R Street
Free markets. Real solutions.

R Street Policy Study
No. 317
March 2025

Leveraging AI and Emerging Technology to Enhance Data Privacy and Security

By Steven Ward

This paper explores how AI can enhance data protection by facilitating and improving privacy-enhancing technology, data impact assessments, data minimization, and data security.

Table of Contents

Executive Summary

With global data breaches compromising billions of records and average data-breach costs significantly increasing year over year, organizations face mounting challenges in protecting data. This paper explores how AI can enhance data protection by facilitating and improving privacy-enhancing technology, data impact assessments, data minimization, and data security. It analyzes specific AI-driven applications, like differential privacy, federated learning, and homomorphic encryption, which enable organizations to extract valuable insights while maintaining strong privacy guarantees. It also examines AI's potential to improve privacy and consent frameworks by demystifying complicated privacy notices and increasing consumer awareness and transparency. The study concludes with policy recommendations emphasizing the need for a comprehensive federal privacy law, flexible AI governance that promotes innovation, and continued support for PET development through public–private partnerships. This analysis arrives at a crucial time as policymakers and organizations continue to seek an optimal balance between fully leveraging AI's capabilities and maintaining strong privacy protections in an increasingly complex technological landscape.

Introduction

A reported 66 percent of the global population uses the internet, and this percentage increases every year.[1] Internet users generate data in everything they do—from internet

1. Simon Kemp, "Internet Use in 2024," DataReportal, Jan. 31, 2024. https://datareportal.com/reports/digital-2024-deep-dive-the-state-of-internet-adoption.

browsing and television streaming to driving cars and running vacuums. On average, each user generates 65 gigabytes of data per day, which is collected, processed, and incorporated into countless business processes, such as operations and delivery services; customer experience improvements; research and development; and marketing.[2]

On average, each user generates 65 gigabytes of data per day, which is collected, processed, and incorporated into countless business processes, such as operations and delivery services; customer experience improvements; research and development; and marketing.

With so much data generation, data breaches and privacy violations have become increasingly prevalent, and organizations face mounting challenges in protecting individual data. Globally, these breaches compromised 17 billion personal records in 2023 alone, and the average cost per breach reached $4.88 million per incident.[3] This is an urgent issue because compromised data can be used for a host of nefarious purposes, such as tracking and surveilling sensitive populations like military and intelligence community members, blackmailing individuals, carrying out more effective cyber incidents, and spreading disinformation, including during warfare.[4] National policy documents, such as the National Cybersecurity Strategy and the Annual Threat Assessment, have highlighted these risks and underscore the need to improve data-security strategies.[5]

With their static policies and manual oversight, legacy privacy-protection methods have struggled to adapt to evolving threats and an ever-increasing data ecosystem. Fortunately, with its scalability and ability to learn patterns, predict potential vulnerabilities, and respond in real time to emerging threats, AI technology offers promising solutions to these limitations. Recent advances in machine learning—particularly in differential privacy, federated learning, and homomorphic encryption—are helping organizations extract valuable insights from sensitive data and enhance data utility while maintaining mathematical privacy guarantees, rather than relying on access controls and anonymization.[6] In addition, AI-driven anomaly-detection systems have demonstrated a remarkable capability for identifying potential privacy breaches before they occur, offering a proactive, rather than reactive, approach to data protection.

This paper explores the potential of AI and emerging technology to enhance how organizations safeguard data. It provides examples of how AI-based technical and policy solutions can effectively enhance data privacy and security components across several overlapping data privacy and security dimensions: data privacy impact assessments (DPIAs), data minimization, data security, privacy-enhancing technology (PETs), and consumer awareness. Our intent is that the information presented herein will help policymakers and organizations assess potential legislative measures and consider how to approach and leverage AI innovations.

Privacy-Enhancing Technologies

Privacy-enhancing technologies, or PETs, are specialized technological solutions that allow organizations to derive value and utilization from sensitive data while maintaining strong

2. Nina Quist, "How Much Data is Created Every Day in 2025?," Business2Community, June 27, 2024. https://www.business2community.com/statistics-pages/how-much-data-is-created-every-day.
3. James Coker, "7 Billion Personal Records Exposed in Data Breaches in 2023," *Infosecurity Magazine*, March 28, 2024. https://www.infosecurity-magazine.com/news/personal-records-exposed-data; "Cost of a Data Breach: Report 2024," IBM, last accessed Jan. 1, 2025. https://www.ibm.com/reports/data-breach.
4. Greg Lindsay et al., "Microtargeting Unmasked: Safeguarding Law Enforcement, the Military, and the Nation in the Era of Personalized Threats," U.S. Secret Service, August 2023. https://www.secretservice.gov/sites/default/files/reports/2023-08/asu-tc-micro-targeting-report_final.pdf; Jessica Dawson and Brandon Pugh, "Ukraine conflict heightens US military's data privacy vulnerabilities," C4ISRNET, April 14, 2022. www.c4isrnet.com/opinion/2022/04/14/ukraine-conflict-heightens-us-militarys-data-privacy-vulnerabilities.
5. "Fact Sheet: Biden-Harris Administration Announces National Cybersecurity Strategy," Archives, The White House, March 2, 2023. https://bidenwhitehouse.archives.gov/oncd/briefing-room/2024/05/07/fact-sheet-ncsip-version-2; "Annual Threat Assessment of the U.S. Intelligence Community," Office of the Director of National Intelligence, Feb. 6, 2023. https://www.dni.gov/files/ODNI/documents/assessments/ATA-2023-Unclassified-Report.pdf.
6. Andrea Gadotti et al., "Anonymization: The imperfect science of using data while preserving privacy," *Science Advances* 10:29 (July 17, 2024). https://www.science.org/doi/10.1126/sciadv.adn7053.

privacy protections.[7] AI is poised to significantly improve PETs by introducing advanced capabilities in anonymization, secure computation, and data analysis. PETs were encouraged by the previous administration.[8]

AI-Enhanced Data Anonymization

Traditional data anonymization techniques often struggle to balance privacy and utility effectively.[9] The risk of re-identification (de-anonymizing data) through attacks such as linkage, inference, and reconstruction is a valid concern.[10] For example, in 1997, a computer scientist re-identified the medical records of Massachusetts' then-Governor William Weld by cross-referencing anonymized medical insurance records with publicly available voter registration information.[11] Soon after doing so, the scientist testified in front of Congress, and her research helped guide the Health Insurance Portability and Accountability Act's (HIPAA's) re-identification provisions.[12] Although data anonymization has improved since then, the threat of re-identifying data still lingers.

