

ARTICLE

Artificial Intelligence in Healthcare: Mapping the Legal and Ethical Boundaries of Privacy

Jennifer Jihae Park*

Abstract

Artificial intelligence (AI) is rapidly transforming healthcare while exposing significant gaps in existing legal and ethical frameworks governing patient privacy, autonomy, and algorithmic accountability. This Article examines the regulation of healthcare AI in the United States, focusing on the Health Insurance Portability and Accountability Act (HIPAA), federal oversight, emerging state privacy laws, and ethical frameworks developed by the U.S. Department of Health and Human Services (HHS), American Medical Association (AMA), World Health Organization (WHO), and U.S. Food and Drug Administration (FDA). It compares this fragmented approach with the European Union's more centralized, rights-based framework under the General Data Protection Regulation (GDPR) and European Union Artificial Intelligence Act (EU AI Act). The analysis demonstrates that ethical principles concerning transparency, accountability, equity, human oversight, and continuous monitoring frequently exceed enforceable U.S. legal requirements. Drawing on European regulation and emerging state approaches, this Article proposes a more harmonized federal framework incorporating comprehensive health data protection, mandatory AI impact assessments, meaningful transparency requirements, and lifecycle-based oversight.

Keywords: Healthcare Artificial Intelligence; Health Data Privacy; AI Governance

A Introduction

Artificial intelligence (AI) has become a buzz word since the launch of ChatGPT by OpenAI in 2022.¹ Since then, there has been growing concerns regarding privacy, accountability and ethics in using AI in diverse sectors including education, business and healthcare.² In fact, AI has rapidly become an integral part of modern healthcare from diagnostic imaging tools that detect cancers with unprecedented accuracy to algorithms predicting patient readmission risks. Data privacy in the healthcare sector is especially critical as it has to do with sensitive patient information.

* Jennifer Jihae Park, Ph.D., Assistant Professor, University of Cincinnati and J.D. Candidate, Indiana University Robert H. McKinney School of Law (2027).

1 OpenAI, *ChatGPT Release Notes* (25 November 2025), <https://help.openai.com/en/articles/6825453-chatgpt-release-notes> [<https://perma.cc/6FH5-KY9V>].

2 Jennifer Jihae Park, *Unlocking Training Transfer in the Age of Artificial Intelligence*, 67 Bus. Horiz. 263 (2024).

AI systems promise to revolutionize medicine by improving efficiency, accuracy and access.³ Yet this technological revolution has also raised profound questions about privacy, accountability and ethics.⁴

The federal law that governs patient information is Health Insurance Portability and Accountability Act (HIPAA). While federal law provides a baseline of privacy protection for “protected health information” (PHI) handled by HIPAA-covered entities, HIPAA’s scope is entity- and transaction-based, leaving large zones of health-related data and many AI use cases outside its reach.⁵ Although federal statutes such as HIPAA establish baseline protections for medical information, much of the regulatory authority over data privacy resides with individual states.⁶ The result is a fragmented legal landscape in which patients’ privacy rights and healthcare providers’ obligations vary depending on jurisdiction.⁷ The current state-law patchwork and HIPAA’s bounded scope together leave consequential AI-driven health data practices under-regulated particularly around secondary use, cross-context data flows and location tracking (e.g. geofencing near healthcare facilities).⁸ HIPAA establishes baseline protections for medical information, and much of the regulatory authority over data privacy resides with individual states.⁹ Additionally, the Federal Trade Commission (FTC),¹⁰ although not a healthcare regulator, plays a central role in the governance of AI.

President Trump’s 2025 executive order¹¹ on AI could be a game changer in the realm of AI law. It imposes broad and strict limits on state-level regulation of AI. The order asserts that AI laws slow innovation, particularly when a patchwork of state regulations makes compliance more challenging and burdens interstate

3 Kirk Stewart, *The Ethical Dilemmas of AI*, USC Annenberg (2024), <https://annenberg.usc.edu/research/center-public-relations/usc-annenberg-relevance-report/ethical-dilemmas-ai> [https://perma.cc/2EM9-A94R].

4 Reuters, *HIPAA-Free Zone? Think Again: The Surprising State Laws Regulating Collection of Health and Fitness Data*, <https://www.reuters.com/legal/litigation/hipaa-free-zone-think-again-surprising-state-laws-regulating-collection-health-2024-10-25/> [https://perma.cc/L7RS-QZDM] (last visited 9 February 2026).

5 U.S. Department of Health and Human Services, *Summary of the HIPAA Privacy Rule* (last updated 2024), <https://www.hhs.gov/hipaa/for-professionals/privacy/laws-regulations/index.html> [https://perma.cc/4VZT-RH86].

6 U.S. Department of Health and Human Services, *HIPAA Privacy Rule – What Is Protected & Covered Entities*, <https://www.hhs.gov/hipaa/for-professionals/privacy/index.html> [https://perma.cc/J2VQ-JSHR].

7 *Id.*

8 HIPAA Journal, *The HIPAA Privacy Rule: Overview and Issues*, <https://www.congress.gov/crs-product/IF12759>.

9 Reuters, *Protecting Reproductive Health Data: State Laws Against Geofencing* (2 January 2025), <https://www.reuters.com/legal/legalindustry/protecting-reproductive-health-data-state-laws-against-geofencing-2025-01-02/> [https://perma.cc/M74F-52R9].

10 Federal Trade Commission Act § 5, 15 U.S.C. § 45(a).

11 The White House, *Ensuring a National Policy Framework for Artificial Intelligence* (11 December 2025), <https://www.whitehouse.gov/presidential-actions/2025/12/eliminating-state-law-obstruction-of-national-artificial-intelligence-policy/> [https://perma.cc/PTN5-3XFD].

commerce. It also establishes an AI litigation task force, and states unwilling to comply with the order may be denied federal funding.¹²

Meanwhile, professional and international organizations, including the U.S. Department of Health and Human Services (HHS),¹³ the American Medical Association (AMA),¹⁴ the World Health Organization (WHO)¹⁵ and the U.S. Food and Drug Administration (FDA),¹⁶ have issued ethical frameworks for the development and deployment of AI in healthcare. These frameworks emphasize transparency, fairness and accountability. However, they often lack enforceability or alignment with existing legal regimes. As a result, the tension between these overlapping ethical principles and inconsistent state privacy laws leaves policymakers, providers and patients without clear guidance.¹⁷

This article asserts that the current patchwork of state privacy laws, when coupled with nonbinding ethical guidance, creates regulatory uncertainty that undermines both innovation and patient trust in AI-enabled healthcare.¹⁸ A systematic comparison of state laws and ethical frameworks reveals significant gaps, particularly in areas concerning data sharing, algorithmic transparency and secondary use of health data.¹⁹ These gaps threaten to erode fundamental principles of medical ethics, including patient autonomy and informed consent.²⁰ However, AI in healthcare is currently in a period of regulatory transition, and President Trump's executive order may create uncertainty regarding the relationship between federal policy and existing state AI laws.

