

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

**PLEASE DO NOT COPY, DISTRIBUTE, OR CITE WITHOUT THE PERMISSION OF
THE AUTHOR**

**The Rise of the Machines:
Legal Risks in the Era of Increasing Health Care AI Use**

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, 2026 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

**The Rise of the Machines:
Legal Risks in the Era of Increasing Health Care AI Use**
By Daniel F. Shay, Esq.

Table of Contents

I.	Introduction	1
II.	Current AI Usage within Health Care	2
A.	A Brief AI Primer.....	3
B.	AI in the Health Care Industry	7
1.	Clinical Use of AI	7
2.	Documentation-Related Uses of AI	12
3.	Claims-Oriented Uses of AI.....	15
4.	Payor Usage of AI.....	17
III.	Legal Risks & Considerations	20
A.	Malpractice Risk	20
B.	False Claims Act Risks	23
C.	HIPAA Risks	25
D.	Reimbursement and Claims Denial Risks.....	29
IV.	Risk Mitigation Strategies in the Era of AI	34
A.	Contractual Clauses and Legal Challenges.....	34
B.	Compliance Plans and Internal Audits.....	39
C.	The Political Process and State Law/Regulation	42
V.	Conclusion	45

I. Introduction

In the last few years, the popularity and seeming omnipresence of artificial intelligence (“AI”) has exploded. Bursting into the cultural consciousness with the launch of ChatGPT in November 30, 2022,¹ AI has quickly grown to become a part of daily life. The health care

¹ “Introducing ChatGPT,” available at <https://openai.com/blog/chatgpt>, November 30, 2022.

industry has likewise seen a dramatic increase in the use of AI across many different aspects of the industry. Uses that at one point were merely speculative are becoming more commonplace.

In January, 2026, it was estimated that some 40 million people had used ChatGPT for some form of health care treatment.² Users have queried the AI tool about all manner of information, including asking ChatGPT to explain medical bills, obtain guidance regarding health insurance, asking about symptoms and whether to contact a doctor or seek emergency care (in much the way a patient might ask an on-call physician), etc. Physicians are likewise interacting more with AI. A survey published by the American Medical Association in 2025 found that two out of three physicians are using some form of health care AI, representing an increase of 78% from 2023.³ The survey examined physicians' attitudes regarding, and use of, AI comparing responses from 2023 and 2024.⁴ For example, the report demonstrated a small increase (+5%) in physician enthusiasm and excitement regarding AI, which corresponded with a drop in physicians who were equally concerned and excited, and those who were more concerned than excited, suggesting that respondents may be shifting their perspectives with regard to AI from a more skeptical or dubious stance to one that is positive and enthusiastic about AI.

² Morrone, Megan, "Exclusive: 40 Million People Turn to ChatGPT for Health Care," Axios, January 5, 2026, available at, <https://www.axios.com/2026/01/05/chatgpt-openai-health-insurance-aca>.

³ Henry, Tanya Albert, "2 in 3 Physicians Are Using Health AI – Up 78% From 2023," AMA News Wire, February 26, 2025, available at, <https://www.ama-assn.org/practice-management/digital-health/2-3-physicians-are-using-health-ai-78-2023>.

⁴ See, "AMA Augmented Intelligence Research: Physician Sentiments Around the Use of AI in Health Care: Motivations, Opportunities, Risks, and Use Cases – Shifts from 2023 to 2024," published February, 2025, available at, <https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf>.

The integration of AI in the health care industry promises great potential, both from an administrative perspective, and from a clinical one. Artificial intelligence may ease the burdens associated with a variety of clerical scutwork, just as it may help clinicians to diagnose and treat their patients more effectively and efficiently. But the use of AI is not without its share of risks, as well. This article will explore current uses of AI in the health care industry, with an eye towards (but not exclusively fixed upon) physician and physician practice usage. It will examine certain legal risks that may arise from different types of AI usage, and explore how physicians, with the assistance of legal counsel, can work to mitigate those risks.