AI-driven algorithms can help reduce the threat of re-identifying data after anonymization because they are able to dynamically identify and mask sensitive information in datasets, preserving privacy while retaining analytical value.

AI-driven algorithms can help reduce this threat because they are able to dynamically identify and mask sensitive information in datasets, preserving privacy while retaining analytical value. For example, generative adversarial networks (GANs) can create synthetic datasets that mimic real data, ensuring privacy without compromising usability.[13] In the re-identification of Gov. Weld's medical records in the example above, the direct identifiers were removed, but demographic data was retained, which created a unique identifier when combined with voter records. Instead of removing identifiers like social security numbers, names, and date of birth from the records, a GAN system would have generative synthetic patient records that would not have any one-to-one relationships with such identifiers, making re-identification extremely difficult, if not impossible.

AI-Enhanced Differential Privacy

Differential privacy (DP) is a mathematical framework for protecting individual privacy while enabling the analysis and sharing of aggregate data introduced in 2006.[14] In the nearly two decades since it was introduced, DP has significantly evolved to provide a quantifiable approach to balancing data utility with privacy protection guarantees.[15] For example, in the 2024-2025 basketball season, the Sacramento Kings had a horrible 3-point shooting percentage: 33.9.[16] If the coach had wanted to publicly share insights with the public or team without disclosing individual players' shooting

7. "Privacy Enhancing Technologies – A Review of Tools and Techniques," Office of the Privacy Commissioner of Canada, November 2017. https://www.priv.gc.ca/en/opc-actions-and-decisions/research/explore-privacy-research/2017/pet_201711.
8. Alexander Macgillivray and Tess deBlanc-Knowles, "Advancing a Vision for Privacy-Enhancing Technologies," NITRD, June 28, 2022. https://www.nitrd.gov/advancing-a-vision-for-privacy-enhancing-technologies; "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence," National Archives, Federal, Nov. 1, 2023. https://www.federalregister.gov/documents/2023/11/01/2023-24283/safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence.
9. Uzair Javaid, "Why Are Legacy Data Anonymization Techniques Failing?," BetterData, Aug. 12, 2024. https://www.betterdata.ai/blogs/why-are-legacy-data-anonymization-techniques-failing.]
10. Ahmed Hilali et al., "Linkage Attack and Protection Mechanism for Social Network From Mobility Profile," *2020 International Conference on Computer, Control, Electrical, and Electronics Engineering* 17 (May 2021), pp. 1-6. https://ieeexplore.ieee.org/document/9429665; Jaideep Vaidya et al., "Identifying inference attacks against healthcare data repositories," *AMIA Joint Summits on Translational Science Proceedings* (March 18, 2013), pp. 262-266. https://pmc.ncbi.nlm.nih.gov/articles/PMC3845790; Aloni Cohen et al., "The Theory of Reconstruction Attacks," Differential Privacy, Oct. 21, 2020. https://differentialprivacy.org/reconstruction-theory.
11. Latanya Sweeney, "k-Anonymity: A Model for Protecting Privacy," *International Journal on Uncertainty, Fuzziness and Knowledge-based Systems* 10:5 (2002), pp. 557-570. https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=2ab47454f59d9d8e55d4d8a69530562a3690794a.
12. Lydialyle Gibson, "When Technology and Society Clash," Harvard Magazine (November-December 2024). https://www.harvardmagazine.com/2024/11/technology-election-impact-harvard-latanya-sweeney-research.
13. Noseong Park et al., "Data Synthesis based on Generative Adversarial Networks," *Proceedings of the VLDB Endowment* 11:10 (July 2, 2018), pp. 1071-1083. https://arxiv.org/abs/1806.03384.
14. Cynthia Dwork et al., "Our Data, Ourselves: Privacy Via Distributed Noise Generation," Advances in Cryptology - EUROCRYPT (2006), pp. 486-503. https://doi.org/10.1007/11761679_29.
15. Cem Dilmegani, "Differential Privacy: How It Works, Benefits & Use Cases in 2025," AI Multiple, Oct. 13, 2024. https://research.aimultiple.com/differential-privacy.
16. Jack Maloney, "Light the beam! Everything you need to know about Kings' unique victory celebration," CBS Sports, April 16, 2023. https://www.cbssports.com/nba/news/light-the-beam-everything-you-need-to-know-about-kings-unique-victory-celebration.

percentages, he could have used DP to add a privacy budget parameter (epsilon) to the dataset, which would control the noise or randomness added to the raw shooting percentage data.[17] The coach could have then reported that the team shot about 30 percent from 3-point range. Although not as accurate as the exact statistic, this degree of variation makes it more difficult to identify specific players' shooting performance, thereby providing stronger protection for individual players. On the other hand, the coach could report 33.5 percent when the actual value is 33.9 percent. This minor variation provides more accurate statistics but less privacy protection for individual players, and it could be easier to figure out which player(s) are contributing to the Kings' 3-point shooting struggles. This example highlights two major challenges with DP. First, too little noise or randomness might compromise privacy, whereas too much can render the data useless. Second, each query search into the players shooting percentage would weaken the privacy guarantee, and there would have to be a set limit on search queries to maintain a privacy guarantee.[18]

