**** Pre-Publication Draft****

Does ChatGPT Dream of Electric Sheep?
Legal Implications of Artificial Intelligence in Health Care

By Daniel F. Shay, Esq.

Alice G. Gosfield and Associates, P.C.

2309 Delancey Place

Philadelphia, PA 19103

(P) 215-735-2384 (F) 215-735-4778

dshay@gosfield.com

www.gosfield.com

Accepted for publication in the Health Law Handbook, 2024 Edition. Alice G. Gosfield, Editor, © Thomson Reuters. A complete copy of the Health Law Handbook is available from Thomson Reuters by calling 1-800-328-4880 or online at www.legalsolutions.thomsonreuters.com

"I know now why you cry. But it's something I can never do."

- *Self-learning T-800 Terminator to John Connor in Terminator 2: Judgment Day*

1. Introduction

Human beings have imagined a range of different possible futures with respect to artificial intelligence ("AI"). Some are decidedly dystopian, envisioning a world in which humans are plugged into a massive computer simulation while their bodies are grown and harvested by machines for bioelectric energy (as in *The Matrix* and its sequels); or a world reduced to a blasted, irradiated wasteland patrolled by killer metal skeletons, and where the last pockets of human resistance are infiltrated by cyborgs who look like Arnold Schwarzenegger (as in the *Terminator* films). Others are more optimistic and imagine the possibility of interacting with a computer simply by speaking naturally to it, where it can respond by providing a range of information about other star systems and species, and even where holographic technology can be used to create a virtual ship's doctor who provides a full range of medical care, to say nothing of fully robotic officers powered by a positronic brain (as in the *Star Trek* franchise). In all of these scenarios, the robots, machines, cyborgs, computers, and/or holodoctors are all based on incredibly powerful AI. The reality of AI, though, is much more grounded.

At the present, AI is nowhere near as advanced as what we see in film and television, nor what we imagine it could become. We are not remotely close to having robotic doctors, for example, no doubt to the disappointment of at least a few *Star Trek* fans. Of course, this has not stopped the public from both fretting about the dangers posed, and salivating at the possible improvements to life offered, by artificial intelligence in the near future. At least some of the concerns regarding AI's functionality now and in the future, including those related to the legal complications posed by AI, are somewhat overblown and based on a lack of understanding of how AI actually functions. Others, however, are legitimate concerns based on a more solid understanding of the technology coupled with observations of the current performance of AI and speculation about how AI could perform in the future. In the midst of all of this, health care attorneys and their clients likely find themselves in a confusing new legal reality, without a real understanding of AI's capabilities or how it functions, yet still having to analyze the risks of using AI in an environment where the law is only gradually catching up to the technology.

This article will explore the impact of AI in the healthcare industry, with an eye towards physicians and physician practices. It will provide an overview of what AI actually is, and how it functions. It will examine current uses and proposed future uses of AI. Finally, it will explore the legal risks associated with its use in the physician context, and what possible problems may arise in the future, and offer practical suggestions on navigating such issues.

2. Artificial Intelligence: An Overview

At present, "AI" and "artificial intelligence" are common buzzwords. While well known to science fiction fans for many decades, "AI" burst into the zeitgeist with the public release of ChatGPT on November 30, 2022.[1] However, many do not understand what artificial intelligence actually is, nor what it is capable of doing in real terms. Moreover, there is much speculation on what artificial intelligence will be able to do and how it will interact with our lives in the future, some of it significant and some purely fanciful. To understand the potential legal issues that can arise in the use of AI, one must have a clearer understanding of certain terminology surrounding "artificial intelligence" as a concept.

[1] "Introducing ChatGPT," available at https://openai.com/blog/chatgpt, November 30, 2022.

There is much debate about what actually constitutes "artificial intelligence." The term itself is reported to have been first coined by professor *emeritus* John McCarthy of Stanford University in 1955, and was defined to mean "the science and engineering of making intelligent machines."[2] However, determining a definition for "artificial intelligence" can be as much a philosophical question as it is a technical one. For example, should we define an "AI" as one that: (1) thinks like humans, (2) acts like humans, (3) thinks rationally, or (4) acts rationally? The "Turing Test," for example, developed by computer scientist Alan Turing, is a test meant to determine whether a computer or machine can act sufficiently human to fool a human interrogator into believing it is one. If the machine can do so, then it is considered intelligent.[3] The Turing Test, therefore, defines an AI based on its ability to act like a human to other humans. If the machine can do so, then it is considered "intelligent."[4] The Turing Test, therefore, defines an AI based on its ability to act like a human to other humans. In recent years, as chatbots have been able to perform in a sufficiently deceptive manner to, at least briefly, fool interrogators (thereby allegedly "passing" the Turing Test),[5] some have called into question just how valuable the Turing Test is and pointed out the ways in which it is limited.

For example, the Turing Test is not specifically focused around different aspects of intelligence; it primarily tests linguistic abilities. Others have proposed presenting AIs with different types of tests that measure, for example, (1) pronoun disambiguation abilities (e.g., when told "The city council member refused the demonstrators a permit because they feared violence," or "…because they advocated violence," being able to correctly infer who the "they" refers to – something which many artificial intelligences find difficult); (2) giving an AI a standardized middle school exams; (3) requiring a robot or other "physically manifested AI" to assemble a structure from a pile of parts using verbal, written, and illustrated instructions (e.g., assembling IKEA furniture); or, (4) presenting an AI with a range of audio-visual content, asking it to summarize an audio file, narrate the story of a video file, etc.[6] Determining the proper test(s) for AI raises questions about the nature of intelligence itself, and may drift into the realm of the philosophical. While fascinating, these are not the focus of this article. Instead, we turn towards the more mundane and legalistic approach.

2.1 Definition of "Artificial Intelligence"

[2] Manning, Christopher, "Artificial Intelligence Definitions," Stanford University Human-Centered Artificial Intelligence, September 2020, available at https://hai.stanford.edu/sites/default/files/2020-09/AI-Definitions-HAI.pdf.

[3] See generally, Turing, A.M., "Computing Machinery and Intelligence," Mind, No. 49, pp. 433-460, available at https://redirect.cs.umbc.edu/courses/471/papers/turing.pdf. The test was also known as "The Imitation Game," giving rise to the biographical film of the same name.

[4] Id.

[5] Specifically, a chatbot named "Eugene Goostman" was able to fool around 1/3 of its interrogators into thinking it was a snarky 13-year-old Ukranian boy, whose second language was English. See, Veselov, Vladimir, "Computer AI Passes Turing Test in 'World First,'" BBC News, June 9, 2014, available at, https://www.bbc.com/news/technology-27762088. Even with respect to this instance, claims of having passed the Turing Test were disputed, noting that the test only lasted about five minutes.

[6] These tests were discussed in Marcus, Gary, "The Search for a New Test of Artificial Intelligence," Scientific American, March 1, 2017, available at, https://www.scientificamerican.com/article/the-search-for-a-new-test-of-artificial-intelligence/.

With the passage of the National Artificial Intelligence Initiative Act of 2020 (the "NAIIA")[7], the Federal government established the National Artificial Intelligence Initiative, which itself is intended to (among other things), "prepare the present and future United States workforce for the integration of artificial intelligence systems across all sectors of the economy and society."[8] The NAIIA defines "artificial intelligence" to mean:

> "*A machine-based system that can, for a given set of human-defined objectives, make predictions, recommendations or decisions influencing real or virtual environments. Artificial intelligence systems use machine and human-based inputs to – (A) perceive real and virtual environments; (B) abstract such perceptions into models through analysis in an automated manner; and (C) use model inference to formulate options for information or action.*"[9]

This provides a legal definition for the concept of "artificial intelligence" broadly, but other terms must be understood to grasp how AI functions.

2.2 "Machine Learning" & "Deep Learning"

Two commonly used terms associated with artificial intelligence are "machine learning" and "deep learning." These terms are often used interchangeably, but have distinct meanings. "Machine learning" is a subset of artificial intelligence whereby machines teach themselves without being specifically programmed.[10] Without the use of AI, for a computer to be able to do anything, it must be specifically programmed to do so. With machine learning, the computer can teach itself new functionality, usually based on the recognition of certain patterns, without being specifically programmed for the new functionality by a human.

To illustrate the difference between "traditional" software and machine learning, imagine software able to play audio files. Under the traditional model, if you wanted the software to play specific favorite songs, you would need to create a playlist, or add songs to a list of favorites. The software could then play the playlist or your tagged favorites, but only because you identified the songs it should play and (in the case of the playlist) the order in which to play them. A simple example of machine learning would be software that "learns" [11] which songs are your favorites and/or automatically creates a list of "favorite songs" based on patterns in how often you play a song and/or how often you skip it, and without requiring other input from you. This is merely a small-scale example.

---

[7] P.L. 116-283, January 1, 2021, Division E. The law was passed as part of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021.

[8] 15 USCA § 9411(a)(3).

[9] 15 USCA § 9401(3).

[10] For excellent introductory discussions of these concepts, IBM has produced a series of web articles that provide explanations of various concepts relating to AI, at https://www.ibm.com/topics/artificial-intelligence.

[11] It is worth noting how we speak about AIs in common conversation. We talk about them as if they are alive, sentient, sometimes self-aware. This tendency to anthropomorphize computer programs notwithstanding, we should remember that, at least for the present, AIs are not alive and do not truly operate independently. They are programmed, created by human hands, and at least in the case of some of the most visible AI platforms, only function based on human prompting. At present, an AI programmed to learn customer shopping preferences, even if using deep learning via neural networks, cannot also learn to play the violin or write thriller novels. Even AIs that function more independently learn only from the data that is made available to them by human hands.

