

October 17, 2023

Privacy Concerns at the Intersection of Generative AI and Healthcare

By Kirk Nahra and Ali Jessani

Share:



Artificial intelligence that can create new texts, images, and other content (or “generative AI”) is revolutionizing every industry, and healthcare is no exception. Doctors are experimenting with using generative AI to improve their patient communications ¹ and to save them time uploading new information to patient records. ² Surgeons (and medical students) are using generative AI to create virtual patient simulations, ³ which allow them to improve their capabilities by practicing treatments and procedures in a risk-free environment. Industry experts predict ⁴ that generative AI could help streamline health insurance claims processing and improve the health insurance industry’s efficiency overall. All of these tools have the potential to transform healthcare services as they become further utilized by the industry.

At the same time, regulators are also paying attention to the privacy and other data-related risks associated with generative AI. This is largely due to the fact that, unlike other AI products that are used internally by companies to improve their products and services, generative AI tools are widely available to consumers, thereby implicating consumer protection concerns. From a data protection standpoint, regulators are concerned with both how generative AI models are trained and how they are used. On the training front, generative AI models rely on large amounts of data for their development, and there are questions as to exactly what data is being used for this purpose (including whether any information that identifies known individuals—or personal

information—is involved). In terms of use cases, both privacy and other regulators are particularly nervous about these generative AI models being used in an unfair or discriminatory manner and want industry to be transparent and accountable with regard to how they use these tools.

Generative AI and HIPAA

The privacy concerns with generative AI are especially complex in the healthcare context, largely due to how health data is regulated in the United States. For healthcare providers, health insurers, and the vendors they use (i.e., business associates), privacy obligations for patient data are generally governed by the Health Insurance Portability and Accountability Act and its implementing regulations (HIPAA). ⁵ HIPAA compliance is also relevant for the development and use of generative AI.

At the development stage, there are potential limitations as to what patient data HIPAA covered entities and their business associates are permitted to use and disclose in order to train their AI models. HIPAA requires covered entities to only use and disclose protected health information (PHI) for certain permitted purposes. These include for the patient's treatment, processing of payments, and the organization's healthcare operations purposes (among others). This last provision (healthcare operations) is defined broadly enough in the HIPAA rules to likely permit a healthcare provider to develop and use generative AI tools as part of offering their overall healthcare services. If a covered entity's use of PHI does not fall within a permitted purpose, it would need an authorization (i.e., consent) from patients in order to use or disclose their identifiable data, which would be extremely difficult or even impossible in the generative AI context (given how much data is needed to train these models). This issue is trickier for business associates, which may also be limited by their contracts in terms of how they can use identifiable patient data for their own purposes.

One way entities regulated under HIPAA can avoid needing a permitted purpose or a HIPAA authorization to train their AI models is by using deidentified data. Deidentified data falls outside the purview of the HIPAA rules and can be just as effective as identifiable patient data from an AI development standpoint. HIPAA does, however, have a specific deidentification standard that must be met before patient data can be considered deidentified ⁶ (it is generally more stringent than what is required under other privacy laws). If a company works for a healthcare provider as a business associate and wants to use deidentified patient data to train its own AI models, it faces additional challenges. Not only must the business associate meet HIPAA's deidentification standard, but it must also ensure that its business associate agreement with the covered entity expressly permits it to deidentify patient data. Given the large amounts of data needed to train and develop generative AI models, this deidentification right can be an important negotiation point between covered entities and their business associates.

HIPAA is also relevant for how healthcare providers and health insurers use generative AI tools. For example, if a healthcare provider regulated under HIPAA uses a chatbot to schedule appointments or to collect information about a patient's symptoms before their visit (or more generally to assist with telehealth services), the patient information collected by the chatbot would be subject to the HIPAA rules. (The same would be true if a health insurance company regulated under HIPAA used a chatbot to expedite its claims processing, for example.) Among other things, this means that the healthcare provider or health insurer in these scenarios would need to ensure that they have an appropriate business associate agreement in place with any third-party service that they use to offer these chatbot services. Additionally, any unauthorized access or disclosure of patient information processed by the chatbot would be subject to the HIPAA Breach Notification Rule (and require notice to the Department of Health and Human Services). Covered entities experimenting with generative AI should properly evaluate these risks.