This article proceeds in four parts. Section B provides background on the current state of artificial intelligence in the healthcare sector, defining AI and examining its integration into clinical practice, data-driven decision-making and health information systems. It then surveys the existing legal landscape governing health data privacy, including HIPAA, emerging state privacy statutes and the role of federal agencies, and recent federal policy developments such as President Trump's executive order on AI, as well as the ethical frameworks developed by HHS, the AMA, the WHO and the FDA. This section also introduces recent legal scholarship on AI governance and health data privacy. Section C analyses the relationship between these binding legal regimes and nonbinding ethical frameworks, identifying points of convergence and divergence and highlighting

¹² *Id.*

¹³ U.S. Department of Health and Human Services, *supra* note 5.

¹⁴ *AMA Principles for Augmented Intelligence Development, Deployment, and Use in Health Care*, American Medical Association (14 November 2023), <https://www.ama-assn.org/system/files/ama-ai-principles.pdf> [<https://perma.cc/ED8K-7NPS>].

¹⁵ World Health Organization, *Ethics & Governance of Artificial Intelligence for Health* (2021), <https://www.who.int/publications/i/item/9789240029200> [<https://perma.cc/Q74Z-Y2RW>].

¹⁶ U.S. Food & Drug Administration, *Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan 5-7* (January 2021), <https://www.fda.gov/media/145022/download> (last visited 16 March 2026).

¹⁷ U.S. Department of Health and Human Services, *AI in Health – HHS Policy & Resources*, <https://www.hhs.gov/programs/topic-sites/ai/index.html> (last visited 16 March 2026).

¹⁸ Reuters, *supra* note 9.

¹⁹ Reuters, *supra* note 4.

²⁰ U.S. Department of Health and Human Services, *supra* note 17.

the regulatory gaps that arise when ethical aspirations exceed existing legal obligations. Finally, Section D concludes by considering the implications of this fragmented governance structure for patient privacy, trust and innovation in AI-enabled healthcare and by offering recommendations for aligning legal standards more closely with ethical principles.

B Background

I Defining AI and Its Relationship to Healthcare

Artificial intelligence (AI) is commonly defined as a set of computational methods that enable machines to perform tasks traditionally associated with human cognition, such as pattern recognition, prediction and decision-making.²¹ Although definitions vary across disciplines, a growing scholarly consensus conceptualizes AI as a socio-technical system composed of algorithms, data inputs, computational models and human oversight.²² In the healthcare context, this definition encompasses an array of applications including machine-learning-based diagnostic tools, natural language processing for clinical documentation, predictive analytics for population health and autonomous or semiautonomous decision-support systems.²³

Understanding AI as a system rather than as a discrete technology is essential for healthcare regulation. Clinical AI tools rely heavily on the continuous acquisition and processing of sensitive medical data, including electronic health records (EHRs), imaging data, genomic datasets and real-time biometric data from wearable devices.²⁴ Machine-learning models depend on these data flows to identify statistical regularities and generate clinically relevant outputs. As a result, AI systems are increasingly embedded within routine clinical workflows, blurring the distinction between traditional health information systems and algorithmic decision-making tools.²⁵

Yet defining AI in healthcare is complicated not only by its technical diversity but also by the inherent struggles associated with its use. Many advanced AI models, especially deep-learning systems, operate as “black boxes” whose internal decision-making pathways are difficult or sometimes impossible to interpret.²⁶ This opacity poses substantial challenges for clinicians who must justify their diagnostic or treatment decisions, for patients seeking to understand how

21 Stuart J. Russell & Peter Norvig, *Artificial Intelligence: A Modern Approach* 1-4 (4th ed. 2021).

22 See Kate Crawford, *Atlas of AI: Power, Politics, and the Planetary Costs of Artificial Intelligence* 7-10 (2021); Jenna Burrell, *How the Machine “Thinks”: Understanding Opacity in Machine Learning Algorithms*, 3 *Big Data & Soc’y* 1, 3-5 (2016).

23 See Eric Topol, *Deep Medicine: How Artificial Intelligence Can Make Healthcare Human Again* 47-75 (2019).

24 See Eric Perakslis, *Using Digital Health to Enable Health Equity*, 389 *New Eng. J. Med.* 236, 236-38 (2023).

25 See Ziad Obermeyer & Ezekiel J. Emanuel, *Predicting the Future – Big Data, Machine Learning, and Clinical Medicine*, 375 *New Eng. J. Med.* 1216, 1216-17 (2016).

26 See Burrell, *supra* note 22, 1, 3-6.

algorithms influence their care and for regulators tasked with evaluating system safety and reliability.

Relatedly, generative and predictive models may produce “hallucinations”, supplying outputs that appear authoritative but are factually incorrect or unsupported by underlying data.²⁷ In a clinical environment, such errors risk undermining trust, distorting medical judgment and potentially causing patient harm.

AI systems are also susceptible to model drift, a phenomenon in which an algorithm’s performance degrades over time as population characteristics, clinical practices or environmental conditions change.²⁸ Drift threatens the reliability of AI tools that regulators have traditionally evaluated as static medical devices, complicating post-market oversight and raising concerns about unnoticed performance deterioration.

Equally significant are persistent issues of algorithmic bias, which can emerge from non-representative training datasets, structural inequities embedded in historical health data or unexamined design choices.²⁹ Bias can lead to systematically worse outcomes for marginalized populations, thereby reinforcing or amplifying existing health disparities.

The complexity of defining AI in healthcare becomes even more pronounced within regulatory frameworks. The FDA distinguishes between “locked” and “adaptive” machine-learning algorithms, recognizing that continuously learning models challenge existing medical device approval pathways.³⁰ The European Union takes a broader regulatory approach: the EU Artificial Intelligence Act defines AI to include a wide range of computational techniques and establishes a risk-based classification system to govern their use.³¹ These divergent approaches highlight the difficulty of crafting a unified definition of AI that is both technically precise and legally workable across jurisdictions, especially as new categories of risks such as opacity, hallucination and drift complicate traditional evaluation models.³²

In practice, AI in healthcare more often augments rather than replaces clinical judgment.³³ Yet, as decision-support tools become more sophisticated, questions arise regarding the scope of clinician responsibility, the contours of algorithmic

27 See Suchin Gururangan et al., *Hallucinations in Neural Machine Translation*, Association for Computational Linguistics 2018; Emily M. Bender et al., *On the Dangers of Stochastic Parrots: Can Language Models Be Too Big?*, ACM Conference on Fairness, Accountability, & Transparency 610, 613-16 (2021).