AI could be a solution for these challenges. Specifically, analysts could use deep neural networks (DNN), which are a subset of machine learning, to generate a privacy budget parameter with sophisticated noise patterns that preserve data utility while maintaining strong privacy guarantees.[19] Going back to the basketball example above, when using a DNN for adaptive noise generation, the AI system might analyze which statistical elements are most crucial for maintaining the data's utility instead of applying uniform noise to all 3-point shooting stats. For example, the system might determine that maintaining the mathematical relationship between successful shots and total attempts is more critical than precise individual game performances. It could then adapt accordingly to the noise generation, preserving crucial relationships while blurring individual identifying patterns that could reveal specific players' 3-point shooting percentages.

Federated Learning

Federated learning (FL) emerged from Google's research in 2016 and incorporates the principles of data minimization and anonymization.[20] It is a machine-learning approach that enables organizations to train AI models (or other research collaborations) across multiple decentralized devices or servers holding local data samples without exchanging the raw data.[21] This approach offers important privacy advantages, as it enables organizations to collaborate while keeping sensitive data within their own secure environments.

FL emerged from Google's research in 2016 and incorporates the principles of data minimization and anonymization.

Using another a basketball analogy, National Basketball Association (NBA) viewership has decreased this season, and some opine this is because star players are missing more games because of injury and load management.[22] If multiple NBA teams wanted to help increase star players' availability by developing an AI model that predicted player fatigue and injury risk without sharing sensitive performance data, FL could help resolve this issue. FL could help address this by having a central server distribute an initial AI model

17. "Privacy budget fundamentals," Tumult Labs, last accessed Jan. 1, 2025. https://docs.tmlt.dev/analytics/v0.11/topic-guides/privacy-budgets.html.
18. Amalie Dyda et al., "Differential privacy for public health data: An innovative tool to optimize information sharing while protecting data confidentiality," *Patterns (NY)* 2:12 (Dec. 10, 2021). https://pubmed.ncbi.nlm.nih.gov/34909703.
19. Richard M. Lueptow and Xiuyang Lü, "Deep Neural Network, "*Computer Aided Chemical Engineering* 49 (2022), pp. 1813-1818. https://www.sciencedirect.com/topics/chemical-engineering/deep-neural-network; Maoguo Gong et al., "Preserving differential privacy in deep neural networks with relevance-based adaptive noise imposition," *Neural Networks* 125 (May 2020), pp. 131-141. https://www.sciencedirect.com/science/article/abs/pii/S0893608020300460.
20. Brendan McMahan and Daniel Ramage, "Federated Learning: Collaborative Machine Learning without Centralized Training Data," Google Research, April 6, 2017. https://research.google/blog/federated-learning-collaborative-machine-learning-without-centralized-training-data.
21. Ibid.
22. Sean Burch, "Why the NBA's Ratings Are Down Big — and Why Its New Media Partners Should Care," Yahoo Sports, Dec. 2, 2024. https://sports.yahoo.com/why-nba-ratings-down-big-141500434.html.

to all participating teams. Each team would then train this model using their local player data, including metrics like minutes played, in-game (and practice) movement patterns, confidential team practice data, player injury data, and biometric readings. Rather than sharing this sensitive and/or confidential information, each team only sends back updates about how the model's predictions have improved based on their local training.

The central NBA server would combine these model updates from each team to create an enhanced global model that benefits from collective learning without any team having to reveal their sensitive or confidential data. In theory, the AI model would improve, and this improved model would then be redistributed to all collaborating teams for another round of local training.

Thus, all participating organizations would receive a trained AI model that could generate actionable insights that each team could implement as necessary. For example, player-specific load management, optimal minutes per a game, required rest between games, and long-term schedule management. The insights would be more reliable because they would be based on a much larger dataset, but they would also still protect each team's confidential data.

The use cases for FL are broad and could positively impact various fields of research, such as advertisement marketing strategies, product safety research and development, and medical research. Society would reap benefits from collaborative research and significantly decreased privacy concerns.

Using AI to Automate Data Privacy Impact Assessments

Data privacy impact assessments, also known as DPIAs or privacy risk assessments, are another tool for which AI technology could be used to improve privacy and security. Organizations conduct DPAs to identify and minimize data protection risks in projects, services, or initiatives that involve processing personal data.[23] These assessments typically involve several components, including data mapping, risk assessments, and mitigation strategies.[24] AI can be leveraged to automate these assessments. For the purpose of this paper, the term DPIA is synonymous with privacy risk assessments and privacy impact assessments.

DPIAs are arguably one of the most critical steps organizations can take to comply with various regulations or laws because they are key to proactively identifying potential privacy risks and fostering accountability and transparency in data processing. The foundation of DPIA principles can be traced back to the Fair Information Practice (FIP), which was developed in the 1970s by the U.S. Department of Health.[25] These principles were further developed in the Privacy Act of 1974, which required federal agencies to publish system of records notices (SORNs).[26] SORNs functioned as an early form of DPIAs in that they served to document how agencies collected, used, and protected personal information per FIP principles.

DPIAs are arguably one of the most critical steps organizations can take to comply with various regulations or laws because they are key to proactively identifying potential privacy risks and fostering accountability and transparency in data processing.

