Generative Artificial Intelligence Leveraged to Deliver Healthcare - Legal Risks and Issues

With daily media reports citing to the explosion of interest in Artificial Intelligence ("AI"), AI start-ups have attracted a huge capital influx. During the last fiscal quarter of 2024 ending on December 31st, investors pumped $42 billion into AI-related start-ups.[1] Why such interest? Because in many respects, AI can be transformative and increase learning exponentially in a vastly compressed time period, with little human capital. Research and analysis of data for repetitive patterns can be dissected and summarized instantaneously.

A subset of AI, generative AI, is the focus of this article. Generative AI is defined to be a class of AI models that emulate the characteristics and structure of the AI training data to generate derived synthetic content (i.e. learned information).[2] Harnessing and analyzing huge amounts of data and information, examples of the use of generative AI in health care are as follows:

- Generative AI can collect and analyze medical scans, highlight abnormalities and recommend potential treating options for physicians.
- Generative AI can collect data from patient monitoring and develop personalized treatment plans.
- Generative AI can send reminders to patients about drug adherence, etc.
- Generative AI can assist providers in completing medical records, allowing for more efficient and speedier care.

Given this interest, and its impact on all aspects of the healthcare delivery system in the United States, it is critical to analyze the legal, regulatory and ethical issues raised by generative AI.

At first glance, the legal framework regulating generative AI is fast changing, as state and federal legislatures, as well as regulatory bodies, try to keep pace with rapid changes in AI. At the state level, to date, only California, Colorado and Utah have enacted AI statutes while pending legislation at both the federal and state level are actively being considered nationwide. Recent enacted and proposed legislation, as well as Medical Boards' guidance, highlight the significant legal issues concerning generative AI:

Clinical Decision Making

In April 2024, the Federation of State Medical Boards ("FSMB"), the supervisory oversight board of state medical boards throughout the United States, issued its guidance to State Medical Boards and their licensees (i.e. licensed physicians and providers) (the "FSMB Guidance") on the use of AI in medical practices. Specifically, the FSMB stated that:

---

[1] January 9, 2025; Crunchbase News; Eye on AI: AI Venture Funding Shattered Quarterly Records in Q4; Chris Metinko.
[2] California Civil Code Section 3110(c) (2024)

> "Physicians may consider AI as a decision-support tool that assists, **but does not replace**, clinical reasoning and discretion. Physicians should understand the AI tools they are using by being knowledgeable about their design, training data used in its development, and the outputs of the tool in order to assess reliability and identify and mitigate bias. Once a physician chooses to use AI, they accept responsibility for responding appropriately to the AI's recommendations (emphasis added)."[3]

As the FSMB Guidance emphasizes and states, it the physician's license which is at risk in the use of generative AI and, as such, the physician must confirm or alter the recommendation of AI to ensure the safety of the patient. Failure for a physician to confirm treatment recommended by AI could result in loss of the physician's license.

Transparency

Separate and apart from Clinical Decision-making, the FSMB Guidance was also clear in stating that: "Licensees should be required to maintain transparency about the use of AI in healthcare."

Such guidance was recently enacted in California (California AB 3030) which requires any "health care facility, clinic, physician's office, or office of a group practice that uses generative artificial intelligence" to ensure that: "a disclaimer that indicates that the communication was generated by generative artificial intelligence" appear on all written communications, as well as audio and visual communications" that such communication was created through generative artificial intelligence.[4]

Overriding Biases in Underlying Data

Given that Generative AI learns from the underlying data it is reviewing, any biases or prejudices in the underlying data can be amplified and reinforced in the generative AI. "Artificial Intelligence systems deployed irresponsibly have reproduced and intensified existing inequities."[5] As such, recent state regulatory and statutory actions impacting healthcare have been enacted. In New York, the NYS Department of Financial Services issued on July 11, 2024, its final Insurance Circular Letter No. 7 on the use of Artificial Intelligence Systems by, amongst other entities, third party payors in New York State, prohibiting such entities from using artificial intelligence unless "the insurer can establish through a comprehensive assessment that the underwriting and pricing guidelines are not unfairly or unlawfully discriminatory…"

---

[3] Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice. Adopted by the Federation of State Medical Board Delegates, April 2024.