28 See Timnit Gebru et al., *Datasheets for Datasets*, 3 Comm. ACM 107, 110 (2021); Christoph Molnar, *Interpretable Machine Learning* 105-08 (2022) (describing concept drift and data drift).

29 See Ziad Obermeyer et al., *Dissecting Racial Bias in an Algorithm Used to Manage the Health of Populations*, 366 Science 447, 447-48 (2019); Cresencia Felicity Par et al., *Bias and Fairness in Artificial Intelligence in Health Care*, 326 JAMA 1926, 1927-28 (2021).

30 U.S. Food & Drug Administration, *Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning-Based Software as a Medical Device* 6-10 (2019).

31 Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June 2024 (Artificial Intelligence Act).

32 See U.S. Food & Drug Administration, *supra* note 16, at 3-6 (referencing challenges posed by adaptive learning and opacity).

33 See Robert M. Wachter & Jeff Goldsmith, *Artificial Intelligence in Health Care: Navigating the ‘Automation–Augmentation’ Paradox*, 45 Health Aff. 552, 552-55 (2023).

accountability and the adequacy of existing privacy protections. Because AI systems require large, diverse and often sensitive datasets, their deployment implicates issues of patient privacy, data ownership, secondary uses of health data and the potential for algorithmic inference beyond the information originally collected.³⁴ The introduction of black-box systems and adaptive learning further complicates these concerns, as opaque or evolving models challenge established norms for documentation, informed consent and regulatory approval.³⁵

Thus, defining AI in the healthcare setting is not merely a descriptive exercise as it shapes the legal landscape. How regulators classify AI tools influences the applicability of statutory privacy frameworks such as the HIPAA,³⁶ state health privacy laws,³⁷ FDA oversight mechanisms³⁸ and emerging AI-specific legislative schemes.³⁹ As subsequent sections of this article explore, the definitional contours of AI play a central role in identifying regulatory gaps, privacy risks and ethical tensions arising from the integration of AI into modern medical practice.

II Federal Laws

1 Health Insurance Portability and Accountability Act (HIPAA)

Health Insurance Portability and Accountability Act of 1996 (HIPAA)⁴⁰ remains the cornerstone of health information privacy law in the United States.⁴¹ HIPAA Privacy Rule governs the use and disclosure of “protected health information” (PHI) by covered entities, health plans, healthcare clearinghouses and healthcare providers conducting certain transactions electronically and their business associates.⁴² However, HIPAA’s entity-based structure limits its reach. It does not apply to many emerging data ecosystems central to AI in healthcare such as mobile health applications, consumer genetic testing platforms or fitness trackers that collect physiological data.⁴³ Nor does it fully address secondary uses of de-identified or aggregated data, which form the backbone of many machine-learning systems.⁴⁴ As health systems and vendors stitch together data streams from EHRs, patient portals, third-party apps, wearables and consumer platforms, privacy guarantees vary widely depending on who holds the data and where the patient resides.⁴⁵

34 See Sandra Wachter & Brent Mittelstadt, *A Right to Reasonable Inferences: Re-Thinking Data Protection Law in the Age of Big Data and AI*, Colum. Bus. L. Rev. 494, 498-505 (2019).

35 See Sara Gerke, Timo Minssen & Glenn Cohen, *Ethical and Legal Challenges of Artificial Intelligence-Driven Healthcare*, 25 AIH 295, 301-04 (2020).

36 Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104-191, 110 Stat. 1936.

37 See, e.g., Cal. Civ. Code §§ 1798.100-1798.199 (California Consumer Privacy Act); Wash. Rev. Code § 19.373 (My Health My Data Act).

38 See U.S. Food & Drug Administration, *supra* note 16, at 3-4.

39 See Algorithmic Accountability Act of 2023, S. 2892, 118th Cong. (2023).

40 Health Insurance Portability and Accountability Act of 1996, § 262(a), 42 U.S.C.A. § 1320d.

41 Health Insurance Portability and Accountability Act of 1996, Pub. L. No. 104-191, 110 Stat. 1936.

42 45 C.F.R. § 160.103 (defining “covered entities”).

43 U.S. Department of Health and Human Services, *supra* note 5.

44 *Id.*

45 U.S. Department of Health and Human Services, *supra* note 5.

HIPAA was enacted in an era that predates algorithmic inference. Even de-identified data may be re-identified through linkage attacks or pattern recognition, exposing latent privacy risks.⁴⁶ The statute also lacks provisions requiring algorithmic transparency, bias auditing or accountability, concepts that underpin contemporary ethical frameworks for responsible AI.⁴⁷ How state privacy laws regulate the use of patient data in AI applications, highlighting the variation across jurisdictions and the emerging trends towards comprehensive privacy statutes.

Because HIPAA provides only limited guidance for privacy protection, states have increasingly legislated in this space.

2 Federal Trade Commission (FTC) Regulation of AI

Although the FTC is not a healthcare regulator, it plays an increasingly central role in the governance of artificial intelligence, particularly in contexts involving consumer data, deceptive claims about AI capabilities, discriminatory algorithmic outcomes and unfair data practices. The agency's authority derives primarily from § 5 of the Federal Trade Commission Act, which prohibits "unfair or deceptive acts or practices in or affecting commerce".⁴⁸ This broad mandate enables the FTC to pursue oversight of AI systems used by hospitals, digital health companies, telemedicine platforms, medical device manufacturers, data brokers and third-party vendors operating outside HIPAA's coverage. As a result, the FTC is emerging as a key federal actor in regulating AI-enabled health technologies and protecting consumer privacy where sectoral statutes fall short.

The FTC has repeatedly signalled that the development, marketing and deployment of AI systems fall squarely within its jurisdiction. The agency treats misleading claims about an AI model's accuracy, safety or performance as potentially deceptive under § 5, and it has warned companies that unsubstantiated claims of algorithmic superiority may constitute prohibited commercial conduct.⁴⁹ The FTC also views certain uses of data such as training AI models on unlawfully obtained or deceptively collected information as an unfair practice. Notably, in *Everalbum*, the FTC's first major AI case, the agency required the destruction not only of the unlawfully retained biometric data but also of algorithms developed using those data, establishing a significant remedy for AI misuse.⁵⁰ This "algorithmic disgorgement" remedy signals the agency's willingness to intervene directly in the lifecycle of machine-learning systems.