23. "Data Protection Impact Assessments," Data Protection Commission, last accessed Dec. 15, 2024. https://www.dataprotection.ie/en/organisations/know-your-obligations/data-protection-impact-assessments.
24. "How do we do a DPIA?," Information Commissioner's Office, last accessed Dec. 15, 2024. https://ico.org.uk/for-organisations/uk-gdpr-guidance-and-resources/accountability-and-governance/data-protection-impact-assessments-dpias/how-do-we-do-a-dpia.
25. Office of the Assistant Secretary for Planning and Evaluation, "Records, Computers and the Rights of Citizens," U.S. Department of Health and Human Services, June 30, 1973. https://aspe.hhs.gov/reports/records-computers-rights-citizens.
26. "System of Records Notices (SORNs)," U.S. Department of the Treasury, last accessed Dec. 3, 2024. https://home.treasury.gov/footer/privacy-act/system-of-records-notices-sorns; "General Data Protection Regulation: GDPR," Intersoft Consulting, last accessed Dec. 3, 2024. https://gdpr-info.eu.

DPIAs continued to evolve over the ensuing decades, driven by federal and global privacy laws, including the General Data Protection Regulation (GDPR).[27] Specifically, Article 35 in the GDPR explicitly mandates risk assessments for high-risk data-processing activities and has influenced other global privacy frameworks, such as Brazil's General Data Protection Law and India's Digital Personal Data Protection Act.[28] DPIAs have also increasingly been added to the most comprehensive privacy and security laws in the United States.

DPIAs' ubiquity in the United States and other global privacy laws reflects the nature of today's data-driven global economy, which is characterized by increasingly complex data processing and increasingly frequent data breaches and privacy incidents.[29]

AI models trained to automate DPIAs are uniquely capable of continuously evaluating and identifying potential privacy vulnerabilities in complex data systems, which improves response times and the quality of response to these increased threats in a number of ways. First, automated DPIAs offer scalability and efficiency to organizations that collect personal data and process a large volume of data.[30] Second, privacy risk experts' opinions on what constitutes a "reasonable risk" differ widely. Manual assessments are therefore prone to inconsistencies and subjectivity because different evaluators might interpret privacy risks and/or company policies differently. Automation can ensure consistency on every assessment and reduce human error. Third, automation can proactively mitigate risks, perhaps even catching risks before they lead to a data breach or regulatory or legal compliance violation.

Automated DPIAs offer scalability and efficiency to organizations that collect personal data and process a large volume of data.

Privacy risk experts' opinions on what constitutes a "reasonable risk" differ widely.

Automation can proactively mitigate risks, perhaps even catching risks before they lead to a data breach or regulatory or legal compliance violation.

Scalability and Efficiency

A modern AI system can efficiently scan terabytes of data in hours or minutes, depending on computational power and optimization.[31] In contrast, the average human can only read and process about 240 words per minute.[32] This makes a compelling use case for an organization looking to automate DPIAs because it could be deployed to scan large volumes of data for compliance checks for existing or future laws and regulations. This is especially important in the United States, which operates without a comprehensive federal data privacy and security law, instead opting to defer to a patchwork of state privacy laws.[33]

Human Error, Inconsistencies, and Subjectivity

Whether due to competency or biological factors, humans are prone to error when protecting sensitive data. A 2023 report revealed that 74 percent of data breaches involve a human element, including errors such as misdelivery, misconfiguration, and phishing scams.[34]

27. Ibid.
28. "Brazilian General Data Protection Law (LGPD, English translation)," International Association of Privacy Professionals, October 2020. https://iapp.org/resources/article/brazilian-data-protection-law-lgpd-english-translation; "The Digital Personal Data Protection Bill, 2023," PRS Legislative Research, last accessed Dec. 30, 2024. https://prsindia.org/billtrack/digital-personal-data-protection-bill-2023.
29. Li Cai and Yangyong Zhu, "The Challenges of Data Quality and Data Quality Assessment in the Big Data Era," *Data Science Journal* 14:2 (May 22, 2015). https://datascience.codata.org/articles/10.5334/dsj-2015-002; Stuart Madnick, "What's Behind the Increase in Data Breaches?," *The Wall Street Journal*, March 15, 2024. https://www.wsj.com/tech/cybersecurity/why-are-cybersecurity-data-breaches-still-rising-2f08866c.
30. Computer Security Resource Center, "Attack surface," National Institute of Standards and Technology, last accessed Dec. 12, 2024. https://csrc.nist.gov/glossary/term/attack_surface.
31. Sejal Sharma, "Samsung unveils the 'world's fastest' data processing AI chip to date with a bandwidth of 1.25 terabytes per second," Interesting Engineering, Feb. 27, 2024. https://interestingengineering.com/innovation/samsung-unveils-fastest-aichip.
32. Marc Brysbaert, "How many words do we read per minute? A review and meta-analysis of reading rate," PsyArXiv Preprints, April 12, 2019. https://doi.org/10.31234/osf.io/xynwg.
33. Steven Ward and Brandon Pugh, "Bipartisanship on comprehensive federal privacy and security legislation continues at House Subcommittee on Innovation, Data, and Commerce hearing," R Street Institute, March 2, 2023. https://www.rstreet.org/commentary/bipartisanship-on-comprehensive-federal-privacy-and-security-legislation-continues-at-house-subcommittee-on-innovation-data-and-commerce-hearing.
34. Jeff Peters, "Human error is responsible for 74% of data breaches," Infosec Institute, Nov. 30 2023. https://www.infosecinstitute.com/resources/security-awareness/human-error-responsible-data-breaches.

A Tessian Security and Stanford University study found that approximately 88 percent of data breaches are caused by employee mistakes, emphasizing the critical role of human error in privacy-related incidents.[35] Some of these studies gathered facts during the COVID-19 pandemic when individuals were learning how to work from home securely and might have been more susceptible to cyber attacks. However, the number of human errors shows that AI systems could be a solution for protecting sensitive data. While AI is not always perfect, it can be a beneficial supplement for humans.

88%

A Tessian Security and Stanford University study found that approximately 88 percent of data breaches are caused by employee mistakes, emphasizing the critical role of human error in privacy-related incidents.