[4] California Health and Safety Code Section 1339.75(a)(1). Note that if a communication that is generated by Generative AI is read and reviewed by a human licensed or certified health care provider, the disclaimer requirements do not apply. See id. At 1339.75(b).

[5] US Presidential Executive Order on the Safe, Secure and Trustworthy Development and Use of Artificial Intelligence, October 30, 2023.

Similarly, California Senate Bill 1120 was enacted in September which restricted the use of AI and other "software tools" in any UM/UR function which is discriminatory. Moreover, SB 1120 requires that physicians, and not AI, make final UR/UM decisions.[6]

Ethical Considerations for Clinicians

In concluding its guidance, the FSMB reminded its member Medical Boards of the ethical issues raised by the use of generative AI. Specifically, the FSMB admonished its members to focus on "[g]overning the use of AI through established ethical principles, including respect for patient autonomy, non-maleficence, beneficence, and justice, that have served as the foundation of professional expectations and demonstrated applicability in a variety of situations."[7] The overriding ethical concern for physicians is to ensure that it is they, and not the generative AI, which recommends treatment of their patients consistent with the physician's ethical standards, something which generative AI will not recognize.

The FDA's Role

Among its many regulatory responsibilities, the US Food and Drug Administration ("FDA") is tasked with ensuring that medical devices are safe and effective for their intended uses. For software developers new to the FDA, it is important to note that the FDA does not judge safety and effectiveness in a vacuum. Safey and effectiveness is always tied to an "intended use" as defined by federal law. It is also important to note that the FDA regulates the function, regardless of platform, of technology that meets the statutory definition of a medical device. This includes software that functions as a medical device, whether that software resides on your lap top computer, phone, watch, or in the cloud.

The Federal Food, Drug and Cosmetic Act ("FDCA") defines a medical device, in pertinent part, as:

An instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent, or other similar or related article, including a component part, or accessory which is:

* * * * *

- *intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease, in man or other animals, or*
- *intended to affect the structure or any function of the body of man or other animals, and which does not achieve its primary intended purposes through chemical action within or on the body of man or other animals and which is not dependent upon being metabolized for*

[6] California Health and Safety Code Section 1367.01.
[7] *Id* at Section VI.

the achievement of its primary intended purposes. The term "device" does not include software functions excluded pursuant to section 520(o)[8]

Section 520(o) became law in 2016 with the passage and signing of the 21st Century Cures Act. The intention of this amendment was to clarify which types of software utilized to deliver healthcare were within the FDA's medical device definition and which types were outside this definition, and therefore beyond the FDA's jurisdiction. Prior to this time, the FDA asserted jurisdiction over all health-care related software and then exercised "enforcement discretion" over certain products. This uncertainty led to this 2016 statutory change.

There are five types of health-related software identified at Section 520(o), four of which are excluded from the definition of a medical device, and the fifth (and the most important for purposes of this article) is bifurcated.

To be complete, and for context, the four health-related software areas excluded from the FDA's oversight are:

- software function intended for administrative support of a health care facility;
- software function intended for maintaining or encouraging a healthy lifestyle (commonly known as "wellness" products);
- software function intended to serve as an electronic patient record;
- software function intended for transferring, storing, converting formats, displaying data and results (commonly known as medical device data systems or MDDS)

The fifth, bifurcated category is clinical decision support software, i.e., software intended to assist a clinician in making a diagnosis or treatment plan for a particular patient. We'll get to the "bifurcated" aspect of the category in a moment.

First, it is worth discussing "intended use" – which is itself defined as the general purpose of the device or its function, which includes the indications for use. These terms of art are often the source of confusion, particularly when they are treated as interchangeable. For our purposes, it is important to understand that the "indications for use" describes, with greater precision, the disease or condition the device is intended to diagnose, treat, prevent, cure or mitigate, which includes a description of the patient population for which the device is intended.