The FTC has also indicated that it views algorithmic discrimination as a form of "unfair" practice, particularly where companies fail to assess or mitigate

46 Paul Ohm, *Broken Promises of Privacy: Responding to the Surprising Failure of Anonymization*, 57 UCLA L. Rev. 1701, 1716-27 (2010).

47 *Id.*

48 Federal Trade Commission Act § 5, 15 U.S.C. § 45(a).

49 Shannon Boettjer, *Unpacking the FTC's Warning To "Keep Your AI Claims In Check"*, N.Y. BUS. L. & COM. LITIG. BLOG (Apr. 11, 2023), <https://www.nybusinesslaw.com/2023/04/unpacking-the-ftcs-warning-to-keep-your-ai-claims-in-check/>.

50 *In re Everalbum, Inc.*, No. C-4749, 2021 WL 268555 (F.T.C. 11 January 2021).

foreseeable harms. In its 2021 business guidance, “Aiming for Truth, Fairness, and Equity in Your Company’s Use of AI”, the FTC urged developers and deployers to conduct pre- and post-deployment testing, ensure representativeness in datasets and avoid models that disproportionately harm marginalized groups.⁵¹ Though the FTC does not enforce civil rights statutes directly, it treats discriminatory outcomes that stem from deceptive data practices or unreasonable failure to adopt safeguards as within its jurisdiction.

This approach is especially relevant to digital health tools, wearable medical devices and AI-driven symptom checkers, where biased outputs may misdiagnose or underserve particular populations. Unlike HIPAA, which focuses on data privacy and security rather than outcomes, the FTC explicitly frames algorithmic fairness as a consumer protection issue.

The FTC’s role is particularly important in healthcare because many AI developers, mobile health apps and digital therapeutics companies fall outside HIPAA’s narrow definition of “covered entities” or “business associates”. The FTC has increasingly used its authority over non-HIPAA health data, especially through the Health Breach Notification Rule (HBNR) to regulate health apps and AI-enabled platforms that handle sensitive information.⁵² Recent enforcement actions indicate that the FTC views sharing or using health data for targeted advertising or undisclosed training purposes as both deceptive and unfair.⁵³ This encompasses many AI training and model-improvement practices.

The FTC has also emphasized transparency and explainability as essential components of lawful AI deployment. The agency has warned companies against overstating the capabilities of AI systems or concealing meaningful limitations, and it has signalled that failure to disclose material facts about how AI systems operate may constitute deception. These concerns are especially salient for medical AI tools that purport to improve diagnostic accuracy, triage patients or optimize care management, areas where unsupported claims may mislead consumers or clinicians about safety or effectiveness.

3 Federal AI Policy and Pre-emption⁵⁴

Recent federal policy developments further illustrate the evolving governance landscape surrounding AI. In December 2025, President Donald Trump issued an executive order titled Ensuring a National Policy Framework for Artificial Intelligence, which directed federal agencies to review and potentially challenge state laws that conflict with national AI policy objectives. The order reflects growing fed-

51 Allen R. Bachman & Derek W. Kelley, *FTC Issues New Guidance, Warning That Bias in Artificial Intelligence Could Create Potential Liability for Enforcement Actions*, K&L GATES (Apr. 23, 2021), <https://www.klgates.com/thought-leadership/FTC-Issues-New-Guidance-Warning-That-Bias-In-Artificial-Intelligence-Could-Create-Potential-Liability-For-Enforcement-Actions-4-21-2021>.

52 See Health Breach Notification Rule, 16 C.F.R. pt. 318; see also Fed. Trade Comm’n, *Policy Statement on the Health Breach Notification Rule* (2021).

53 See, e.g., *In re GoodRx Holdings, Inc.*, No. C-4793 (F.T.C. 1 February 2023) (alleging unauthorized disclosure of consumer health data to advertising platforms); *In re BetterHelp, Inc.*, No. C-4802 (F.T.C. 2 March 2023).

54 The White House, *supra* note 11.

eral interest in preventing regulatory fragmentation across states that may impede technological innovation or national competitiveness in AI development. At the same time, federal intervention raises important questions regarding the balance between national policy coordination and states' traditional role in regulating health data privacy and consumer protection. In the healthcare context, such federal efforts to limit state regulation could significantly influence how states develop privacy protections for AI-driven health technologies and may reshape the existing patchwork governance structure described above.

The executive order represents a significant shift towards federal centralization of AI governance. At its core, the order seeks to impose broad and strict limitations on state-level regulation of AI, premised on the view that fragmented regulatory regimes hinder innovation and undermine the United States' global competitiveness. The order characterizes the current landscape as a "patchwork of 50 different regulatory regimes", which it argues makes compliance more difficult, particularly for start-ups, and, in some cases, results in impermissible extraterritorial regulation that impinges on interstate commerce.

To operationalize these goals, the executive order directs the Attorney General to establish an AI Litigation Task Force with broad authority to challenge state AI laws that may unconstitutionally burden interstate commerce. In addition, it grants the Secretary of Commerce and the FTC the authority to withhold federal funding from states that fail to align with the order's directives. These enforcement mechanisms signal a clear federal intent to deter expansive state-level AI regulation and promote a more uniform national framework.

Despite its sweeping language, the executive order does not automatically invalidate existing state AI laws. Current and pending state and local regulations remain enforceable unless and until they are enjoined by a court or pre-empted by federal legislation. Nevertheless, the order exerts significant practical pressure on states to reconsider their regulatory approaches.

One key lever is federal funding. States risk losing access to programmes such as the Broadband Equity, Access, and Deployment (BEAD) Program if their AI laws are deemed "onerous". This creates a strong incentive for states to scale back or avoid aggressive AI regulation, including in sensitive sectors such as healthcare.

If the order is actively enforced, state healthcare laws governing AI, such as those addressing algorithmic bias, patient privacy or clinical decision-making, may face legal challenges. This could result in diminished state-level protections, leaving patients increasingly reliant on federal oversight, which may be slower to adapt or less tailored to local concerns.

At the same time, private sector stakeholders may welcome the shift towards a streamlined federal framework. A uniform regulatory environment could reduce compliance costs and uncertainty, encouraging more rapid adoption of AI technologies. In the healthcare context, this may accelerate the deployment of AI tools in clinical settings, support advances in precision medicine and spur innovation in research and development.

However, this potential for accelerated innovation must be weighed against the risks of reduced regulatory diversity and experimentation at the state level. As states

recalibrate their strategies in response to federal pressure, the long-term balance between innovation, patient safety and regulatory oversight remains uncertain.