Machine learning usually requires providing large amounts of data to the software to "train" it, often with human intervention to (1) select the model or algorithm that will be used to train the computer to recognize patterns or make predictions, (2) adjust the models to increase accuracy, and (3) withhold information from the data set to use to test the computer at a later time for accuracy. Different approaches to machine learning can affect how the computer performs, but the three subcategories of machine learning are supervised learning, unsupervised learning, and reinforcement learning. Supervised learning involves using labeled data to train the software, such as training a computer to recognize motorcycles by providing it with images labeled by humans as showing "motorcycles." Unsupervised learning has the program attempt to discern patterns within unlabeled data, such as being able to spot trends in customer purchases based on demographic data. Reinforcement learning involves "rewarding" the software by indicating when it has behaved appropriately, such as attempting to train a self-driving car to recognize pedestrians or ignore certain objects like a floating plastic bag.

"Deep learning" is a *subset* of machine learning that generally requires less human intervention and using automation instead, and which learns from extremely large data sets. Deep learning tends to rely on unsupervised learning, and is provided with unstructured data. By way of example, if one took a set of photographs of cars, and wanted to train an AI to group them into categories of "sedan," "SUV," and "truck," a deep learning algorithm would learn which features are important to distinguishing between the types of vehicles. For example, the overall height of the vehicle, whether it has a covered cargo bay that is part of the body of the vehicle, etc. In other machine learning – supervised learning, usually – these features and their relative importance would be manually determined by humans. In deep learning, the computer itself would distinguish these features, adjusting over time to increase its accuracy. Deep learning uses "neural networks," another buzzy phrase associated with discussions of AI.

2.3 Neural Networks

"Neural networks" are a form of machine (deep) learning that attempts to mimic the capabilities of the human brain, albeit in a much less sophisticated manner, from which they get their "neural" moniker. Neural networks are made up of "node layers," which consist of an input layer (where one enters data), one or more hidden layers (which analyze data), and an output layer (which outputs results), with each layer made up of individual nodes. The input-layer nodes each feed information they receive to the topmost hidden-layer nodes, which in turn feed information to the next highest hidden-layer nodes, and so on, until the final hidden-layer nodes feed information to the output layer nodes. Most neural networks feed information in only one direction, but some can be programmed to feed information in the opposite direction (a process known as "back propagation").[12] The hidden-layer nodes also do not always pass on information; information is passed on only if it meets certain criteria programmed into the node.

The "deep" in "deep learning" means that the program is using a neural network with more than three layers (including the input and output layers). Otherwise, the program is a basic neural network. Put another way, all deep learning models use neural networks; not all neural networks operate as deep learning models.

2.4 Generative AI

[12] For a visual representation of what is an admittedly confusing concept, see the graphic on "AI vs. Machine Learning vs. Deep Learning vs. Neural Networks: What's the difference?", available at https://www.ibm.com/blog/ai-vs-machine-learning-vs-deep-learning-vs-neural-networks/.

One way in which deep learning AI is influencing popular culture is through "generative" AIs such as ChatGPT, Google's Bard, or image generation sites like DALL-E or Midjourney. It is these AI models that have prompted the most attention from the public, and which have jumpstarted the public's collective imagination about the potential for AI technology and its uses. "Generative" AIs are, as implied by the name, AIs that generate outputs based on input by the user. By entering a simple prompt, the software can produce outputs, which often may appear creative. "Generative AI refers to deep-learning models that can take raw data — say, all of Wikipedia or the collected works of Rembrandt — and 'learn' to generate statistically probable outputs when prompted. At a high level, generative models encode a simplified representation of their training data and draw from it to create a new work that's similar, but not identical, to the original data."[13]

Generative AIs are trained on enormous amounts of data to generate outputs derived from that data, as instructed by prompting. As described by OpenAI, creators of ChatGPT, "To train a generative model, we first collect a large amount of data in some domain (e.g., think millions of images, sentences, or sounds, etc.) and then train a model to generate data like it."[14] The results created by generative AIs are meant to be similar to (indeed, derived from) the dataset upon which they are trained, but different and not simple reproductions of the data. It is debatable how "new" the output of generative AIs actually is.[15]

Some generative AIs are what are referred to as "large language models" or "LLMs." These are a subset of generative AIs, and are meant to understand and generate text the way humans do, again, based on the data upon which the LLM was trained. While this might sound similar to the function of a search engine, the key difference is the generative output. With a search engine, the user enters a prompt, and the engine attempts to reproduce the information requested. But the search engine is literally only reproducing data to which it already has access; it does not generate new content based on that data nor rearrange the data and present it differently. With generative AIs, like ChatGPT or Google Bard, the AI is capable of outputting something based on, but which did not previously exist within, the dataset upon which the AI was trained. This is accomplished through the use of deep learning models, and the AI's pattern recognition capabilities, coupled with immense amounts of unstructured data.[16]

---

[13] Martineau, Kim, "What is generative AI?", April 20, 2023, available at https://research.ibm.com/blog/what-is-generative-AI?utm_content=SRCWW&p1=Search&p4=437000780779089 73&p5=e&gclid=CjwKCAiA75itBhA6EiwAkho9ewESWU8r8EPGJLHPhX4iLCKc2EC4y4MgdEhrRzBYGjUX5rOLq-Xe7RoCJuYQAvD_BwE&gclsrc=aw.ds.

[14] "Generative Models," OpenAI blog, June 16, 2016, located at https://openai.com/research/generative-models.

[15] Debates surrounding the legalities surrounding the originality of generative AIs is perhaps best left to copyright attorneys.

[16] In fact, the vast amounts of information being processed are also analyzed through the use of so-called "transformer" models of machine learning, which allow the processing of information in parallel, rather than in sequence, thereby allowing the machine to process blocks of information much faster. The term "transformer" is one also seen in discussions of AI capabilities, albeit less often in popular discussions than other terms addressed in the body of this article. That said, most well-known chatbots and image or video generation sites use transformer technology. For a good starting point in the discussion of transformers, see Towes, Rob, "Transformers Revolutionized AI. What Will Replace Them?", Forbes, Sept. 3, 2023, located at https://www.forbes.com/sites/robtoews/2023/09/03/transformers-revolutionized-ai-what-will-replace-them/?sh=5aa0d9919c1f.

The end result is an AI that can understand and create text the way a human would, or at least to provide as close an approximation to human behavior as is possible within the limits of the AI.

For example, asking Chat GPT-3.5[17] to "describe the experience of eating a cheeseburger" produced an output within seconds that began with the following statement[18]:

> *Eating a cheeseburger can be a delightful experience. The moment you pick up the burger, you're greeted by the aroma of grilled beef, melted cheese, and various toppings. The first bite is a symphony of flavors and textures – the juicy beef patty, the creamy melted cheese, the crisp freshness of lettuce and tomatoes, and the tangy kick of ketchup or other condiments.*

Entering the identical prompt a second time produced the following opening paragraph:

> *Eating a cheeseburger can be a delightful experience for many people. Imagine biting into a perfectly grilled beef patty, topped with melted cheese, nestled between two soft and slightly toasted burger buns. The first sensation is often the savory and juicy flavor of the beef, complemented by the creamy richness of the melted cheese.*

A third identical prompt produced the following as the opening paragraph:

> *Eating a cheeseburger can be a delightful experience. Imagine biting into a juicy and flavorful patty, cooked to perfection, nestled between soft and slightly toasted burger buns. The combination of the savory meat, melted cheese, crisp lettuce, ripe tomatoes, and tangy condiments creates a harmonious symphony of tastes and textures.*

Each output had similar words and phrasing, but presented somewhat differently each time, including focusing on different word combinations or adding descriptive text associated with different elements. Still, there are repeated phrasings or couplings of words and concepts, such as the "symphony of tastes/flavors and textures," or the "tangy" condiments/ketchup, or the statement "imagine biting into a…patty…" To some extent, results like this may call into question how "new" the content is that is created by the AI. Put another way, is the content truly new, or is it simply a reshuffling of concepts, words, or phrases that were designated as important in the training of the AI? Is the AI actually *creating,* or is it merely regurgitating in a highly sophisticated fashion? At least in the case above, the AI's responses are easily distinguished from human responses due to the AI's tendency to repeat phrases, concepts, or word associations in a way that seems strange or artificial.

For most readers, AI is already a factor in their own lives, possibly in ways they do not even realize, usually through some form of machine learning. For example, many streaming entertainment platforms now include systems whereby a user can watch a film or TV show, or listen to a song, indicate whether they like that film, show, or song, and the streaming service will customize recommendations for the user based on that feedback, coupled with data gathered from other similar users and the use of various data indicators within the content. Thus, a platform like Pandora might be able to distinguish between genres (likely thanks to some degree of supervised learning), and common aspects between songs and/or artists, and thereby continue playing songs that the user is more likely to want to hear thanks to the use of machine learning algorithms. Similar algorithms determine what content is

[17] While OpenAI does offer GPT-4, use of the GPT-4 model required a paid subscription at the time of this article's writing. Thus, the author has only used the freely available GPT-3.5 model.

[18] The full output was multiple paragraphs regarding eating a cheeseburger.

presented to users on social media sites, based on a range of information, likely including whatever demographic information the user shares, and especially based on the user's activity and engagement where the more you interact with certain content, the more the algorithm will feed you that same kind of content and make recommendations of new content in the same vein.