One more source of confusion to address: clinical decision support technology, or "CDS", by definition, plays a supportive role to the clinician's decision-making. It is not required to make the decisions for the clinician to qualify as a medical device (but if it does, then it is). This is the case because the phrase "use in" prior to diagnosis, etc. in the statutory definition, provides the FDA the authority to regulate these software products when playing a supportive role to inform a clinician's medical judgment; hence the phrase clinical decision *support*. Perhaps, most importantly, "intended use" and "indication for use" is defined by the software developer as

[8] Section 201(h)(1), FDCA. If the product achieves its primary intended purpose through chemical action within or on the body, or by being metabolized, then it is a drug.

objectively measured through marketing materials and other public statements about the technology.

The FDA has put out a September 2022 guidance on Clinical Decision Support Software. Most importantly, the key distinction between CDS that is regulated as a medical device, and CDS that is not regulated as a medical device, turns on the clinician's ability to independently verify how the CDS produces its output. It comes down to transparency. If the software developer discloses the medical inputs and otherwise describes in plain language how the algorithm evaluates patient-specific information against those medical inputs, such that the clinician can (in theory) recreate the process to reach their medical conclusion or judgment, then the FDA considers the software as "non-device CDS" outside the scope of its jurisdiction.

Most frequently, however, software developers treat this information as proprietary, and therefore are not transparent enough to have their software judged to be non-device CDS. Such software, therefore, is a medical device, and will need a clearance or approval before it can be marketed or sold in the US. For software developers new to the FDA world, "clearance" is the term of art that means "approval" for moderate risk or Class II medical devices. "Approval" is the term of art used to mean just that for significant risk or Class III medical devices. Risk is defined as risk to the patient and risk to the user of the technology (when the device is hardware). In the case of CDS, the risk profile will turn on the disease or condition being evaluated.

How does generative AI fit into this framework? This is the key question. Anticipating the AI explosion, the FDA began building its Digital Health Center of Excellence in 2020. Over the past few years it has published several guidance documents on the topic to inform regulated industry and the broader public of its current thinking on these topics.[9] Most recently, the FDA published its latest draft guidance on point, highlighting its view that it intends to apply its Total Product Lifecycle approach to regulating AI-based devices.[10] It continues to evolve and will likely be the focal point of the FDA's efforts to regulate generative AI-based CDS without intending to inhibit it development or use. For now, however, the agency will need to rely upon its risk-based system to regulate the ever-changing world of AI-based medical devices.

More likely than not, therefore, we anticipate generative AI-based CDS tools will require an FDA 510(k) clearance or a PMA approval. The next issue for the FDA to grapple with will be the extent to which software developers with cleared generative AI-based CDS will need to seek a new clearance or approval. The historic framework requires 510(k) holders to seek a new clearance for an already cleared device when the holder (think manufacturer or software developer) makes a material change to the technology that can affect its safety and effectiveness. At what stage will the FDA judge that generative AI triggers the need for a new clearance? We are faced with the classic situation of law and regulation catching up to technology. We will have to wait and see.

---

[9] See FDA's website on Good Machine Learning Practice for Medical Device Development: Guiding Principles. See FDA's website on Transparency for Machine Learning - Enabled Medical Devices: Guiding Principles. See FDA's website on Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations, and the FDA's Virtual Public Workshop - Transparency of Artificial Intelligence/Machine Learning-enabled Medical Devices.

[10] See FDA's website on Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations .

Perhaps with the incoming administration's focus on making the government more efficient, there will be some fresh thinking about how to regulate these software uses to make our system nimble to keep up with technology advancements. One model that may be worthwhile to consider is the CLIA approach to evaluate the quality of lab developed tests by evaluating the quality of the labs themselves.

In the meantime, software developers focused generative AI-based CDS will not only need to satisfy the FDA's risk-based requirements, but they will also need to factor in medical practice standards, as influenced by the FSMB, around equity and eliminating bias to be eligible for use by clinicians in everyday practice.

Data Privacy Issues

Training and improving Generative AI tools generally requires large amounts of data and this need presents unique data privacy and security challenges for tools designed to be used in the health care industry. This is the case because the tools require large amounts of health data, which is considered to be one of the most sensitive types of personal data, and use of such data in this context presents heightened privacy risks for patients and consumers. Although the regulation of Generative AI technology is still in its infancy, current data privacy and security laws regulate much of the individually identifiable health information processed by these tools, with more such laws anticipated.