III State Laws

Because federal privacy law provides only a limited baseline for the protection of health information, states have increasingly become primary regulators of health-adjacent data implicated in AI. This article focuses specifically on California, Colorado and Washington because these jurisdictions illustrate distinct and influential approaches to regulating non-HIPAA health data. California's Privacy Rights Act represents the most mature comprehensive consumer privacy regime in the United States; Colorado's Privacy Act reflects a risk-based model influenced by European data protection principles, and Washington's My Health My Data Act uniquely defines and protects "consumer health data" regardless of whether it originates in a traditional healthcare setting. Together, these statutes capture the range of state-level strategies shaping AI-driven health data governance and highlight the compliance challenges and normative inconsistencies facing healthcare organizations and AI developers operating across state lines.

1 California Privacy Rights Act (CPRA)⁵⁵

The CPRA became effective on 1 January 2023. This act expands the California Consumer Privacy Act (CCPA) by granting consumers the right to limit the use of "sensitive personal information", including health data, even when HIPAA does not apply.⁵⁶ The CPRA also creates the California Privacy Protection Agency to enforce compliance and issue regulations, signalling a shift towards dedicated privacy enforcement.⁵⁷ It adds opt-out and sensitive-data limits that can still touch health-adjacent data even where HIPAA exemptions apply.

2 Colorado's Privacy Act⁵⁸

Colorado's Privacy Act modelled partly on the European Union's General Data Protection Regulation (GDPR). It exempts HIPAA-regulated data but covers non-HIPAA consumer health data.⁵⁹ However, it does not automatically cover all health-related consumer data outside HIPAA contexts. Its provisions require data protection assessments for processing that presents a "heightened risk of harm", which may encompass algorithmic profiling in healthcare.⁶⁰

Colorado's Privacy Act⁶¹ expressly regulates "automated decision-making" involving personal data and grants consumers a qualified right to opt out of such

55 California Privacy Rights Act of 2020, 'Text of the CPRA' (unannotated version), <https://www.caprivacy.org/cpra-text/>.

56 Cal. Civ. Code § 1798.100 et seq. § 1798.121.

57 California Privacy Protection Agency, *Regulations Overview* (2023).

58 Colorado Privacy Act (summary/exemption context).

59 Colo. Rev. Stat. § 6-1-1301 et seq. (2023).

60 Colo. Rev. Stat. § 6-1-1309 et seq. (2023).

61 Colo. Rev. Stat. §§ 6-1-1301 to 6-1-1313.

processing. Controllers must provide notice when automated decision-making produces legal or similarly significant effects and must honour opt-out requests when decisions are made solely through automated means. However, the statute permits controllers to deny opt-out requests if a human is meaningfully involved in the review of the decision. When an opt-out request is denied, controllers are required to provide the consumer with specific information regarding the logic involved in the automated processing and the data used. As a result, Colorado's framework offers conditional transparency rather than categorical control, allowing many AI-driven systems to remain operative so long as a human-in-the-loop is maintained.

3 *Washington's My Health My Data Act (MHMD Act)*⁶²

Washington's MHMD Act goes one step further compared to other state laws by defining "consumer health data" broadly. It includes any personal information that can be linked to a person's physical or mental health status, regardless of whether it originates in a healthcare setting.⁶³ The Act mandates affirmative consent for collection and sharing, prohibits geofencing near healthcare facilities and applies to out-of-state entities targeting Washington consumers.⁶⁴ These state-level differences create a compliance puzzle for AI developers and health systems operating nationally.⁶⁵ Expressly targets personal "consumer health data" outside HIPAA's ambit by imposing consent, deletion and geofencing prohibitions aimed at closing post-Dobbs surveillance gaps.⁶⁶

IV *Key Ethical Frameworks*

As AI tools increasingly shape clinical decision-making and population health management, numerous professional and governmental bodies have developed ethical frameworks to guide their responsible deployment. Although these frameworks do not carry the force of law, they play a crucial role in shaping institutional policy, professional standards of care and, indirectly, regulatory expectations. Ethical guidance has proliferated faster than binding legal standards, and understanding these frameworks is essential to mapping the normative landscape in which healthcare AI operates. This section analyses four influential sources: the HHS, the AMA, the WHO and FDA.

Each offers distinct principles but collectively reflects broad commitments to transparency, accountability, equity, autonomy and safety. HHS, the AMA and the WHO offer ethical frameworks, while the FDA provides a regulatory framework. The

⁶² Wash. Rev. Code Ann. § 19.373 (2023).

⁶³ Wash. Rev. Code Ann. § 19.373 (2023).

⁶⁴ Wash. Rev. Code Ann. § 19.373.030 (2023).

⁶⁵ Reuters, *supra* note 4.

⁶⁶ Washington State Attorney General, *Protecting Washingtonians' Personal Health Data and Privacy: My Health My Data Act*, <https://www.atg.wa.gov/protecting-washingtonians-personal-health-data-and-privacy>.

similarities among these frameworks demonstrate a growing international ethical consensus, even as gaps remain between ethical ideals and enforceable legal obligations.

1 *U.S. Department of Health and Human Services (HHS)*

HHS' most recent guidance promotes the development of AI systems that are "ethical, effective, and secure", emphasizing robust governance structures, explainability of algorithmic outputs and clear lines of accountability throughout the AI lifecycle.⁶⁷ HHS encourages health systems deploying AI to develop internal risk management processes, perform continuous monitoring and ensure meaningful human oversight.⁶⁸ The agency highlights that AI must enhance, not undermine, the trust central to healthcare delivery. Notably, its emphasis on explainability and accountability reflects concerns that opacity in machine-learning models could compromise informed consent or obscure responsibility for clinical decisions.

While HHS's guidance is largely advisory, it intersects with several existing legal duties. For example, the call for secure AI systems aligns with HIPAA's Security Rule requirements to ensure the confidentiality, integrity and availability of electronic protected health information (ePHI).⁶⁹ Likewise, the expectation of meaningful human oversight reinforces professional standards of care and may influence negligence inquiries concerning clinician reliance on AI decision-support systems. At the same time, HHS acknowledges that formal regulation of adaptive machine-learning systems is still emerging, leaving a gap between ethical aspirations and enforceable obligations.

2 *American Medical Association (AMA)*

The AMA's policy on "augmented intelligence" reflects a professional ethics perspective grounded in the physician-patient relationship. The AMA urges that AI systems used in medicine must be transparent, clinically validated, equitable in their performance across populations and designed to preserve physician autonomy and patient-centred care.⁷⁰ The organization uses the term "augmented intelligence" deliberately, signalling that AI should enhance rather than replace clinical judgment.⁷¹ It also emphasizes the importance of understanding training datasets, mitigating algorithmic bias and ensuring that physicians retain ultimate responsibility for patient care decisions.