3. Current Uses of Artificial Intelligence in Health Care

The deployment of artificial intelligence within the health care industry is still in its infancy. Nevertheless, there are some current users, and more are certain to follow. Artificial intelligence is currently being used (or at least offered by vendors) in a variety of functions, including clinical notetaking, image analysis and diagnostics, interactive chatbots, and medical coding functions.

For example, as of this writing, Amazon Web Services is advertising its own generative AI capability and how it could be used by health care providers, promoting a wide range of potential uses including: digital scribing/transcription, medical image interpretation, automated medical coding, an "intelligent health assistant" (essentially an LLM-based chatbot), and health care document summarization (e.g., for summarizing research, performing document standardization, etc.).[19] Similarly, Nuance Communications is offering Dragon Ambient eXperience, described as a "clinical intelligence solution" that is meant to document patient encounters at the point of care.[20]

3.1 AI, Image Analysis, and Other Diagnostics

Perhaps unsurprisingly, the most active integration of artificial intelligence in health care is in radiology through AI-assisted image analysis. As of March, 2023, the FDA had approved roughly 500 different medical AI tools, with nearly 400 of them applying to radiology.[21] By January, 2024, the total number of FDA-approved AI/machine learning enabled devices has increased to nearly 700.[22] Driving this wide adoption of AI tools are, arguably, two factors: (1) that artificial intelligence pairs much more naturally with a field oriented around the analysis of discrete images from which patterns may be discerned, and (2) "a shortage of radiologists – with a demand for imaging services growing and the complexity of imaging services increasing…Something has got to give."[23] Radiologists are already employing the use of AI-assisted readouts of mammograms (for which patients must currently pay out-of-pocket – sometimes as much as $100 per screening, due to payors not yet covering the AI-assisted service).[24] Recent studies have demonstrated that artificial intelligence significantly improved cancer

---

[19] See, https://aws.amazon.com/health/gen-ai/.

[20] See, https://www.nuance.com/healthcare/ambient-clinical-intelligence.html.

[21] Hudnall, Chad E., "Choosing AI," ACR Bulletin, March 23, 2023, available at https://www.acr.org/Practice-Management-Quality-Informatics/ACR-Bulletin/Articles/April-2023/Choosing-AI.

[22] Johnson, Mark, "How Doctors are Using AI to Diagnose a Hidden Heart Condition in Kids," Washington Post, January 16, 2024, available at https://www.washingtonpost.com/science/2024/01/16/ai-diagnose-rheumatic-heart-disease/.

[23] Gloria L. Hwang, M.D., associate chair for clinical performance improvement in the radiology department of Stanford University, as quoted in "Choosing AI," ACR Bulletin, March 23, 2023, available at https://www.acr.org/Practice-Management-Quality-Informatics/ACR-Bulletin/Articles/April-2023/Choosing-AI.

[24] DePeau-Wilson, Michael, "Is Adding an AI Readout to Mammography Worth It?", MedPage Today, January 24, 2024, available at https://www.medpagetoday.com/special-reports/features/108405?xid=nl_mpt_DHE_2024-01-

detection rates. In interim results from a randomized, controlled trial of 80,000 women in Sweden, cancer detection rates were 20% higher when an AI performed an additional read of a mammography.[25]

Artificial intelligence has similarly proven useful in diagnosing children with rheumatic heart disease.[26] Specifically, the Uganda Heart Institute, working with Children's National Hospital in Washington, designed an AI-powered tool to diagnose the disease well before surgical intervention would be needed. The disease, caused by a bacterial infection, can be detected by an echocardiogram, but many impoverished countries lack the number of cardiologists necessary to interpret the images. The new AI would allow trained nurses to screen children early in the disease's progression, when penicillin could be prescribed as treatment for less than $1 per year. The device is made up of a probe about the size of a computer mouse, which connects to a tablet using the AI component. The AI itself is capable of interpreting images rapidly, and classifying them either as normal or instructing the clinician to review and consider for the presence of rheumatic heart disease. If the AI detects the possibility of the disease, the test would be followed up with a cardiology examination which could be conducted in person or over telemedicine. In interpreting echocardiograms from over 500 patients, the tool was able to reach almost 90 percent accuracy.

Computer-aided imaging software has existed in one form or another for at least 30 years. But whereas traditional computer-aided detection and computer-aided diagnosis (referred to as CADe and CADx, respectively) used older machine-learning techniques that required human intervention and manual pre-definition of lesions to learn to detect abnormalities, new deep-learning AIs are capable of learning to identify patterns rapidly automatically without requiring human pre-labeling and programming.[27] Thus, AI, especially when empowered by deep-learning models, is especially well suited to the field of radiology and image analysis where the AI's development can be oriented around determining patterns in both "normal" and "abnormal" images. This capacity can then be leveraged into recognizing the presence of specific diseases by differentiating between features of abnormal images. In a sense, the data are somewhat circumscribed. To the extent that the AI is distinguishing simply between "normal" and "abnormal" images, or even if the AI is recognizing patterns that distinguish between common presentations of specific diseases, the AI may be able to make that determination based solely on the image itself.

The pattern-recognition capabilities of artificial intelligence have also been put to use in detecting potentially faulty research. In January, 2024, the Dana-Farber Cancer Institute in Boston,

---

24&eun=g2230312d0r&utm_source=Sailthru&utm_medium=email&utm_campaign=Daily%20Headlines%20Evening%202024-01-24&utm_term=NL_Daily_DHE_dual-gmail-definition.

[25] DePeau-Wilson, Michael, "Is Adding an AI Readout to Mammography Worth It?", MedPage Today, January 24, 2024, available at https://www.medpagetoday.com/special-reports/features/108405?xid=nl_mpt_DHE_2024-01-24&eun=g2230312d0r&utm_source=Sailthru&utm_medium=email&utm_campaign=Daily%20Headlines%20Evening%202024-01-24&utm_term=NL_Daily_DHE_dual-gmail-definition.

[26] Johnson, Mark, "How Doctors are Using AI to Diagnose a Hidden Heart Condition in Kids," Washington Post, January 16, 2024, available at https://www.washingtonpost.com/science/2024/01/16/ai-diagnose-rheumatic-heart-disease/.

[27] Brady, Adrian P., et al., "Developing, Purchasing, Implementing and Monitoring AI Tools in Radiology: Practical Considerations. A Multi-Society Statement from the ACR, CAR, ESR, RANZCR & RSNA," Journal of the American College of Radiology, January 22, 2024, p.2, available at https://www.jacr.org/article/S1546-1440(23)01020-7/fulltext.

Massachusetts moved to retract six research papers, and correct multiple others, following allegations of faulty data in 58 studies.[28] Specifically, the individual who discovered the faulty data, Dr. Sholto David, a British molecular biologist, alleged irregularities with western blot testing results that appeared to have been copied and pasted across different experiments. In other instances, Dr. David noted images that appeared copied, rotated, or stretched to suggest intentional manipulation. In one instance, he found an image of mice from the first day of an experiment that apparently reappeared in the results on day 16, in a separate portion of the experiment. Dr. David, operating by himself, used AI software to identify some of the problems before personally checking the images.

3.2 Health System Deployment of AI

The capabilities of artificial intelligence are being put to broader use than image analysis alone, however. Phoenix Children's Hospital uses multiple AIs developed internally with local data, to perform tasks such as detecting indicators of malnutrition in patients based on medical records. The hospital reports that between 60 and 80 percent of patients identified by the AI are proven to have malnutrition after follow-up with a nutritionist. The hospital estimates that roughly seven patients per week are identified as being malnourished who might otherwise have gone undiagnosed. Other AI-powered technology has been used by the hospital to monitor vital signs and detect circumstances where the patient's condition worsens rapidly and then alert the care team of the patient's condition. Since first being deployed in 2021, the technology identified more than 100 children who needed immediate transfer to the ICU.[29]

Other health systems have used AI to perform a variety of different tasks. CHI Health, based in Omaha, Nebraska, has made use of population health management tools to scan patient medical records and find patients who are due for a lung or colon cancer screening, and then provide the patient's doctor with that information for the doctor to determine whether it is appropriate for the patient to have the screening service. Eskenazi Health system has developed an AI tool to find patients likely to develop Alzheimer's disease. The tool reviews a range of information about patients (e.g., memory issues, vascular conditions, other factors that could indicate dementia) and then predicts which patients are likely to develop Alzheimer's disease based on the responses. Since July, 2022, the tool has examined over 5,400 primary care patients who are 65 or older and have no pre-existing cognitive impairment; the AI is reported to have up to 80 percent accuracy in identifying patients who will develop dementia within between one and three years.[30]

3.3 AI Deployment in Physician Practices

Health systems and larger institutions, however, are not the only entities deploying AI within the health care setting. The American Academy of Family Physicians ("AAFP") has developed its own AI assistant tool designed to assist in clinical review of patient charts. In a report published in January, 2024, the AAFP noted that family physicians spend more than 13 percent of their time (over one and a

[28] Mueller, Benjamin, "Dana-Farber Cancer Institute Seeks to Retract or Correct Dozens of Flawed Studies," New York Times, January 24, 2024, A13.

[29] Deveraux, Mari, "How Health Systems are Using AI to Augment Patient Safety," Modern Healthcare, January, 2024, pp. 28-29, 31.