Inconsistencies across an organization's communication, DPIA deployment, or policies can result in unmitigated risk. Even if an organization has a detailed and thorough DPIA policy in place, that policy may not be applied correctly at every human-led deployment, which can cause overlooked issues, resulting in noncompliance with regulations and laws. Further, inconsistent communications across an organization, like IT, legal, or operations/product teams, might result in incomplete identification of data-processing activities.

Subjectivity can also significantly influence an organization's DPIA process and potentially lead to data-protection vulnerabilities. For example, different employees will perceive a threat differently based on experience or interpretation. Even if employees correctly identify a privacy risk, it may not be swiftly mitigated due to a low-risk categorization or lack of attention. Subjectivity may also lead to an overreliance on deploying familiar security measures while neglecting more effective, but lesser-known or newer, measures.

Whether human-error issues arise in the identification or safeguarding stage of tackling a threat, AI's speed and adaptability are an increasingly viable option for mitigation. AI could more effectively support human activities and data protection than relying on fallible humans alone.

Real-Time Risk Mitigation

DPIAs are designed to identify and mitigate risks to individual privacy before they materialize. They are an essential principle of "privacy by design," a concept introduced in 1995.[36] A strength of AI systems is that they can provide real-time identification and mitigation of risks through several technological techniques and approaches, like continuous monitoring and analysis of system and user behavior, identifying deviations from standard behavior patterns that might indicate privacy violations. Further, the ability of machine learning algorithms allows AI systems to learn from historical data and predict or adapt to new and emerging privacy threats without human intervention.[37] This provides organizations with proactive solutions for the threat landscape that historically have relied on reactive solutions.

Data Minimization and AI

Data minimization is a fundamental and critical privacy principle that like many privacy principles, originated from FIP and can also be enhanced with AI applications.[38] It requires that organizations limit their collection and retention of personal data to what is relevant and necessary to accomplish a specific purpose.[39] For good reason, data minimization provisions are included in most state, federal, and global privacy laws. Some proposed

35. Jeff Hancock "The Psychology of Human Error: 2022 Report," Tessian, 2022. https://f.hubspotusercontent20.net/hubfs/1670277/%5BCollateral%5D%20Tessian-Research-Reports/%5BTessian%20Research%5D%20Psychology%20of%20Human%20Error%202022.pdf.
36. Ann Cavoukian, "Privacy by Design: Foundational Principles," University of California Santa Cruz, last accessed Nov. 8, 2024. https://privacy.ucsc.edu/resources/privacy-by-design---foundational-principles.pdf.
37. "What is Self-Learning AI and How is it Applied in Cybersecurity," MixMode, last accessed Jan. 15, 2025. https://mixmode.ai/what-is/self-learning-ai.
38. "Cost of a Data Breach: Report 2024." https://www.ibm.com/reports/data-breach.
39. Ibid.

laws, like the American Privacy Rights Act of 2024, had stringent data minimization provisions tied to permissible purposes.[40] It is critical that stakeholders come together to determine how best to meet the stringent data-minimization provisions embedded in privacy legislation, and finding the right balance is essential because too much data-minimization can negatively affect future equality rights and innovation.[41] Finding the right balance will take a collective effort from all stakeholders.

A notable challenge around data minimization is convincing organizations to limit data collection when certain data could potentially help with current or future business needs.

A notable challenge around data minimization is convincing organizations to limit data collection when certain data could help with current or future business needs. However, advanced AI models could demonstratively convince organizations that the data they retain and process fulfills all their current and future business needs. By leveraging AI in this way, organizations might be less compelled to retain irrelevant data that serves only to increase their threat landscape.

AI can be particularly helpful in this regard in its ability to classify data in complex ways. Data classification is the process of organizing data into categories based on its sensitivity, importance, and other specific criteria, like whether data is public, internal, or confidential.[42] It supports data-minimization efforts in several ways, such as identifying sensitive data, prioritizing data protections, assisting in regulatory compliance, and reducing risk.[43]

Manually classifying large volumes of data is time consuming and resource intensive.[44] Most organizations struggle to maintain consistent classification standards across different teams and departments, leading to inconsistencies in how data is categorized and protected. Furthermore, classification accuracy can be challenging, and misclassification can lead to unnecessary or inadequate security controls, reduced efficiency, and noncompliance.[45]

Some organizations attempt to automate data classification by implementing data crawlers to scan through the organization's file systems and data-storage reservoirs to locate and classify data.[46] They also implement a rules-based classification system that organizes data based on predefined rules, such as laws or regulations defining sensitive data. But these systems have challenges and limitations. For example, if an individual submitted an organization complaint or feedback form while also discussing their specific medical conditions, sensitive patient data could unwittingly enter an organization's data flow. This data might be missed as sensitive when analyzed by a rules-based or data-crawler classification system.

An AI model designed for data classification could enhance responses to data-classification tasks by addressing the contextual and nuanced challenges that non-AI systems, including rule-based systems, struggle to handle. Through natural language processing, an AI system can understand the use of context and linguistic subtleties that a rigid rule-based classification system typically misses.[47]

