turner Lee, Paul Resnick, and Genie Barton, “Algorithmic bias detection and mitigation: Best practices and policies to reduce consumer harms,” The Brookings Institution, May 22, 2019.
21. John Stephens, “California Consumer Privacy Act,” American Bar Association, July 2, 2019.

---

Copyright 2024 The Brookings Institution