[30] Id.

half hours) per clinic day conducting chart reviews.[31] The role of the AI assistant, developed in partnership with Navina, is to review a patient's entire chart and available records, including lab results, diagnostics, referrals, scanned documents, etc., and provide the physician with a problem-oriented summary of the patient's chart. The AAFP reports that the assistant can produce a report within minutes, which will identify missing diagnoses, suspected conditions, and gaps in care, and can also provide recommendations to the physician. The study comprised approximately 683,200 patient visits.[32] The assistant was tested for a period of roughly 6 months, across ten different physician practices ranging in size from under 50 providers to more than 300 that were either already participating in value-based care programs, or transitioning into them, and surveyed individual physicians to ask about their results.[33] Roughly 30 percent of the 58 total respondents reported a reduction in chart review burden; 45 percent reported an increase in feeling prepared for the visit; 23 percent reported a decreased sense of burnout; and almost 22% reported an improvement in satisfaction. On average, approximately 9 minutes were saved per visit, and preparation time was reduced by about 38 percent.[34] These may seem like disappointing numbers, but 82 percent of respondents indicated that they would recommend an AI assistant, with over half indicating that they strongly agree with such a recommendation. Moreover, 98% of respondents thought that the AI assistant enabled easier chart review and more accurate capture of diagnoses.[35] Ultimately, the AAFP determined that the study indicated that an AI assistant could significantly save physicians time, ease burden, and reduce burnout. Interestingly, the report indicated that the AI assistant was used by respondents in 85 percent of patient visits, and that the assistant's clinical recommendations were followed 84 percent of the time.[36]

4. Future Uses of Artificial Intelligence in Health Care

---

[31] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, p. 2, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html. The relevant report is the Navina Phase 2 Lab Report.

[32] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, p. 1, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html.

[33] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html. A chart indicating the precise breakdown for practices can be found on p. 3 of the report.

[34] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, pp. 4-5, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html.

[35] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, p. 6, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html.

[36] Waldren, Steven E., M.D., M.S. and Edmund Billings, M.D., "Artificial Intelligence Assistant for Clinical Review and Value-Based Care," American Academy of Family Physicians, January, 2024, p. 7, available at https://www.aafp.org/family-physician/practice-and-career/managing-your-practice/health-it/innovation-lab.html.

While adoption and implementation of AI in health care is still somewhat slow, one can easily envision a future where AI is far more integrated into health care providers' day to day life. The Brookings Institute has noted that generative AI, for example, "Performs optimally in environments characterized by high repetition and low risk. This effectiveness stems from the technology's reliance on historical data to identify patterns and make predictions, under the premise that future conditions will mirror those of the past."[37] At the very least, we can therefore expect to see artificial intelligence deployed in ways that will automate these repetitive tasks, likely in a manner that makes their performance more efficient both in terms of time and cost. However, as AI develops and improves, the risks posed by its use may be mitigated through enhancements in the technology, leading to its use in performing or assisting in clinical functions.

4.1 "Front Desk" and Triage Uses of AI

One of the easiest ways to imagine generative AI performing repetitive tasks is in the performance of certain "front desk" and administrative duties. For example, one need not strain to imagine chatbots being used for patient intake. Automated kiosks are already used for patient check-ins at various laboratory specimen collection sites. Expanding that functionality to use natural language processing interaction via a chatbot is a natural progression of such technology. This capability could also be broadened to perform triage functions at various levels. Consider how an automated, AI-driven call center might function, where a computerized (but still human-sounding) voice could take a patient's complaint and forward the relevant details to the on-call physician. Some of the technological backbone for such a system already exists. AI-enabled voice assistants are already prevalent on phones and other devices, such as Apple's Siri or Amazon's Alexa. As the technology develops and can be trained on larger amounts of data (including health care data), the risks posed by using it with patients calling in to a practice – either during the day or during "on call" hours – may be reduced (but likely never eliminated altogether). In an interview with the Journal of the American Medical Association, Gary Marcus, Ph.D. posits a scenario where an app on one's phone, powered by AI, could be used to determine how urgent a case is. He states, "The most useful thing in the short term will be dermatology apps for people who can't easily get to a dermatologist, and you really need to triage who should come in."[38]

In fact, at larger institutions, AI use in triage is already happening. Johns Hopkins University developed and implemented a nurse-assistance triage AI in 2022 with positive results.[39] The AI in question, deployed in the hospital's emergency department, is used by triage nurses to rapidly evaluate patients and do so in a more standardized manner. When a patient presents, the nurse asks the patient about their condition and takes vital signs. This information is combined with the patient's medical history, and then run thru the AI algorithm to determine the patient's risk of certain acute outcomes, after which the software recommends a triage level of care and provides an explanation for the decision. The nurse has the final say on the patient's triage level, however.

---

[37] Yaraghi, "Generative AI in Health Care: Opportunities, Challenges, and Policy," Brookings Institute, January 8, 2024, available at https://www.brookings.edu/articles/generative-ai-in-health-care-opportunities-challenges-and-policy/.

[38] Desai, Angel N., M.D., MPH, "Artificial Intelligence: Promise, Pitfalls, and Perspective," JAMA, June 30, 2020, available at, https://jamanetwork.com/journals/jama/fullarticle/2766942.

[39] "Digital Health Startup That Assists Emergency Department Decision Making Acquired," Johns Hopkins Technology Ventures, October 11, 2022, available at https://ventures.jhu.edu/news/stocastic-beckman-coulter-acquisition-digital-health/.

To some extent, computers (albeit not necessarily powered by AI) already perform some automated functions relating to scheduling and intake. For example, many patients already receive automated messages from their doctors' patient portals in preparation for office visits. These notices often instruct a patient to complete various forms to save time at check-in, and to take certain aspects of a patient's history, all with an eye towards improving throughput in the office. Artificial intelligence may likewise be developed to enhance these processes, as well as to communicate with physicians.

For example, patient portals may develop interactive chatbot-style tools for patients to request medication updates, referrals, or to present test results to patients in a more easily understood manner. Likewise, electronic health records software may develop to provide more effective patient monitoring, such as informing a physician when a patient has not requested a refill of a particular prescription within a given timeframe, or to notify the physician for things like glucose monitor results that are above a given range or abnormal for that specific patient. The Brookings institute proposes a potential use where an AI may cross-reference patients' medication list and current health complaints, and then "verify whether patients are adhering to their prescribed regimens or have discontinued any conflicting medications in light of new prescriptions. This process aids in assembling a more comprehensive medical history for the patient, which can then be used by the physicians to provide better care."[40]

4.2 Billing, Administrative, and Practice Management Uses of AI

Billing and coding provide another use case for AI in health care, where AI could be used to standardize the preparation and submission of claims. This, however, would likely need to be coupled with standardized chart generation, which could also be performed by AI. Especially in the case of a deep learning model, the software could be trained to interpret the shorthand used in notetaking by the physicians. The patient's record could be further enhanced by providing transcriptions of interactions between patient and physician, which would likely require the use of an LLM-powered AI. If the two functions were similarly connected by AI, the software could, for example, alert the physician to an element in the patient transcript that the physician had failed to reference in their own post-visit notes. This could then be used to produce more standardized notes throughout a single institution or practice, which would aid in the preparation and submission of claims, and further assist in matters such as claims appeals or responding to payor audits. Relatedly, because of its ability to recognize patterns in large volumes of data, AI may also prove useful in compliance efforts, such as conducting internal self-audits of billing patterns, as well as HIPAA compliance audits such as recognizing data access and/or transfer patterns necessary to spot hacks or hacking attempts.

A range of administrative tasks necessary for the function of physician offices can likely be automated and operated by AI-driven software. For example, managing prior authorizations is one of the major time-sinks in a primary care physician's day,[41] and AI may eventually be able to perform this service instead of requiring physician involvement. Especially for primary care physicians,[42] developers

---

[40] Yaraghi, "Generative AI in Health Care: Opportunities, Challenges, and Policy," Brookings Institute, January 8, 2024, available at https://www.brookings.edu/articles/generative-ai-in-health-care-opportunities-challenges-and-policy/.

[41] A fact recently lamented by James DomDera, M.D., the Medical Editor of FPM in a "From the Editor" piece titled "Can AI Solve Our Bigger Problem?", FPM, November/December 2023, p. 3.

[42] The author wishes to thank LeeAnn Tanaka, D.O. (the wife of the author, in the interests of full disclosure) for her invaluable contributions and insights in discussing the types of issues she faces in her work as a primary care physician. Dr. Tanaka provided several suggestions and observations on ways that AI might alleviate her daily clinical and non-clinical burdens in the primary care setting.

may create an AI to assist in minimizing "peer to peer" interactions where primary care physicians are required to call an insurance company medical director to inquire as to why a given order was denied by the insurance company. The potential use for AI in this scenario would be to monitor the specific coverage and documentation requirements necessary to order a given drug, device, or service for a patient and inform the physician before submitting the order as to whether the requirements had been met. This would also avoid checkbox or click-thru "fatigue" where physicians, irritated by a barrage of pop-up windows and boxes that must be ticked or clicked to proceed, simply speed thru a process; in this case, the AI would be reviewing the patient's record itself, within seconds, compare that to the insurance company's requirements, and then inform the physician if anything necessary was missing.

Separate from triage and call-in center functionality, AI may develop to assist with a range of scheduling issues both for patients and for health care providers themselves. For example, AI could be used to monitor patterns in patients who tend to "no-show" to better ensure that physicians are not over-booked, while still being able to treat walk-ins and maintain productivity. This, in turn, could help reduce physician burnout. On the practitioner side, AI could monitor scheduling and productivity patterns to shift coverage by medical assistants and nursing staff to ensure adequate coverage, speedy rooming of patients, and efficient workflow, while also ensuring that such staff are not themselves overscheduled.