The risks highlighted above are not exclusive to the development of new AI products; they are also relevant if entities regulated under HIPAA use existing (publicly available) generative AI tools. Healthcare providers, health insurers,

and any companies that process PHI as business associates should have policies and procedures in place that outline how their employees are permitted to use widely available generative AI tools (such as ChatGPT). Using these tools to process identifiable patient data may violate the HIPAA rules and lead to potential data security concerns, as there have already been a number of security issues associated with these tools. [7](#)

FTC Concerns

If a company processes health information but falls outside of the scope of HIPAA (such as companies that offer direct-to-consumer health services, including personal health record vendors (PHRs), wearables providers, and certain health-based mobile apps), there are other privacy rules that are relevant to the processing of personal information for generative AI purposes. Indeed, both the Federal Trade Commission (FTC) and state regulators are increasingly paying attention to health data that falls outside the purview of HIPAA.

The FTC has been particularly active as a privacy enforcer in recent months and is increasingly focusing on companies that process consumer health data (such as GoodRx [8](#) and BetterHelp [9](#)). The agency has been especially concerned with companies disclosing health data to advertisers without consumer consent. Most notably, the agency has been using the “unfairness” prong of Section 5 of the FTC Act [10](#) to bring enforcement actions against companies that engage in this activity. This means that, no matter what a company potentially says about its data practices in its privacy policy or other disclosures to consumers, it may potentially be in violation of Section 5 if it does not obtain consumer consent for the disclosure of health data to advertisers. The FTC has also actively been enforcing the Health Breach Notification Rule (HBNR) [11](#) against companies that engage in this activity (but note that the HBNR only applies to certain PHRs). For companies developing and using generative AI tools for the processing of health data, potential FTC enforcement creates a compliance risk. These entities should be particularly careful about what data they are generating from AI tools and who

that data is being disclosed to. For example, if a health app uses a chatbot to help someone self-diagnose their symptoms, any information collected about that person may constitute health information from the FTC's perspective. This may be true even if the person has not logged in to any account associated with the company because that information may still be associated with their IP address and other unique trackers (information that is still considered personal information from the FTC's perspective). If a company then shares that information to the advertisers whose trackers it uses on its website or app (without first obtaining proper consent), it may catch the FTC's eye as potentially being in violation of Section 5 of the FTC Act.

Outside of the agency's focus on health data, the FTC has also been concerned with the use and development of generative AI specifically. In recent months, the agency has issued guidance on potential deceptive market claims, ¹² consumer trust issues, ¹³ competition concerns, ¹⁴ and copyright considerations ¹⁵ related to the use and development of generative AI. The agency has also repeatedly emphasized fairness and discrimination concerns ¹⁶ associated with the use of AI. These guidance documents clearly indicate that the FTC is keeping track of market developments involving generative AI and paying attention to relevant consumer protection concerns.

State Privacy Laws

As if this regulatory landscape were not already complicated enough, there are new state privacy laws that create specific obligations both for the processing of health data and for the use of personal information for certain AI-related use cases. There are now 12 states that have "comprehensive" privacy laws (four of which are already in effect—the rest will go into effect in the coming months and years). All of these laws regulate health data in some capacity (that is not already regulated under HIPAA) as "sensitive" data and require entities to either obtain opt-in consent prior to processing such data or to provide consumers with an opt-out right. A number of these states (such as Colorado and Connecticut) provide consumers with transparency and opt-out rights in relation to the use of their personal information for automated decision-

making related to “decisions that produce legal or similarly significant effects.” This includes the use of personal information for healthcare services.