The AMA's principles align with core legal concepts such as informed consent, competency-based medical practice and non-discrimination. For example, transparency supports a patient's ability to understand how AI contributes to diagnosis or treatment, a key component of informed consent doctrine. Likewise, the AMA's focus on equity resonates with federal civil rights protections, including Title VI of the Civil Rights Act and Section 1557 of the Affordable Care Act, which

67 U.S. Department of Health and Human Services, *Trustworthy AI Playbook* 3-5 (2023).

68 *Id.*, at 6-10.

69 45 C.F.R. §§ 164.302-318 (HIPAA Security Rule).

70 Am. Med. Ass'n, *Augmented Intelligence in Health Care* H-480.940 (2023).

71 *Id.*

prohibit discriminatory impacts in federally funded health programmes.⁷² Yet AMA guidance often goes beyond existing law, particularly in addressing structural bias embedded in training datasets; current legal frameworks do not expressly require algorithmic auditing or impact assessments in clinical settings.

3 World Health Organization (WHO)

The WHO's *Ethics and Governance of Artificial Intelligence for Health* offers one of the most comprehensive global frameworks, articulating six guiding principles, which includes (1) protecting autonomy; (2) promoting well-being and safety; (3) ensuring transparency; (4) fostering accountability; (5) ensuring inclusiveness and equity; and (6) promoting responsive and sustainable AI.⁷³ The WHO emphasizes global health equity, noting that AI systems risk reinforcing existing disparities if deployed without attention to access, cultural context or uneven data representation. It also stresses the need for participatory development processes involving diverse communities and stakeholders.

These principles often exceed the scope of domestic legal regimes. For example, the WHO's focus on sustainability and long-term monitoring has no clear analogue in U.S. law, which still lacks statutory mandates for ongoing post-deployment evaluation of AI in healthcare. Nonetheless, WHO principles complement emerging regulatory impulses, such as the FDA's exploration of total product lifecycle oversight for adaptive AI systems, and inform debates about the global harmonization of AI governance standards.

4 U.S. Food and Drug Administration (FDA)

Although the FDA does not issue "ethics" guidelines in the way HHS, AMA or WHO do, it has developed a *regulatory* framework that incorporates several ethical principles including transparency, accountability, safety and continuous oversight into its approach to artificial intelligence and machine-learning-based software as a medical device (AI/ML SaMD). Central to the FDA's policy development is the recognition that adaptive algorithms pose unique challenges because their performance may evolve after deployment.

The FDA's 2019 *Proposed Regulatory Framework for Modifications to AI/ML-Based Software as a Medical Device* introduced the concept of a "predetermined change control plan", describing in advance (1) the types of modifications a manufacturer intends to make as the model learns over time and (2) the methods for assessing the safety and effectiveness of such changes.⁷⁴ This approach reflects an effort to balance innovation with patient safety by ensuring that adaptive AI systems remain subject to ongoing oversight rather than one-time approval.

In 2021, the FDA expanded its vision in the AI/ML SaMD Action Plan, which emphasized transparency, real-world performance monitoring, algorithmic bias

72 Civil Rights Act of 1964, Pub. L. No. 88-352, § 601, 78 Stat. 241, 252 (codified at 42 U.S.C. § 2000d); Patient Protection and Affordable Care Act, Pub. L. No. 111-148, § 1557, 124 Stat. 119, 260 (2010).

73 World Health Organization, *Ethics and Governance of Artificial Intelligence for Health* (2021).

74 U.S. Food & Drug Administration, *supra* note 30, at 5-9.

mitigation and patient-centred communication.⁷⁵ The FDA signalled that trustworthy AI requires “a total product lifecycle” model of regulation, integrating pre-market review, post-market surveillance, continuous risk assessment and mechanisms for updating patients and clinicians about changes to system performance.⁷⁶

The FDA’s framework intersects with ethical concerns raised by both domestic and international actors. Its focus on real-world performance speaks to accountability; its emphasis on bias relates to equity; and its lifecycle-based governance align with WHO’s call for responsive and sustainable AI systems. However, unlike AMA and WHO frameworks, the FDA grounds its guidance primarily in statutory and regulatory obligations rather than normative professional duties. Conversely, FDA guidance stops short of addressing autonomy and informed consent in the depth seen in HHS and WHO frameworks; these issues remain governed by existing clinical and institutional policies rather than device-specific regulation.

V. *Recent Scholarship on AI Governance and Health Data Privacy*

Recent legal scholarship has similarly highlighted the structural gaps between existing privacy law and the governance challenges posed by AI in healthcare. Daniel Solove argues that AI does not create entirely new privacy problems but instead amplifies longstanding issues involving data collection, inference, surveillance and automated decision-making, exposing the inadequacy of current privacy frameworks built around individual notice and consent.⁷⁷ Frank Pasquale likewise critiques control-based privacy models, demonstrating how health-related data increasingly circulates through data brokers and reputational analytics systems outside the traditional HIPAA framework, raising risks of discrimination, stigmatization and opaque decision-making.⁷⁸ Fazal Khan extends this discussion specifically to clinical AI governance, examining legal issues including liability, informed consent, privacy protection, professional licensure and FDA oversight.⁷⁹ Khan emphasizes that the opacity of algorithmic “black-box” systems and the potential for regulatory capture require stronger accountability mechanisms and ongoing oversight. Together, these scholars underscore that existing legal structures are poorly equipped to govern the emerging ecosystem of AI-driven health data and that more comprehensive regulatory approaches are necessary to align innovation with ethical principles.

⁷⁵ U.S. Food & Drug Administration, *supra* note 16, 3-4.

⁷⁶ *Id.*, at 5-7.

⁷⁷ Daniel J. Solove, *Artificial Intelligence and Privacy*, 77 Fla. L. Rev. 1 (2025).

⁷⁸ Frank Pasquale, *Redescribing Health Privacy: The Importance of Information Policy*, 14 Hous. J. Health L. & Pol’y 95 (2014).

⁷⁹ Fazal Khan, *Regulating the Revolution: A Legal Roadmap to Optimizing AI in Healthcare*, 25 Minn. J.L. Sci. & Tech. 49 (2023).


VI. Summary Points

Across HHS, AMA, WHO and FDA frameworks, several normative commitments consistently appear promoting transparency, fostering accountability, preventing discrimination and protecting autonomy. These shared principles suggest a de facto ethical baseline for healthcare AI. However, tensions remain between these ethical recommendations and binding legal obligations.