4.4 Population Health Uses of AI

While it is already being used in some aspects of diagnostics, AI will inevitably be developed to become more accurate and versatile in its ability to, if not outright diagnose patients the way a physician would, at least serve as an extremely useful tool for physicians to spot potential issues.

Population health offers still another area in which AI will likely develop over time, both on a large and small scale. Through the use of deep learning models, AI may be deployed to detect patterns in health care data and spot potential issues as they are developing and potentially before they develop. This may require the AI to learn on large amounts of public health data and specific patient records, such as those contained in state-wide health information exchanges. For example, effective use of such AI might detect increases in conditions associated with lead poisoning, which could potentially be used to spot tainted drinking water within communities. Initially these tools will be most useful at much larger scales, such as by state Departments of Health, or within larger health systems. But as AI develops greater precision and improves its ability to recognize patterns within health care data at the smaller scale, these tools may likewise prove useful within smaller physician practices.

5. The Legal Response to AI

Artificial intelligence is almost certain to be an increasingly prevalent part of human existence in both the near- and long-term future. When the internet first burst into the public consciousness in the early 1990s, it was not merely because the idea was exciting, but rather that it was poised to become an omnipresent aspect of our lives. Some thirty years later, it remains such, and unsurprisingly this has led to the passage of new laws and regulations, as well as adapting other laws and regulations, to govern use of the internet. The same is currently needed with respect to AI, and will be needed even more in the future. While still a novel field, both the federal government and state governments have taken steps to regulate the use and development of AI, although it is still very much early days with regard to regulating AI.

5.1 Federal Action

At the federal level, as of this writing, there is relatively little by way of law governing artificial intelligence. Although there is no comprehensive law governing AI, some comprehensive legislation has been proposed and continues to be worked on.[43] Other proposed legislation has been targeted towards more specific issues, dealing with them in a discrete manner rather than attempting to create an overarching law to govern all aspects of AI. For example, Sen. John Thune (R-SD) proposed the Artificial Intelligence Research, Innovation and Accountability Act of 2023, which would require NIST to develop standards to authenticate online content as human- or AI-generated, among other things.[44] Another example is the No Robot Bosses Act, proposed by Sen. Robert Casey, Jr. (D-PA), which would prohibit some uses of automated decision systems by employers.[45]

What legislation has been passed in recent years seems more focused around establishing task forces and agencies in the Executive Branch. For example, as referenced above at Section 2.1, in 2021, Congress passed the National Artificial Intelligence Initiative Act of 2020 ("NAIIA").[46] In addition to providing the general definition of "artificial intelligence" previously discussed, the NAIIA instructs the Director of the Office of Science and Technology Policy to establish and appoint a director of an office to be known as the National Artificial Intelligence Initiative Office. This office is tasked with providing technical and administrative support to certain committees; serving as point of contact for Federal agencies, industry, academia, etc.; engaging in public outreach to stakeholders, and promoting access to technologies, innovations, best practices, and expertise from the Initiative's activities across the federal government.[47] The Act also requires the creation of an Interagency Committee[48] and Advisory Committee,[49] the latter of which will be made up of advisors from various perspectives such as academic institutions, companies across multiple sectors, nonprofit and civil society organizations, etc., and which will have the duty to advise the President on matters relating to the Initiative itself.

---

[43] For example, Senate Leader Chuck Schumer's proposed legislation, known as the Security, Accountability, Foundations Explainability (SAFE) law. See, "Majority Leader Schumer Delivers Remarks to Launch SAFE Innovation Framework For Artificial Intelligence At CSIS," June 21, 2023, available at https://www.democrats.senate.gov/news/press-releases/majority-leader-schumer-delivers-remarks-to-launch-safe-innovation-framework-for-artificial-intelligence-at-csis.

[44] S. 3312 of the 118th Congress 1st Session, available at, https://www.congress.gov/bill/118th-congress/senate-bill/3312/text.

[45] S. 2419 of the 118th Congress 1st Session, available at, https://www.congress.gov/bill/118th-congress/senate-bill/2419/text. For a list of AI legislation introduced in the 118th Congress, see the Brennan Center's Artificial Intelligence Legislation Tracker, available at https://www.brennancenter.org/our-work/research-reports/artificial-intelligence-legislation-tracker.

[46] P.L. 116-283, January 1, 2021, Division E. The law was passed as part of the William M. (Mac) Thornberry National Defense Authorization Act for Fiscal Year 2021.

[47] 15 USCA § 9412.

[48] 15 USCA § 9413.

[49] 15 USCA § 9414.

The Initiative itself was created for specific purposes outlined in the statute, namely (1) to ensure the United States' continued leadership in AI research and development; (2) to lead the world in the development and use of trustworthy AI systems in both public and private sectors; (3) to prepare the U.S. workforce, both in the present and the future, to integrate with AI systems across all sectors of the economy and society; and, (4) to coordinate ongoing AI research, development, and demonstration activities in the various federal civilian agencies, as well as the Department of Defense, and the intelligence community.[50] It does not establish guidelines for the use or development of AI itself, nor does it specifically reference the Department of Health and Human Services, nor provide any guidance on AI in health care, although the National Institute of Standards and Technology for the Future ("NIST") has had legislation passed requiring it to support the development of AI and data science, and carry out the activities of the NAIIA.[51] Still, it represents an initial step with respect to developing the tools to potentially regulate AI.

Within the Executive Branch, the Biden Administration has taken two recent, major steps in addressing the issue of AI in a broad-spectrum manner. First, in October, 2022, the Biden Administration, through its Office of Science and Technology Policy, published a Blueprint for an AI Bill of Rights.[52] The Blueprint is actually a 73-page white paper, oriented around a short 2 and 1/3 page statement of principles. The bulk of the document is made up of definitions of terms, and how these statements of principle would be applied on a practical level. As an initial matter, the document itself is not legally binding in any sense, but rather provides a window into the Biden Administration's views with respect to artificial intelligence and what constitutes responsible and appropriate development and use of such technology, with an eye towards ensuring the safety and privacy of Americans. For example, the Blueprint generally states that individuals "should be protected from unsafe or ineffective systems," meaning that AI systems should be designed with the safety of both end-users and the general public in mind. It further states that individuals should be protected from algorithmic discrimination, and that AI systems should be designed to protect against such discrimination, exhorting developers and those who deploy such systems to "take proactive and continuous measures to protect individuals and communities from algorithmic discrimination and to use and design systems in an equitable way."[53]

Unsurprisingly, data privacy also appears in the Blueprint, and it states that individuals should only have their data used with permission, and that protection from violations of privacy should be based on "design choices that ensure such protections are included by default, including ensuring that data collection conforms to reasonable expectations and that only data strictly necessary for the specific context is collected."[54] The Blueprint also states that individuals should know when an automated

---

[50] 15 USCA § 9411.

[51] 42 USCA § 18937.

[52] The full title of the document is "Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People." The full text is available at https://www.whitehouse.gov/wp-content/uploads/2022/10/Blueprint-for-an-AI-Bill-of-Rights.pdf. An overview of this document can be found at https://www.whitehouse.gov/ostp/ai-bill-of-rights/.

[53] "Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People," October, 2022, p. 5.

[54] "Blueprint for an AI Bill of Rights: Making Automated Systems Work for the American People," October, 2022, p. 6.

system is being used and understand how and why it impacts them. Finally, it states that individuals should be able to opt out when possible, and have access to a human who can address the problems the individual encounters. Again, these are not requirements, or even actual protections, but rather a kind of "mission statement" for how AI should be used and designed with a particular eye towards the experience of individuals.

A year later, in October 22, 2023, the Biden Administration published a sprawling Executive Order titled "Executive Order on the Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence."[55] The Executive Order is a government-wide effort, directing more than 50 federal entities to engage in over 100 specific actions to implement guidance across eight overarching policy areas: (1) safety and security, (2) innovation and competition, (3) worker support, (4) consideration of AI bias and civil rights, (5) consumer protection, (6) privacy, (7) federal use of AI, and (8) international leadership.[56]

For example, the Executive Order instructs the Secretary of Health and Human Services to "identify and, as appropriate and consistent with applicable law [and other portions of the Executive Order], prioritize grantmaker and other awards, as well as undertake related efforts, to support responsible AI development and use…"[57] Specifically, the Secretary is instructed to collaborate with private sector actors through the use of HHS programs that may support advancing AI-enabled tools to help develop personalized immune-response profiles for patients, and to accelerate grants awarded through certain National Institutes of Health programs. Interestingly, another element of the HHS instructions include providing certain awards to initiatives looking to improve health care data quality for the purpose of development of AI tools for use in clinical care, population health, public health, and other similar research.

The Executive Order further instructs HHS to publish a plan to address the use of AI by states and localities in distributing public benefits and services that are funded by HHS. As examples, the Executive Order cites the promotion of notice to recipients about the presence of such systems, processes to appeal denials to human reviewers, and regular evaluation to detect unjust denials.[58] The Secretary of HHS is also instructed to, in consultation with the Secretaries of Defense and Veterans Affairs, respectively, establish an HHS AI Task Force to develop a strategic plan to create policies and frameworks for the responsible deployment and use of AI in the health and human services sector, and to identify appropriate guidance and resources to promote deployment of AI in certain areas. These include: use of predictive and generative AI in health care delivery and financing (including quality measurement, program integrity, benefits administration, etc.); incorporation of equity principles in AI-enabled technologies used in the health and human services sectors; and incorporation of safety, privacy, and security standards into the software development process for protection of personally identifiable

---

[55] Executive Order 14110, available at https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/. The full scope of the Executive Order is beyond the focus of this article, but it is an exceptionally wide-ranging document.