Each Generative AI use case involving health data, including development of such tools or implementation of such a tool within a hospital or other health care setting, must be assessed under the Administrative Simplification Provisions of the Health Insurance Portability and Accountability Act, including the Privacy Rule ("HIPAA"), as well as other applicable state laws governing health data. The following outlines the current data protection legal and regulatory scheme.

Note that HIPAA and the other privacy laws discussed below apply to identifiable health information, meaning information that directly identifies the data subject, e.g., name, address, as well as information that indirectly identifies the individual, e.g., IP address, unique numbers or other identifiers. Although de-identified information is generally not regulated by privacy laws, data that has been de-identified may not be as valuable or useful. In addition, as technology advances, the risk of re-identification increases and should be considered.

The HIPAA Privacy Rule

HIPAA regulates certain health care providers, along with health plans and healthcare clearinghouses.[11] It also regulates the vendors to such entities to the extent those vendors have access to covered identifiable health information (defined as Protected Health Information or PHI) for performance of certain types of services.[12] The HIPAA Privacy Rule contains a set of rules

[11] 45 CFR Section 160.103.
[12] Id.

governing use and disclosure of PHI.[13] As a result, both developers and users of Generative AI tools must determine underlying purpose(s) for any uses or disclosures of PHI to determine if there is an applicable exception to permit the use case. Oftentimes, there is confusion or lack of awareness regarding the applicability of HIPAA and other privacy laws when developing or using complex technologies like Generative AI. It is always important to "go back to basics" and remember that the privacy analysis under HIPAA and applicable privacy laws does not change just because Generative AI is being used. Use of such technology may make the analysis more complex as it requires a clear understanding of how the technology work and a map of all data flows. However, the basic rules still apply.

Other Privacy Laws and Regulations Governing Health Data

In addition to HIPAA, there are a number of other laws and regulations at the federal and state level intended to protect health data. 42 CFR Part 2, the federal regulations governing substance use disorder treatment records, have been recently amended to better align with HIPAA, however, these regulations contain a number of requirements that are significantly more stringent than HIPAA and must be considered in any Generative AI use case that may involve this type of information. At the state-level there are also long-standing laws governing HIV, alcohol and drug use data, genetic data and mental health data. These laws are also often more stringent than HIPAA and restrict uses and disclosures that would otherwise be permitted under HIPAA.

In addition to the above state laws that regulate specific types of health data, there has also been a strong effort at the state level to enact comprehensive privacy laws that include regulation of health data,[14] as well as state laws that specifically regulate all types of health data.[15] These laws generally apply to entities that are not governed by HIPAA and the laws were intended to fill the gaps where personal data, including health data, was unprotected by other privacy laws and regulations.

Health Data Privacy Risks and Mitigation Strategies

In addition to other risks associated with Generative AI, such as reliability issues and risk of bias discussed above, there are privacy risks associated with impermissible uses or disclosures of health information to train or improve this technology. Given the large and sensitive nature of the necessary data, there are also risks of cyberattacks and other types of data breaches. Such impermissible uses or disclosures, including data breaches could result in regulatory enforcement, imposition of fines and penalties, as well as litigation and reputational damage.

Existing data privacy and security compliance frameworks can be leveraged to help mitigate these risks, including existing privacy and data security policies, procedures and processes. Robust training on Generative AI risks and appropriate uses of this technology can also help to reduce risks associated with impermissible uses and disclosures of health data.

---

[13] See e.g., 45 CFR Section 164.502 for general rules governing uses and disclosures of PHI.
[14] See e.g., California Consumer Privacy Act, Cal. Civ. Code § 1798.100 et seq.
[15] See e.g., the Washington My Health My Data Act, RCW 19.373.

Authors:

Michael Gaba

Food and Drug Vice Chair

Email | Bio

+ 1 202.772.8496

Lisa Acevedo

Shareholder

Email | Bio

+1 312.463.6322

Paul Squire

Shareholder

Email | Bio

+1 646.289.6513