40. H.R.8818 - 118th Congress (2023-2024): American Privacy Rights Act of 2024, (June 25, 2024). https://www.congress.gov/bill/118th-congress/house-bill/8818/text; Brandon Pugh and Steven Ward, "Is the Third Time the Charm? Analyzing APRA 3.0," R Street Institute, June 21, 2024. https://www.rstreet.org/commentary/is-the-third-time-the-charm-analyzing-apra-3-0.
41. Orly Lobel, "The Law of AI for Good," *Florida Law Review* 75:6 (2023). https://ssrn.com/abstract=4338862; Tal Roded and Peter Slattery, "What drives progress in AI? Trends in data," FutureTech, March 19, 2024. https://futuretech.mit.edu/news/what-drives-progress-in-ai-trends-in-data.
42. Morgan Sullivan, "Understanding Data Classification: Enhance Security & Efficiency," Transcend, Dec. 8. 2023, https://transcend.io/blog/data-classification.
43. Ibid.
44. "Auto Classification in Records Management," OpenText Blogs, March 28, 2023. https://blogs.opentext.com/auto-classification-in-records-management.
45. Ibid.
46. J.C. Miraclin Joyce Pamila et al., "Ensemble classifier based big data classification with hybrid optimal feature selection," *Advances in Engineering Software* 173 (November 2022). https://www.sciencedirect.com/science/article/pii/S0965997822000928.
47. Salvatore Claudio Fanni et al., "Natural Language Processing," in *Introduction to Artificial Intelligence* (Springer, 2024), pp. 87-99. https://doi.org/10.1007/978-3-031-25928-9_5.

Data Security and AI

Data security is yet another area where AI models and applications can be leveraged to gain efficiencies and improvements. It encompasses the technical controls, systems, and processes that protect data from unauthorized access, breaches, and cyber threats. These protective measures, such as encryption, access controls, and network security, create the secure framework upon which data privacy can be maintained. Without these fundamental safeguards, even the most robust privacy policies and procedures would be ineffective. Thus, data security and data privacy are inextricably linked. This is why data security provisions, while usually short and sweet in many privacy laws, act as the backbone of every significant privacy law.[48]

When data security fails, the consequences can be severe and wide ranging—from eroded consumer trust to legal liability to compromised national security.[49] Moreover, data has emerged as a strategic resource with significant military and economic implications, making data security innovations all the more important.[50] In fact, nations are increasingly viewing data control and security as essential elements of national power, similar to traditional military assets or natural resources.[51] This was reflected in Congress' 2023 hearing in which TikTok's CEO was asked a range of questions over nearly 5 hours aimed at assessing data privacy and security concerns surrounding the popular platform.[52] Congress was not satisfied with the CEO's responses, so a subsequent law was passed forcing TikTok to divest or be banned in the United States because of national security concerns around China's access to American TikTok users' data.[53]

Nations are increasingly viewing data control and security as essential elements of national power, similar to traditional military assets or natural resources.

Two key areas of data security that could be bolstered by AI innovations are data encryption and homomorphic encryption, both of which are discussed in more detail below.

Data Encryption

Data encryption technology is a critical data-security measure that converts readable data (plaintext) into an encoded format (ciphertext) that can only be decoded with the correct encryption key.[54] Modern encryption has become essential for protecting various types of data and is used across industries to protect processes and data including financial transactions, medical records, law enforcement records, and military communications. Encryption is implemented to protect data both when it is in storage (at rest) and when it is being transmitted across networks (in transit).[55] For example, when browsing a nonconsequential website, a user might notice that the URL might begin with "http," whereas the URL they see when browsing their banking information begins with "https."[56] Because of the sensitive data involved, the user's bank encrypts the communication

48. Steven Ward, "New Series: The Quest for 'Reasonable Security,'" R Street Institute, Jan. 29, 2024. https://www.rstreet.org/commentary/new-series-the-quest-for-reasonable-security.
49. Ward. https://www.rstreet.org/commentary/new-series-the-quest-for-reasonable-security; R Street Institute, "Data Privacy and Security as a National Security Imperative," YouTube, Oct. 26, 2023. https://www.youtube.com/watch?v=-1YprkJdGiA; Brandon Pugh and Chris Riley, "R Street Institute Comments on FTC's ANPR on Commercial Surveillance and Data Security," R Street Institute, Oct. 24, 2022. https://www.rstreet.org/outreach/r-street-institute-comments-on-ftcs-anpr-on-commercial-surveillance-and-data-security.
50. "DOD Issues New Data Strategy," U.S. Department of Defense, Oct. 8, 2020. https://www.defense.gov/News/Releases/Release/Article/2376629/dod-issues-new-data-strategy.
51. Ibid.
52. House Energy and Commerce Committee, "TikTok: How Congress Can Safeguard American Data Privacy and Protect Children from Online Harms," 118th Congress, March 23, 2023. https://www.govinfo.gov/content/pkg/CHRG-118hhrg53839/pdf/CHRG-118hhrg53839.pdf.
53. Sapna Maheshwari and David McCabe, "Congress Passed a Bill That Could Ban TikTok. Now Comes the Hard Part," *The New York Times*, April 23, 2024. https://www.nytimes.com/2024/04/23/technology/bytedance-tiktok-ban-bill.html.
54. "What Is Data Encryption?," Digital Guardian, last accessed Jan. 1, 2025. https://www.digitalguardian.com/resources/knowledge-base/what-data-encryption.
55. "Data Protection: Data In Transit vs. Data At Rest," Digital Guardian, May 6, 2023. https://www.digitalguardian.com/blog/data-protection-data-in-transit-vs-data-at-rest.
56. "What is a URL?," Mozilla Developer Network, last accessed Jan. 1, 2025. https://developer.mozilla.org/en-US/docs/Learn_web_development/Howto/Web_mechanics/What_is_a_URL.

between the user's web browser and servers via a transport layer security, more commonly referred to as "TLS."[57]

The modern threat landscape poses three significant challenges to organizations that must implement and maintain effective encryption. First, encryption-key management can be challenging, particularly for large or fast-growing organizations that must manage numerous encryption keys to keep data secure yet accessible to the correct individuals.[58] Second, small and medium-sized businesses often lack the resources to implement robust encryption standards. Third, interoperability and compatibility become difficult when systems must work across different platforms and adhere to different security standards. For example, HIPAA requires medical record portability.[59] If a patient chooses to go to a different medical provider, their previous provider must securely transfer that patient's medical records to their new provider. Realistically, a healthcare system must share patient records with multiple external providers. However, each provider might use different electronic health record systems with varying encryption standards.