[56] These areas have been categorized by the Congressional Research Service. See, Harris, Laurie A., and Chris Jaikaran, "Highlights of the 2023 Executive Order on Artificial Intelligence for Congress," CRS R47843, November 17, 2023, p.1, available at https://crsreports.congress.gov/product/pdf/R/R47843.

[57] Executive Order 14110 § 5.2(e).

[58] Executive Order § 7.2(b)(i).

information.[59] Other instructions to the Secretary of HHS include developing a strategy to determine if AI-enabled technologies maintain appropriate levels of quality[60] and evaluating compliance with federal nondiscrimination laws in the health and human services sector as they relate to the use of AI.[61]

Put simply, at least for the foreseeable future, HHS and other related agencies are likely to issue a range of new regulations and guidance surrounding the use of AI.[62] Some of this will come slower than others, as necessary taskforces and committees must be convened to perform the work, but at least some regulations have already been issued to take AI into account.

For example, the Office of the National Coordinator for Health Information Technology ("ONCHIT" or simply the "ONC") has already published regulatory updates to its EHR certification regulations, as well as its information blocking regulations that impose new requirements for EHRs with "Predictive Decision Support Intervention" (or "Predictive DSI"). This term is defined in the new regulations to mean "Technology that supports decision-making based on algorithms or models that derive relationships from training data and then produces an output that results in prediction, classification, recommendation, evaluation or analysis."[63] The requirements themselves apply more to the functionality of certified EHR software, and thus are more relevant for developers of such software. But they still represent one of the first sets of requirements regarding AI that many health care providers are likely to encounter.

5.2 The States' Response[64]

In the absence of much Congressional action, states have also taken steps to regulate the use of AI. States including California, Colorado, Connecticut, Delaware, Illinois, Indiana, Iowa, Louisiana, Maryland, Montana, New York, Oregon, Tennessee, Texas, Vermont, Virginia, and Washington have all enacted laws focused on AI regulation. Several states have created task forces or similar groups to study

[59] Executive Order § 8(b)(i).

[60] Executive Order § 8(b)(ii).

[61] Executive Order § 8(b)(iii).

[62] Truly, the full scope of the Executive Order is incredibly broad, and includes many different dates by which various agencies must comply with the Order. For a timeline of deadlines and implementation dates, see "Timeline: Biden Administration Executive Order on AI – Key Deadlines and Implementation Dates," Ernst & Young, available at https://www.ey.com/en_us/public-policy/key-takeaways-from-the-biden-administration-executive-order-on-ai. The Congressional Research Service's "Highlights of the 2023 Executive Order on Artificial Intelligence for Congress", available at https://crsreports.congress.gov/product/pdf/R/R47843, is also extremely helpful in navigating the various requirements.

[63] 45 CFR §§ 170.102, et seq. The full Final Rule can be found at 89 Fed. Reg. 1192 (January 9, 2024).

[64] For a general overview of state legislative actions and laws, see Wright, Rachel, "Artificial Intelligence in the States: Emerging Legislation," Council of State Governments, December 6, 2023, available at https://www.csg.org/2023/12/06/artificial-intelligence-in-the-states-emerging-legislation/, and the state-by-state chart provided at "Artificial Intelligence 2023 Legislation," National Conference of State Legislatures, January 12, 2024, available at https://www.ncsl.org/technology-and-communication/artificial-intelligence-2023-legislation.

the impact of AI and to recommend legislation.[65] Others have passed laws to protect individuals from unsafe or ineffective AI systems.[66] Of the states that have enacted AI laws, several have passed laws regarding data privacy, such as requiring that consumers be provided with an option to opt out of data profiling/harvesting, depending on the reason for such harvesting and how it will be used.[67] State laws have also been passed to require transparency regarding how and when an AI system is being used.[68]

In several states, governors have issued executive orders, usually creating executive branch entities to address the usage of AI by state agencies and departments. For example, the governors of

[65] Texas, for example, has passed a law creating an AI advisory council to study the impact of AI use in state government, and make recommendations to state agencies regarding the use of AI. For the full text of the law, HB 2060, see, https://capitol.texas.gov/BillLookup/Text.aspx?LegSess=88R&Bill=HB2060.

[66] For example, Colorado insurance laws prohibit the use of "any external consumer data and information sources, as well as any algorithms or predictive models that use external consumer data and information sources, in a way that unfairly discriminates based on race, color, national or ethnic origin, religion, sex, sexual orientation, disability, gender identity, or gender expression." C.R.S.A. § 10-3-1104.9(b). Other states have taken broader approaches to have state agencies analyze AI systems and provide reports to their respective governors regarding the impact of these systems. See, Wright, Rachel, "Artificial Intelligence in the States: Emerging Legislation," Council of State Governments, December 6, 2023, available at https://www.csg.org/2023/12/06/artificial-intelligence-in-the-states-emerging-legislation/

[67] Montana's Consumer Data Privacy Act, for example, requires that consumers be given the right to (among other rights): (1) confirm whether an entity that owns or controls a company that processes personal data (including by using AI) is processing or accessing the consumers' data, (2) correct inaccuracies in such data, (3) delete data about the consumer, (4) obtain a copy of the data previously provided by the consumer, and (4) opt out of processing such data for targeted advertising, sale of the data, or profiling the consumer "to assist solely automated decisions that produce legal or similar significant effects concerning the consumer." MCA § 30-14-2808(1). Interestingly, the Montana law does not apply to PHI. MCA § 30-14-2804(2)(a).

[68] For example, California law prohibits "bots" from being used to communicate or interact with another person online, when the purpose is to mislead the other person about the bot's artificial identity to deceive the persona bout the content of the bot's communication (e.g., to suggest purchase or sale of goods or services or to influence votes in elections). Cal. Bus. & Prof. Code § 17941(a). If the person using the bot discloses the identity of the bot, they are not liable under this law. The term "bot" is defined to mean "an automated online account where all or substantially all of the actions or posts of that account are not the result of a person." A generative AI operating online would meet this definition. Cal. Bus. & Prof. Code §17940(a).

California,[69] Maryland,[70] Pennsylvania,[71] and Virginia[72] have each issued executive orders instructing state agencies to investigate or form task forces to provide recommendations or other guidance regarding the use of AI in government agencies. In addition to establishing these entities, these orders also often include statements of principles not unlike the Biden executive order, which require state government agencies using AI to ensure privacy, lack of bias, employee empowerment, and transparency.[73]

State attorneys general have also begun to take steps to address artificial intelligence. In September, 2023, the National Association of Attorneys General issued a letter on behalf of 45 state and territorial attorneys general to urge Congress to study how AI could be, and is being, used to exploit children through child sexual abuse material, specifically such material that is created by AI.[74] The letter further requested that Congress propose legislation to protect children from such abuses. In January, 2024, 26 attorneys general, led by Pennsylvania Attorney General Michelle Henry, submitted and signed a comment letter in response to the Federal Communications Commission's notice of inquiry regarding the potential impact of AI on efforts to protect consumers from illegal robocalls and robotexts.[75] In California, Attorney General Rob Bonta began an inquiry into racial and ethnic bias in health care AI algorithms, specifically looking at ways in which software used by health care providers to make decisions affecting access to health care might have discriminatory impacts based on race and ethnicity, and requesting information from hospital CEOs regarding how health care facilities and providers identify

---

[69] Executive Order N-12-23, available at https://www.gov.ca.gov/wp-content/uploads/2023/09/AI-EO-No.12-_-GGN-Signed.pdf.

[70] Executive Order 01.01.2024.02, available at https://governor.maryland.gov/news/press/pages/governor-moore-announces-action-to-transform-maryland-executive-branch-digital-services.aspx#:~:text=The%20executive%20order%20also%20establishes,and%20talent%20in%20state%20government.

[71] Executive Order 2023-19, available at https://www.governor.pa.gov/newsroom/governor-josh-shapiro-signs-executive-order-on-commonwealth-use-of-generative-artificial-intelligence-taking-proactive-step-to-harness-new-technology-while-empowering-state-workforce%EF%BF%BC/#:~:text=The%20Executive%20Order%20enumerates%20ten,safety%20and%20security%2C%20and%20transparency.

[72] Executive Order No. 30 (2024), available at https://www.governor.virginia.gov/newsroom/news-releases/2024/january/name-1019979-en.html#:~:text=RICHMOND%2C%20VA%20%E2%80%93%20Governor%20Glenn%20Youngkin,individual%20data%20of%20all%20Virginians.

[73] See, Maryland Executive Order 01.01.2024.02, pp. 2-3; Pennsylvania Executive Order, pp. 2-3.

[74] "54 Attorneys General Call on Congress to Study AI and Its Harmful Effects on Children," NAAG, available at https://www.naag.org/press-releases/54-attorneys-general-call-on-congress-to-study-ai-and-its-harmful-effects-on-children/.

[75] The letter was signed by the attorneys general for Alabama, Arizona, California, Colorado, Connecticut, Delaware, the District of Columbia, Hawai'i, Illinois, Maine, Maryland, Massachusetts, Michigan, Minnesota, Mississippi, New Jersey, North Carolina, North Dakota, Ohio, Oklahoma, Oregon, Pennsylvania, South Dakota, Tennessee, Vermont, and Washington. The letter is available at https://www.attorneygeneral.gov/wp-content/uploads/2024/01/2024-01-16-Reply-Comment-of-State-AGs-FCC-AI-NOI.pdf.

and address potential disparities created by such AI tools.[76] Thus far, state legislatures have proven far more nimble than Congress[77], and it should be expected that they will continue to regulate AI usage in the years to come. Health care providers and their legal counsel alike will need to monitor these developments, especially if their clients conduct business across state lines or work with those who do.