For companies using generative AI tools in the healthcare space, both the inputs and outputs associated with these tools may be subject to these compliance obligations. With regard to data inputs used to train these generative AI models, companies may need to ensure that they have consumer consent (or are providing consumers with an opt-out right) if they are using identifiable data that is considered sensitive health information under these laws. With regard to data outputs, companies may need to also provide consumers with transparency and opt-out rights if their AI tools are being used to make health-related decisions regarding consumers. For example, if a health start-up is making its new experimental product only available to consumers that meet a certain criterion and is using an AI tool to determine eligibility, that company may be subject to state privacy law obligations related to automated decision-making.

Even outside of these health- and AI-specific obligations, other elements of comprehensive privacy laws may be relevant for companies experimenting with generative AI. For example, all of these state laws require entities to provide consumers with certain data rights (such as the right to access or delete their data), which may be relevant for companies using identifiable data to develop AI tools. Like HIPAA, these laws also have a specific standard for “deidentified” data (though the standard is less stringent than what is required under HIPAA). There are also service provider contract requirements (similar to business associate agreements under HIPAA) that may be relevant if a company is using a vendor for their generative AI capabilities.

In addition to comprehensive privacy laws, there are three states (Washington, Nevada, and Connecticut) that have privacy laws specifically regulating “consumer health data.” The most notable element of these laws is their strict consent requirement for the collection and sharing of health data.

Washington’s law should be a particular concern for companies operating in this space because of its private right of action (all of the other state privacy

laws (both comprehensive and consumer health data specific) can only be enforced by state regulators).

Conclusion

The bottom line is that, for companies developing and using generative AI tools in the healthcare space, privacy compliance should be top of mind. This is true for companies that offer traditional healthcare services (such as hospitals), companies that work for healthcare providers and health insurers as business associates, start-ups experimenting with bringing cutting-edge health tools directly to consumers, and everyone in between. While regulators are learning about generative AI issues in real time, there is an existing legal framework that they can enforce against entities that are using health data in ways that they believe to be harmful to consumers. This makes proper compliance all the more necessary.

Authors



Kirk Nahra

WilmerHale, Washington, DC

Kirk J. Nahra is a partner with WilmerHale in Washington, DC, where he is co-chair of the firm's Cybersecurity and Privacy Practice and the Artificial Intelligence and Data Innovation Practice. He assists companies in a wide range of industries and of all sizes in analyzing and implementing the requirements of privacy and security laws across the country and internationally. He teaches privacy and security law issues at several law schools, including serving as an adjunct professor at the Washington College of Law at American University and at Case Western Reserve University. He received the 2021 Privacy Vanguard Award from the International Association of Privacy Professionals in recognition of his "exceptional leadership, knowledge and creativity in privacy and data protection." He also serves as a mentor to college students, law students and other young privacy and security professionals with more than a dozen organizations. He can be contacted at kirk.nahra@wilmerhale.com.

Ali Jessani

WilmerHale, Washington, DC

Ali A. Jessani is a senior associate in the Cybersecurity and Privacy Practice at WilmerHale. He counsels clients on privacy, cybersecurity, and other regulatory risks related to data protection and has particular experience advising companies on US state privacy laws, health data issues, and compliance matters related to emerging technologies, including Big Data, biometrics, and artificial intelligence. He serves on

the Publications Advisory Board for the International Association of Privacy Professionals and as an adjunct law professor at George Mason University. He can be contacted at ali.jessani@wilmerhale.com.

Endnotes



ENTITY:

HEALTH LAW SECTION

TOPIC:

HEALTH

The material in all ABA publications is copyrighted and may be reprinted by permission only. Request reprint permission [here](#).

ABA American Bar Association |

/content/aba-cms-dotorg/en/groups/health_law/publications/aba_health_esource/2023-2024/october-2023/privacy-concerns-at-the-intersection-of-generative-ai-and-healthcare