AI can provide the technology to resolve these challenges. AI algorithms can provide enhanced key-management solutions by automating key generation, distribution, and rotation, reducing human error and improving overall security.[60] Future AI research could also assist with encryption interoperability and compatibility across different legacy systems.[61]

Another key innovation is AI-powered adaptive encryption in data-protection technology.[62] This approach enables encryption systems to adjust their security protocols automatically in response to evolving threats, providing advantages over conventional encryption methods.[63] Adaptive systems continuously evaluate potential risks and modify security measures accordingly, ensuring adequate protection at all times.[64] This technology is quite valuable for small and medium-sized businesses that often lack security resources, making advanced data protection more accessible and less resource intensive for organizations of all sizes.

Homomorphic Encryption

Full homomorphic encryption (FHE) is an advanced cryptographic technique that allows computations to be performed directly on encrypted data without the need for decryption.[65] In the context of AI, this technology allows organizations to process sensitive data while maintaining confidentiality. Unfortunately, FHE requires significant computational resources, which limits its accessibility for organizations that lack such resources.[66] The FHE market is expected to reach only $358 million by 2028, growing 8.3 percent from 2023 to 2030, which likely indicates that the computational-cost factor is limiting market adoption.[67] The hope is that with AI advances and decreased computational power cost, FHE will become more accessible over time.

FHE is an advanced cryptographic technique that allows computations to be performed directly on encrypted data without the need for decryption.

57. "Why is HTTP not secure?," Cloudflare, last accessed Jan. 1, 2025. https://www.cloudflare.com/learning/ssl/why-is-http-not-secure.
58. "Managing Risks in Cryptographic Key Management," Cryptomathic, Jan. 21, 2022. https://www.cryptomathic.com/blog/cryptographic-key-management-the-risks-and-mitigations.
59. "Individuals' Right Under HIPAA to Access their Health Information 45 CFR § 164.524," U.S. Department of Health and Human Services, last accessed Jan. 1, 2025. https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/access/index.html.
60. Samuel Omokhafe Yusuf et al., "Analyzing the efficiency of AI-powered encryption solutions in safeguarding financial data for SMBs," *World Journal of Advanced Research and Reviews* 23:3 (2024), pp. 2138-2147. https://wjarr.com/sites/default/files/WJARR-2024-2753.pdf.
61. Ibid.
62. Ibid.
63. Ibid.
64. Ibid.
65. "What Is Homomorphic Encryption?," IEEE Digital Privacy, last accessed Jan. 1, 2025. https://digitalprivacy.ieee.org/publications/topics/what-is-homomorphic-encryption.
66. Ibid.
67. "Homomorphic Encryption Market," Next Move Strategy Consulting, last accessed Jan. 1, 2025. https://www.nextmsc.com/report/homomorphic-encryption-market.

AI and Notice-and-Choice Privacy Frameworks

In the United States, the notice-and-choice framework is prevalent in most states' comprehensive privacy laws. For notice, covered entities must provide a privacy notice to consumers that is readily available and written clearly enough for consumers to understand how to exercise their rights. For choice, consumers can make "informed decisions" about whether and how their data is processed. For example, when a user visits a website, a cookie-consent banner or "pop-up" box alerts the user of that organization's privacy policy. These alerts attempt to gain the user's informed consent for data collection and processing. In theory, the user has a choice to accept the organization's privacy practices or opt-out of data collection. But notice-and-consent frameworks have received criticism:

> Hardly anyone reads privacy notices, those who try to read them struggle to understand them, the statements in privacy notices are often vague and ambiguous, and the effort to read privacy notices does not scale because there are too many to read. As a result, a remarkably low percentage of people opt out, which allows organizations to use personal data with only the vague self-imposed limits stated in the privacy notices.[68]

AI could be essential in alleviating such criticism by improving user transparency via revealing significant deviations from privacy norms and potential privacy concerns with an organization's privacy notice. Further, AI could make complicated privacy notices more digestible and alert users of significant changes in data use or privacy policies.

Ideally, AI deployment in notice-and-consent frameworks could allow users to quickly and effectively understand how their data will be used and processed, reaching a state of true informed consent, rather than the far more common gesture toward informed consent. For example, an AI browser plugin could quickly read a privacy notice and inform the user of the particular privacy policy and real-world alerts for privacy failures or previously deceptive practices investigated by a consumer protection agency like the Federal Trade Commission.

Ideally, AI deployment in notice-and-consent frameworks could allow users to quickly and effectively understand how their data will be used and processed, reaching a state of true informed consent.

Policy Recommendations

As the United States has ushered in its next administration and the 119th Congress, new leadership must take a balanced approach to AI and emerging technologies. Our country's strength lies in supporting innovation, and it is critical that this continue by leveraging AI technologies in inventive ways to advance data privacy and security. Below are key recommendations that we suggest lawmakers consider to ensure that the United States is at the forefront of these advances.

Pass a Federal Comprehensive Data Privacy and Security Law

Policy Recommendation 1

Passing a comprehensive federal data privacy and security law is one of the best ways to mitigate data privacy risks before data is collected and used to train AI.[69]

- A federal privacy law should protect the rights of all Americans.[70]
- Individuals should be able to control their data, and entities should be transparent

68. Daniel J. Solove, "Murky Consent: An Approach to the Fictions of Consent in Privacy Law," *Boston University Law Review* 104:593 (2024). https://ssrn.com/abstract=4333743.
69. Brandon Pugh and Steven Ward, "Key Data Privacy and Security Priorities for 2025," R Street Institute, Jan. 15, 2025. https://www.rstreet.org/research/key-data-privacy-and-security-priorities-for-2025.
70. Steven Ward, "Privacy Should Be a Fundamental Right for All, Not Just for Elites," Inkstick Media, Sept. 18, 2023. https://inkstickmedia.com/privacy-should-be-a-fundamental-right-for-all-not-just-for-elites.

about the type of data they collect; how they intend to use the data; whether they will sell, share, or transfer the data to a third party; and how long they will retain the data. This includes individual rights to delete, access, and port data.