6. Current and Potential Future Pitfalls

The promise of artificial intelligence in health care is broad. Because of the sheer power to process data – without requiring constant human intervention – deep learning AIs will undoubtedly become more prevalent in an industry that, by its nature, already grapples with massive amounts of data on a regular basis. Artificial intelligence may improve efficiency and quality of life for health care providers, and may help provide better outcomes for patients. But use of AI does not come without risks, and these risks carry with them potential legal consequences. While legislation specifically oriented around AI is still in its infancy, we can expect that many situations involving AI usage will be governed by existing law applied to the new technology.

6.1 "Hallucinations"

One of the greatest risks posed by the use of generative AI in a space as highly regulated as health care is the problem of so-called "hallucinations." Hallucinations in AI occur when a large language model AI (e.g., a chatbot) "perceives patterns or objects that are nonexistent or imperceptible to human observers, creating outputs that are nonsensical or altogether inaccurate."[78] A generative AI "hallucinates" when it is prompted by a user and produces outputs not based on training data, or that do not follow identifiable patterns. While perhaps an overly anthropomorphized explanation, the AI is behaving in a manner that is "similar to how humans sometimes see figures in the clouds or faces on the moon."[79]

By this point, many attorneys are familiar with the mind-boggling story of the law firm that submitted an AI-generated brief, where the AI itself created case citations out of whole cloth.[80] Setting aside the gross professional irresponsibility involved in such a course of action, the incident provides a sobering example of AI hallucination. In my own practice, clients likewise have been the victims of

---

[76] See, "Attorney General Bonta Launches Inquiry into Racial and Ethnic Bias in Healthcare Algorithms," August 31, 2022 press release, available at https://oag.ca.gov/news/press-releases/attorney-general-bonta-launches-inquiry-racial-and-ethnic-bias-healthcare.

[77] Admittedly, this is a low hurdle to clear, given Congress' inability to pass any legislation pertaining to artificial intelligence thus far.

[78] "What Are AI Hallucinations?" IBM, available at https://www.ibm.com/topics/ai-hallucinations.

[79] "What Are AI Hallucinations?" IBM, available at https://www.ibm.com/topics/ai-hallucinations.

[80] Bohannon, Molly, "Lawyer Used ChatGPT In Court – And Cited Fake Cases. A Judge Is Considering Sanctions," Forbes, June 8, 2023, available at https://www.forbes.com/sites/mollybohannon/2023/06/08/lawyer-used-chatgpt-in-court-and-cited-fake-cases-a-judge-is-considering-sanctions/?sh=686925407c7f. In fact, other instances of this same phenomenon have occurred, in spite of the widespread publicity of the first story. See, Merken, Sara, "Another NY Lawyer Faces Discipline After AI Chatbot Invented Case Citation," Reuters, January 30, 2024, available at https://www.reuters.com/legal/transactional/another-ny-lawyer-faces-discipline-after-ai-chatbot-invented-case-citation-2024-01-30/; and, "Michael Cohen Says He Unwittingly Sent AI-Generated Fake Legal Cases to His Attorney," Associated Press, December 30, 2023, available at https://www.npr.org/2023/12/30/1222273745/michael-cohen-ai-fake-legal-cases.

generative AI hallucinations. In one instance, a client was researching a complex issue regarding Medicare's billing and participation requirements and attempted to use two different generative AI programs. The AIs each created Medicare manual sections, including specific citations, that did not exist. The citations themselves either pointed to non-existent sections (e.g., pointing to a chapter in a manual that literally does not exist), or pointed to sections that covered other material (e.g., pointing to the Medicare Benefit Policy Manual Chapter 15, sec. 40.3 and claiming that it covers radiology supervision requirements, when that section in fact addresses the effective date of Medicare opt-outs). In addition, both generative AIs created nonexistent manual text in which the actual substance of the cited material was completely fabricated. This was discovered after attempting to conduct a search for quotes from the fabricated material; the information could not be reproduced either in current manuals or even in prior editions of the manuals. This has occurred more than once with clients.

It is not difficult to imagine the risks posed by using similar technology in a clinical treatment context. If a physician relies on generative AI, either by using a general chatbot like ChatGPT or Google Bard, or by using a generative AI specifically designed for use in a clinical setting, that physician is still potentially risking that the AI will hallucinate an answer to a query. This, in turn, could expose the physician to significant malpractice risk. Just as lawyers who fail to check the veracity and accuracy of citations generated by AI risk potential malpractice actions and imperil their licenses, so too may physicians and other licensed health care providers who fail to check the accuracy of an AI's information in a clinical context.

In practical terms, this risk is best mitigated by first having a clearer understanding of how these tools work. First, generative AIs are designed to *generate* output (either as text or images) based on their training data. They are not, as Michael Cohen discovered to his lament, "super-charged search engine[s]."[81] While popular generative AIs like ChatGPT and Bard are *capable* of behaving *like* a search engine by reproducing specific information upon which they may have been trained, their *primary* functions are (1) to recognize patterns in their training data, and (2) produce output based upon those patterns. As generative AI becomes incorporated into clinical practice tools, presumably it will be trained to minimize and correct hallucinations when they occur, but unless and until someone can create an AI that is functionally immune to the phenomenon, hallucinations will remain a risk. Therefore, the best approach for attorneys and health care providers alike is to check the results that any generative AI produces against whatever sources are cited before relying upon them.

As a legal matter, one other aspect of the malpractice risk issue is that most software licenses already disclaim liability for malpractice (and, indeed, most other defects[82]). Thus, without statutory intervention to impose liability on the creators of AI programs, health care providers likely will be unable to find relief at court from a developer if an AI powering their clinical practice software hallucinates, and

[81] Quote from Michael Cohen, in "Michael Cohen Says He Unwittingly Sent AI-Generated Fake Legal Cases to His Attorney," Associated Press, December 30, 2023, available at https://www.npr.org/2023/12/30/1222273745/michael-cohen-ai-fake-legal-cases. As a separate matter, Cohen's own attorney, David M. Schwartz, failed to review the material that Cohen sent over, and submitted the information without checking it, which ultimately exposes him to licensure risk (since Cohen is himself no longer an attorney and therefore cannot be subjected to further licensure sanctions).

[82] For a longer discussion of this type of language and how it appears in electronic health record software license agreements, see, Shay, Daniel, "A Primer on Electronic Health Record License Agreements," Health Law Handbook, 2006 edition, pp. 425-457; "Downstreamed Physician EHR License Agreements: Understanding the Ebb and Flow," Health Law Handbook, 2008 edition, pp. 45-76.

the hallucination causes a patient to be harmed. Even beyond the risk to the provider's license, it is unlikely that they will be able to shift damages in a malpractice case to the developer. Given the software industry's traditional reliance upon disclaimers of warranties and liability, health care providers and their attorneys will likewise want to push back and demand that the software industry assume more responsibility and stand good behind their products by indemnifying health care providers at least for harm caused by improperly functioning AIs.

6.2 Privacy and HIPAA

Privacy poses another risk with respect to use of AI in the clinical setting, depending on how protected health information ("PHI") is used by the AI itself. At present, most popular deep-learning software (e.g., ChatGPT) are trained using publicly available information. As health-care-specific AIs are developed, especially those used in clinical applications, they will need to be trained on deidentified data at the very least, and potentially on identifiable PHI. The HIPAA implications surrounding the training of AIs will depend heavily on the nature, method, and purpose of using patient health information to train the AI. For example, if the AI being trained is itself limited to only data available within a practice's own records, and the AI will only be used by the practice, and the AI would not be deployed by the developer to other health care providers, then (1) the practice could likely train the AI on identified PHI, and (2) such training would likely qualify as "health care operations," which would not require obtaining a patient authorization.[83] Yet this is not usually how AIs function. Because AIs perform best and recognize patterns more reliably when they can analyze large volumes of data, an AI trained solely on one practice's patient data would not be nearly as useful as an AI that was trained on large volumes of patient information from across multiple practices.

Most developers would likely want access to PHI so that they could aggregate the practice's data along with data provided by other health care providers, and thereby better educate the AI. On the surface, this might seem to be permitted under HIPAA. The Privacy Rule's regulations governing business associate agreements ("BAAs") allow covered entities to include in a BAA the ability for the business associate (in this case, the developer of the AI) to provide "data aggregation services relating to the health care operations *of the covered entity*."[84] "Data aggregation" is itself defined as a business associate combining PHI from one covered entity with the PHI from another covered entity to allow data analysis that relates to the health care operations *of the respective covered entities*.[85] Where this creates difficulties for developers of AI and other software is that, on its face, the regulations only permit the developer/business associate to provide such services to the entities from which it already obtained the PHI used in the data aggregation services; it does not permit them to perform data aggregation to provide services to *future* covered entities.

The workaround for this apparent limitation is to instead use de-identified information, which removes those identifiers specified in regulations, and thereby renders the information no longer within the definition of PHI.[86] Most BAAs include language that permits business associates to use de-identified

[83] See, the definition of "health care operations" at 45 CFR § 164.501, and requirements for authorizations at 45 CFR § 164.504.

[84] 45 CFR § 164.504(e)(2)(i)(B), emphasis added.