As shown in Figure 1, AI practice in healthcare is shaped by overlapping layers of binding law and nonbinding ethical guidance. Ethical frameworks often impose proactive duties such as bias auditing, sustainability planning or ensuring global fairness that current legal standards do not yet require. Conversely, legal frameworks may impose strict privacy or security obligations that ethical statements reference only at a high level. The evolving interplay between ethics and law underscores the need for a governance model capable of operationalizing ethical commitments through enforceable regulatory mechanisms.

FTC oversight complements other ethical and regulatory frameworks by creating enforceable boundaries around data practices. While HIPAA provides rules governing PHI held by covered entities, the FTC fills gaps for non-HIPAA data holders, such as wellness apps, AI research platforms and consumer-facing medical devices. Moreover, the FTC's explicit attention to fairness, discrimination and deceptive claims parallels ethical principles articulated by HHS, AMA and WHO, albeit through an enforcement-driven lens. In this way, the FTC serves as a bridge between ethical AI norms and binding legal oversight, especially in parts of the AI ecosystem not reached by healthcare-specific laws.

Figure 1 Layered Governance of AI and Privacy in Healthcare.



Author's illustration based on HIPAA, state privacy statutes and ethical frameworks issued by HHS, AMA, WHO and the FDA.

C Analysis

Although ethical frameworks from HHS, AMA, WHO and the FDA articulate ambitious principles for responsible AI in healthcare and their relationship to binding privacy laws, the analysis reveals significant gaps particularly with the HIPAA and emerging state privacy statutes. These divergences demonstrate that while ethical guidance is expansive and forward-looking, legal obligations remain narrower, more rigid and focused on data stewardship rather than holistic governance of AI systems.

I Alignment: Where Ethical Frameworks Reinforce Existing Legal Duties

Several commitments embedded in major AI ethical frameworks parallel obligations already present in HIPAA and state privacy laws.

1 Data Security and Integrity

HIPAA's Security Rule requires covered entities and business associates to implement administrative, physical and technical safeguards to ensure the confidentiality, integrity and availability of the PHI.⁸⁰ Ethical frameworks similarly emphasize secure AI design, continuous monitoring and risk management, particularly HHS and FDA, both of which link trustworthy AI to robust security controls.⁸¹ These requirements align closely in that ethical guidance reinforces compliance with existing law by encouraging advanced safeguards appropriate for algorithmic systems, such as monitored model updates and cyber-resilience testing.

2 Minimization of Bias and Non-discrimination

State privacy statutes such as the CCPA and the Washington My Health My Data Act prohibit discriminatory uses of health data in certain contexts.⁸² Ethical frameworks from AMA and WHO likewise stress equity, inclusiveness and the mitigation of algorithmic bias. While HIPAA does not explicitly address discrimination, civil rights laws like Title VI and ACA § 1557 prohibit discriminatory outcomes in healthcare programmes.⁸³ AMA's insistence that AI must not exacerbate health disparities reflects this legal constraint, creating a point of substantive convergence.

⁸⁰ 45 C.F.R. §§ 164.302-318.

⁸¹ U.S. Department of Health and Human Services, *supra* note 67, at 3-10; U.S. Food & Drug Administration, *supra* note 16, at 3-6.

⁸² See Cal. Civ. Code §§ 1798.100-199; Wash. Rev. Code § 19.373.030(1).

⁸³ 42 U.S.C. § 2000d; 42 U.S.C. § 18116.

3 *Accountability and Governance*

Ethical frameworks uniformly call for systems of accountability across the AI lifecycle. HIPAA's Privacy Rule similarly places responsibility on covered entities for the lawful handling, disclosure and safeguarding of PHI.⁸⁴ State privacy statutes add accountability mechanisms such as consumer access rights, notice obligations and enforcement through attorneys general. FDA's lifecycle regulatory model aligns with these legal expectations by conditioning AI deployment on post-market monitoring and documentation.⁸⁵

II. *Areas of Divergence: Ethical Ambitions Beyond Current Legal Requirements*

Despite overlaps, ethical frameworks routinely exceed what HIPAA and state statutes require.

1 *Transparency and Explainability*

HIPAA does not mandate that covered entities explain how PHI is used to generate algorithmic recommendations; nor does it require model-level transparency.⁸⁶ Ethical frameworks, however, especially WHO, AMA and HHS emphasize explainability as essential to patient autonomy and clinician trust. Many state privacy laws contain transparency requirements, but these usually concern data practices (collection, sale, sharing), not the internal logic of AI systems.⁸⁷ Thus, ethical transparency exceeds legal transparency.

2 *Autonomy, Informed Consent and Algorithmic Understanding*

Ethical frameworks highlight autonomy as a central principle, particularly WHO and AMA, which argue that patients should understand when AI is used and how it influences care. HIPAA, by contrast, does not create a general requirement to disclose AI use in treatment; nor does it impose duties regarding algorithmic explainability in informed consent. State privacy statutes similarly do not extend informed consent doctrine to algorithmic decision-making in clinical settings. Although some laws, such as Colorado's Privacy Act,⁸⁸ require notice and provide a conditional right to opt out of certain forms of automated decision-making, these protections apply only when decisions are made solely through automated means and may be displaced by human review. As a result, the ethical requirement for consent-aligned transparency remains largely aspirational.

84 45 C.F.R. §§ 164.500-534.

85 FDA, *AI/ML SaMD Action Plan*, Park, *supra* note 2, at 5-7.

86 See 45 C.F.R. §§ 164.520 (Privacy Rule notice requirements), 164.506 (permitted uses), 164.524 (right of access).

87 See Cal. Civ. Code § 1798.110 (requiring transparency about categories of data collected and purposes).

88 Colo. Rev. Stat. §§ 6-1-1301 to 6-1-1313.

3 *Post-Deployment Monitoring and Sustainability*

HHS and WHO encourage extensive post-deployment model monitoring, bias auditing and continuous validation. FDA's regulatory approach partially incorporates this for AI/ML SaMD through its total product lifecycle model.⁸⁹ HIPAA and state laws do not impose comparable duties. Their obligations end at safeguarding data, not overseeing algorithmic performance. As AI systems evolve over time, this legal gap becomes increasingly consequential.

4 *Equity Beyond Non-Discrimination*

State laws and civil rights statutes prohibit discriminatory impacts, but ethical frameworks call for proactive equity promotion requiring designers and deployers to consider structural determinants of health and data representativeness. WHO's inclusiveness principle and AMA's equity commitments extend well beyond what U.S. law mandates. Legal requirements focus on preventing discrimination; ethical frameworks aim to correct inequity.