- Strong preemption remains a critical feature of a federal privacy proposal, given the increasing number of state laws and the challenges with compliance, especially for small and medium-sized businesses. Preemption language must be carefully crafted to limit the proliferation of burdensome state-level privacy laws.[71]
- A federal privacy law should have data-security provisions for collecting and processing data with a flexible, nonprescriptive approach.
- A flexible data-minimization provision could productively limit the data collected and force entities to understand why they need specific data and how it can be used, thus eliminating superfluous data collection and limiting its abuse.

Establish Flexible and Pro-innovation AI Policy

Policy Recommendation 2

The United States faces challenges and pressure to develop AI governance frameworks that protect public interests while maintaining technological leadership. As AI continues transforming industries and society, it will be critical that related policy support innovation to ensure that the United States maintains global AI competitiveness. As a previous R Street study notes, "[g]etting governance balance right—and ensuring that it remains flexible, responsive and pragmatic—is essential if the United States hopes to remain at the forefront of global AI innovation and competitiveness."[72]

Policymakers should follow similar free-market and limited-government principles when addressing AI frameworks and development. This will support AI innovation in the United States, which, as a bipartisan congressional AI Task Force noted in a recent report, is vital to maintaining U.S. leadership in AI deployment. The report suggests this can be obtained through a flexible, agile regulatory approach and public–private partnership.[73]

Offer Incentives and Guidance for PET Development

Policy Recommendation 3

Funding opportunities are essential for public–private partnerships that could advance AI innovation. The United States could fund PET developments through grant programs offered by the U.S. National Science Foundation, National Institute of Standards

and Technology funding opportunities, and the Department of Homeland Security cybersecurity grant program.[74]

Encouragingly, last year saw bipartisan support for PET developments, including legislation introduced by Rep. Haley Stevens (D-MI) and co-led by Rep. Tom Kean Jr. (R-NJ), which was overwhelmingly passed by the House but stalled in the Senate.[75] This legislation

71. Tatyana Bolton et al., "Preemption in Federal Data Security and Privacy Legislation," R Street Institute, May 31, 2022. https://www.rstreet.org/commentary/preemption-in-federal-data-security-and-privacy-legislation.
72. Adam Thierer, "Flexible, Pro-Innovation Governance Strategies for Artificial Intelligence," R Street Institute, April 20, 2023. https://www.rstreet.org/research/flexible-pro-innovation-governance-strategies-for-artificial-intelligence.
73. House AI Task Force, "Bipartisan House Task Force Report on Artificial Intelligence: Guiding principles, forward-looking recommendations, and policy proposals to ensure America continues to lead the world in responsible AI innovation," U.S. House of Representatives (2024). https://republicans-science.house.gov/_cache/files/a/a/aa2ee12f-8f0c-46a3-8ff8-8e4215d6a72b/E4AF21104CB138F3127D8FF7EA71A393.ai-task-force-report-final.pdf.
74. "Security, Privacy, and Trust in Cyberspace (SaTC 2.0)," U.S. National Science Foundation, last accessed Jan. 1, 2025. https://new.nsf.gov/funding/opportunities/satc-20-security-privacy-trust-cyberspace; "2020 Differential Privacy Temporal Map Challenge," National Institute of Standards and Technology, last accessed Jan. 1, 2025. https://www.nist.gov/ctl/pscr/open-innovation-prize-challenges/past-prize-challenges/2020-differential-privacy-temporal; "DHS Announces $279.9 Million in Grant Funding for the Fiscal Year 2024 State and Local Cybersecurity Grant Program," U.S. Department of Homeland Security, Sept. 23, 2024. https://www.dhs.gov/news/2024/09/23/dhs-announces-2799-million-grant-funding-fiscal-year-2024-state-and-local.
75. H.R.4755 - 118th Congress (2023-2024): Privacy Enhancing Technology Research Act. (April 30, 2024). https://www.congress.gov/bill/118th-congress/house-bill/4755/all-actions.

would have driven the advancement of PETs through research funding, workforce training initiatives, standards development, and coordinated government efforts. Legislation like this should be a priority for the next administration and the 119th Congress.

Conclusion

As the United States navigates the future of AI and other emerging technologies, we can acknowledge the inherent privacy risks and benefits that could come with such innovations. Rather than allowing privacy concerns to quell innovative progress, however, we should embrace technological advancement—as the United States has historically done—by implementing a comprehensive federal data privacy and security law, developing flexible AI governance frameworks, and continuing to offer support for PETs.

The path forward requires carefully crafted policies that balance innovation with responsive and agile guardrails. Industry leaders and policymakers have an opportunity to establish privacy standards and best practices. Taking proactive steps in this process can ensure that technological progress and privacy protection advance together, serving society's best interests while fostering continued innovation.

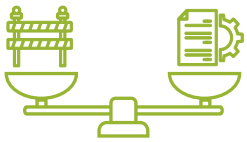

The path forward requires carefully crafted policies that balance innovation with responsive and agile guardrails. Industry leaders and policymakers have an opportunity to establish privacy standards and best practices.

About the Author

Steven Ward is a privacy and security fellow for R Street's Cybersecurity and Emerging Threats team. He focuses on data security, data privacy, and cybersecurity at the federal and state levels; emerging cyber and national security threats; and cybersecurity impacts on the private sector.