[85] 45 CFR § 164.501, emphasis added.

[86] See, 45 CFR § 165.514(a),(b), and (c) for the requirements surrounding de-identified information.

information for their own purposes (and software licenses in particular tend to include such language). However, some AIs have been shown to be capable of re-identifying de-identified information.[87] If AI is capable of reidentifying de-identified information, then HIPAA is in drastic need of an update to take into account the capabilities of this new technology. Some of this may be handled by agencies like ONCHIT, which can mandate certain standards for ensuring privacy that must be maintained within AI-powered EHR software, but other protections may prove necessary. At the time of this writing, the Office for Civil Rights of the Department of Health and Human Services, the government agency tasked with enforcing HIPAA, has not published HIPAA guidance relating to the use or development of AI by covered entities or business associates, but one should expect such guidance to be forthcoming. Likewise, in response to the Biden Administration's Executive Order, other federal agencies will also issue regulations that will necessarily address some of these issues. While many of the requirements that are likely to come will fall on the shoulders of AI software developers, health care providers must be prepared both for regulatory requirements that directly apply to them, as well as those which indirectly affect them.

In the meantime, there are already state laws and caselaw examples that both show ways in which AI usage may implicate such laws, and the limits of those laws. For example, under Illinois' Biometric Information Privacy Act[88], plaintiffs have attempted to file suit for the use of information pertaining to them in the training of various AIs. In Vance v. Amazon, Inc.,[89] two plaintiffs attempted to sue Amazon (as well as Microsoft[90]) for use of images of the plaintiffs without their consent in the training of facial recognition AIs. The plaintiffs had uploaded images of themselves to an image-sharing website (Flickr) which at the time had been operated by Yahoo!, and which was later made publicly available by Yahoo! The images in the Yahoo! dataset were then used in the development of a subsequent dataset by IBM to develop the "Diversity in Faces" ("DiF") dataset that could be used to train AIs to perform facial recognition more accurately by providing a more diverse set of faces. Amazon, in turn, used the DiF dataset to do just this, and the plaintiffs sued. Both the Amazon and Microsoft cases brought by the plaintiffs ultimately failed to withstand motions for summary judgment, because both companies pointed out that the Illinois law in question – which might otherwise have protected the plaintiffs' biometric information – did not include an extraterritoriality provision that would expressly state that the law should apply to conduct occurring outside of Illinois. Because both Amazon and Microsoft's activities had been conducted outside the bounds of Illinois, and because neither company (including the specific researchers engaged in training the AIs) had connections to Illinois, the court upheld the companies' motions for summary judgment.

In Meehan v. VIPKID; VIPKIDS INTERNATIONAL, INC., et al.,[91] the plaintiff brought suit against several companies, alleging – among other facts – that while working for VIPKid (an online teaching

---

[87] For an overview of the study, see Cohen, Jessica Kim, "AI Can Re-Identify De-Identified Health Data, Study Finds," Becker's Health IT, January 3, 2019, available at https://www.beckershospitalreview.com/healthcare-information-technology/ai-can-re-identify-de-identified-health-data-study-finds.html?oly_enc_id=8942E8823912A4S. For the study itself, see Na, Liangyuan, et al., "Feasiblity of Reidentifying Individuals in Large National Physical Activity Data Sets From Which Protected Health Information Has Been Removed With Use of Machine Learning," JAMA Network Open, December 21, 2018, available at https://jamanetwork.com/journals/jamanetworkopen/fullarticle/2719130.

[88] 740 ILCS § 14/1, et seq.

[89] 2022 WL 12306231 (W.D.Wa. 2022).

[90] See, Vance v. Microsoft Corp., 2022 WL 9983979 (W.D.Wa. 2022).

[91] 2024 WL 277846 (E.D.N.Y. 2024).

platform allowing American teachers to connect remotely with Chinese students to teach them English), the plaintiff and other teachers had their biometric data harvested to develop AI and facial recognition technology. The plaintiff alleged various violations of New York common law (specifically fraud, fraud in the inducement, and conversion), as well as New York's General Business Law §§ 349 and 350. Unfortunately for the plaintiff, the many defendants named in the case were able to secure motions to dismiss, many on lack of jurisdiction grounds. However, with respect to VIPKid itself, because of the presence of an arbitration clause in the plaintiff's contract, VIPKid was able to move to compel arbitration under the Federal Arbitration Act ("FAA"). More specifically, the court found that the conduct was within the scope of the arbitration provision in the plaintiff's contract and therefore arbitration was required. Thus, whatever the merits of the plaintiff's claim with respect to (for example) conversion, the matter was ultimately subject to arbitration and the court had no jurisdiction.

These cases illustrate two general issues with the application of existing laws surrounding privacy to circumstances involving AI. First, in the absence of national law, state laws will represent an imperfect patchwork to protect individuals' privacy. In the Vance case, because the training of the AI did not occur in Illinois, Illinois' own law was inapplicable to the conduct. Thus, at least in the analysis of the Western District of Washington, as long as individuals acted outside of Illinois' borders and did not otherwise have any contact or connection with Illinois, individuals who uploaded their images to the original Flickr website could not avail themselves of the law's protections. Second, in cases involving contractual violations or conduct that is itself covered by a contract, if that contract contains an arbitration clause, courts are likely to uphold motions to compel arbitration under the FAA; this will continue, regardless of what other laws or regulations are passed, unless Congress sees fit to create exceptions to FAA with respect to cases dealing with AIs.

Nevertheless, one can imagine other scenarios in which a health care provider runs afoul of such laws. For example, even in the absence of extraterritorial application of laws, one can envision a situation in which a health care provider takes patient information (not de-identified) and provides it to a business associate who, in turn, improperly uses that information in the development of data sets or training of an AI in violation of a state privacy law. In such circumstances, while there may be no private right of action under HIPAA, there would at least be a sufficient state connection to bring suit under a state privacy law because of the health care provider's actions in providing access to such information. With this in mind, health care providers and their legal counsel should be mindful of state privacy laws as they continue to develop with respect to AI.

6.3 Fraud and Abuse

One final area of concern for health care providers is fraud and abuse. As artificial intelligence is used more frequently for the preparation and submission of claims, as well as in the generation of documentation, there is risk that an AI will "hallucinate" and submit claims for the wrong codes, either by upcoding or by failing to generate sufficient documentation to support an otherwise valid code. When this happens, who will bear the burden under the Federal False Claims Act ("FCA")?[92] One would normally expect that the health care provider would be most in harm's way. In most cases, if the provider's name is on the claim, they bear the liability for false claims; even if the claims were prepared and submitted by staff or a billing company that habitually miscodes, the provider is ultimately responsible for the submission of the claims. One would expect that this would continue as AI becomes

[92] 31 U.S.C.A. § 3729.

more integrated into practice management software packages, including for documentation and claims submission purposes.

There is some limited precedent for FCA liability of software vendors, but such cases have been relatively rare up to the present. In 2017, the Department of Justice entered into a settlement with eClinicalWorks ("ECW"), one of the largest vendors of EHR software in the country.[93] The $155 million settlement resolved allegations that ECW had misrepresented its software's capabilities and falsely obtained certification for its EHR software by concealing from the certifying entity that it failed to comply with certification requirements, such as failing to record user actions in an audit log and failing to record diagnostic imaging orders and perform drug interaction checks. This, in turn, generated false claims under the EHR Incentive Program established as part of the American Recovery and Reinvestment Act of 2009[94] by marketing the software to potential purchasers who would buy it because it was certified software under the program, and in turn receive incentive payments.

In the scenario where the AI-powered EHR itself malfunctions, however, the developer may face its own liability (e.g., for causing the situation in which the false claim was generated), but that will not protect the health care provider, nor allow them to substitute the developer in their place. Instead, the health care provider's best hope is to be able to, after repaying whatever overpayments are generated by the software, sue the EHR developer. Yet this once again puts the health care provider at the mercy of the text of their EHR software license agreement. If, as is common, the agreement includes disclaimers of liability, and does not include indemnification language, the health care provider will simply have to bear the expense of the overpayments. In light of this, as AI becomes even more integrated into practices, the health care industry ought to demand that where generative AI is performing so many automated tasks for the provider, that the developers stand good behind their products. This will require pushing back on developers in license negotiations and insisting upon indemnification language and carve-outs for liability disclaimers in the event of overpayment generation.

7. Conclusion

Artificial intelligence is the proverbial genie that cannot be put back into the lamp. Like the genie, AI also can prove tricky and capricious as well as occasionally fanciful as it hallucinates results. Nevertheless, it will increasingly become part of the health care industry both in clinical and administrative contexts. Given the potential risks posed by deep-learning models, and especially with generative AIs, health care providers and their legal counsel must remain vigilant about how such AIs are used in the health care context, what contractual protections exist between the provider and the AI developers. Likewise, both providers and legal counsel will need to closely monitor both state and federal efforts to regulate AI in the health care context and beyond. The pace of these developments – both the adoption of AI technology, and the growth of a regulatory framework – will likely accelerate in the coming years, as more providers adopt and use the technology, regulators recognize the need to control or at least shape such adoption and use, and that, in turn, provides guidance to providers and technology developers to further implement AI, creating a kind of feedback loop.

---

[93] "Electronic Health Records Vendor to Pay $155 Million to Settle False Claims Act Allegations," May 31, 2017, available at https://www.justice.gov/opa/pr/electronic-health-records-vendor-pay-155-million-settle-false-claims-act-allegations.

[94] P.L. 111-5 (2009).