III *Regulatory Gaps: Ethical Obligations Without Legal Counterparts*

The gap between ethical guidance and binding legal requirements becomes most visible when examining the expectations placed on algorithmic transparency. Ethical frameworks consistently emphasize the need for explainability, requiring clinicians and patients to understand, at least in broad terms, how an AI system generates recommendations. Yet neither HIPAA nor most state privacy statutes impose any duty on healthcare organizations to disclose the logic or functioning of AI tools. Transparency obligations under U.S. privacy laws focus on data practices, not algorithmic processes. As a result, ethical expectations for explainability exist largely without legal counterparts, leaving patients and clinicians dependent on institutional discretion rather than regulatory guarantees.

Similarly, a significant gap exists in the oversight of adaptive or continuously learning models. Ethical frameworks, especially those issued by HHS, WHO and incorporated into FDA's lifecycle-based regulatory vision, call for ongoing monitoring, periodic evaluation and continuous risk assessment to ensure that AI systems remain safe and effective over time.

HIPAA and state privacy statutes, however, are static. They regulate the handling of data but do not impose duties to evaluate the performance of algorithms that rely on those data. This creates a structural disjunction in which ethical advisers promote dynamic governance while legal frameworks remain tied to traditional, fixed compliance models.

A third gap emerges around equity and representativeness. Ethical frameworks require proactive measures to ensure that AI systems do not replicate or exacerbate existing health disparities, such as by mandating diverse training data or encouraging community engagement in system design. Legal requirements, however, remain focused on preventing discriminatory outcomes rather than ensuring equitable development. Civil rights laws prohibit discrimination, but

89 FDA, *AI/ML SaMD Action Plan*, Park, *supra* note 2, at 3-5.

they do not mandate the forward-looking fairness assessments or bias audits contemplated by AMA and WHO guidance. Thus, the ethical commitment to equity significantly outpaces current legal mandates.

Finally, the divergence between ethical principles and legal requirements is evident in matters of autonomy and consent. Ethical frameworks stress that patients should be informed when AI tools are used in clinical decision-making and should understand how those tools influence their care. HIPAA's consent provisions, by contrast, apply only to uses and disclosures of PHI and do not require disclosure of algorithmic involvement in treatment. State privacy statutes also stop short of extending informed consent doctrines to the use of AI. As a result, patients may unknowingly receive algorithmically informed care even though ethical frameworks uniformly treat notice and comprehension as foundational to autonomy.

Collectively, these regulatory gaps underscore that ethical frameworks envision a more expansive set of obligations rooted in transparency, equity, continuous oversight and autonomy than those currently imposed by U.S. privacy laws. Bridging these divergences will require either substantial statutory reform or the adoption of regulatory mechanisms that more closely mirror the principles articulated by healthcare ethics organizations and international bodies.

IV Proposed Policy Recommendations

A harmonized definition would ensure that privacy protections attach to the data itself, rather than to the institution holding it, thereby eliminating loopholes that permit secondary or inferential uses without consent. Requiring AI impact assessments would impose proactive accountability, compelling organizations to evaluate risks related to privacy, fairness and bias prior to deployment. Incorporating ethical standards into law would further transform abstract principles into enforceable obligations, bridging the gap between voluntary professional norms and formal regulatory oversight.

The following policy recommendations are proposed:

- Adopting a model definition of consumer health data applicable across jurisdictions;
- Mandating AI impact assessments and transparency reporting for healthcare AI systems;
- Embedding ethical norms within statutory and regulatory frameworks to create a more coherent governance regime.

Taken together, these measures would enable policymakers to construct a governance ecosystem in which AI innovation coexists with human dignity and public trust. Law would provide the structural foundation of enforceable rights and duties, while ethics would offer normative guidance, ensuring that algorithmic efficiency does not eclipse the foundational commitments of medicine – respect for autonomy, beneficence and justice.

Ultimately, the future of AI in healthcare will depend not only on technological capability but also on whether its deployment aligns with the core values of the

healing professions. Harmonizing privacy law with ethical governance is therefore not merely a regulatory necessity; it is a moral imperative in an era where machines increasingly participate in human care.

D. Conclusion

The integration of AI into healthcare exposes a widening gap between ethical aspiration and legal reality. Despite the transformative potential of AI for improving diagnosis, treatment and administrative efficiency, the governing legal framework has not evolved at a comparable pace. HIPAA's narrow scope is limited to covered entities and their business associates and thus fails to encompass the broad ecosystem of health data now fuelling algorithmic systems.⁹⁰ AI development increasingly relies on datasets collected outside traditional healthcare settings, such as fitness apps, social media platforms or consumer genomics companies. These forms of data, though predictive of health status, fall outside HIPAA's jurisdiction and remain inconsistently regulated by state law.

The uneven reach of state privacy statutes compounds this problem. While some states, such as California, Colorado and Washington, have enacted privacy laws addressing non-HIPAA health data, these regimes differ sharply in scope, terminology and enforcement mechanisms. The CPRA extends opt-out rights and limits the use of "sensitive personal information", yet its applicability to medical AI depends on whether an entity qualifies as a business subject to the Act. Colorado's Privacy Act and Washington's MHMD Act go further in recognizing "consumer health data" as a distinct category, but they diverge in definitions, consent standards and exemptions.

For AI developers and healthcare organizations that operate nationally, this patchwork creates both compliance uncertainty and ethical inconsistency. Identical uses of patient data may be lawful in one state but restricted in another. As a result, core aspects of AI governance, algorithmic transparency, bias mitigation and cross-context data use remain largely unregulated, even as these features determine whether AI advances equity or exacerbates disparities in care.

Within this regulatory vacuum, however, there still exist ethical frameworks promulgated by the HHS, the AMA, the WHO and the FDA that articulate principles of transparency, accountability, autonomy and fairness, norms intended to guide responsible AI adoption in healthcare.

The AMA emphasizes that AI must "enhance, not replace" physician judgment. The WHO underscores that AI should promote human well-being and equity. HHS stresses that health AI must remain ethical, effective and secure. Collectively, these frameworks fill part of the conceptual void left by fragmented law. Yet their lack of legal enforceability limits their practical effect. Ethical guidance can inspire

90 See 45 C.F.R. § 160.103 (defining "business associate").

professional best practices, but, without statutory backing or regulatory teeth, adherence is voluntary.

This reliance on “soft law” allows significant variance in implementation and weakens patient trust, precisely the opposite of what ethics aims to achieve. To close the gap between principle and practice, legal reform must integrate ethical commitments into enforceable standards. Trump’s December 2025 executive order may serve as a starting point for accelerating legal reform and advancing federal AI laws in healthcare, or it may instead introduce further uncertainty. We are entering a transformative era in the regulation of AI in healthcare.