II. Current AI Usage within Health Care

Within the context of health care, one need no longer merely speculate about how AI may be used; there are multiple examples of current uses of AI within the health care industry. These uses tend to fall into certain general categories, such as clinical uses (e.g., for diagnostic imaging), documentation-related uses (e.g., AI scribing), claims-oriented uses (e.g., automated billing and claims submission), and usage by payors (e.g., AI's influence in various payment control mechanisms like prior authorizations). Before exploring these different uses, however, it may be helpful for the reader to have a better understanding of how AI actually functions, so as to better grasp why and how AI is being used in these different ways.

A. A Brief AI Primer

“Artificial intelligence” is not merely a marketing buzzword; the concept itself derives from computer science and theories about how computers could potentially function.⁵ In broad

⁵ For a deeper exploration of the background of AI, see, Shay, Daniel F., “Does ChatGPT Dream of Electric Sheep? – Legal Implications of Artificial Intelligence in Health Care,” Health Law Handbook, 2024-2025 ed., pp. 135-176.

strokes, one of the more common definitions of AI is software that is capable of fooling a human into believing that they are talking to another human instead of interacting with software. Put another way, true “AI” is software that is capable of human-like presentation of information and information processing. Much of the AI that currently exists is based on “machine learning,” which is done through the use of “neural networks” and “deep learning” software models. Without diving too deeply into the terminology here, the software is capable of very rapidly processing information, often through the use of multiple parallel data processing nodes, and then providing a user with processed data.

What artificial intelligence is truly best at, though, is pattern recognition and pattern reproduction. The software itself is not creative or capable of inductive reasoning. In spite of much of the anthropomorphic language used to describe AI (e.g., “neural” nodes, software “hallucinations,” “learning,” etc.), the software is not alive nor capable of what humans can do. The anthropomorphic language is simply convenient shorthand used to describe complicated software processes in a way that people can more easily understand.

Most generative AI, such as ChatGPT or Google Bard for “chatbot” style programs, or DALL-E or Midjourney for image generation, are not actually creative, nor are they super-powered search engines. Instead, what they do is to generate statistically probable results in response to user prompts (be they queries or other inputs), when compared to the repertoire of data that has been programmed into the AI (i.e., data on which the AI has been “trained”). This helps explain why certain AI-drawn images can have a very uniform look to them, and why text generation may end up seeming generic or bland. This is the software functioning correctly; the responses end up being perceived as generic because, in a manner of speaking, they are. Those aspects that we perceive as “generic” are simply the most statistically probable results of the

prompt, because those results likely appear most often within the data on which the AI has been “trained.”

To illustrate this point, consider the results when I prompted ChatGPT v5.2 in its free access version⁶ to do nothing more than “Describe a sunset” three separate times. The results are as follows:

Response 1:

“The sun eases itself down, unbuttoning the day.

“Light softens and stretches, turning the sky into a slow spill of color—gold thinning into amber, amber deepening to rose, then bruised purple at the edges. Clouds catch fire for a moment, glowing as if lit from within, before dimming into silhouettes.

“The air feels different, cooler, quieter, like the world is holding its breath. Shadows lengthen, sounds mellow, and the last light lingers just long enough to feel deliberate—an unhurried goodbye before night takes over.”

Response 2:

“The sun leans toward the horizon like it’s tired of being bright. Light loosens, softens, spills into colors it never uses at noon—amber, rose, bruised purple, a last flare of gold. Shadows stretch and yawn. The sky looks painted, then repainted, then quietly erased at the edges

“For a moment, everything feels held in place: the air cooler, sounds gentler, time slowed just enough to notice. Then the light slips away, not dramatically, just honestly, leaving behind a hush and the sense that something beautiful passed through and didn’t ask to stay.”

Response 3:

“The sun drifts toward the edge of the world and the day exhales. Light thins, stretches, turns liquid. The sky blushes—first gold, then peach, then a deeper embered orange—like it’s remembering warmth as it lets go of it. Clouds catch fire for a moment, glowing from within, before cooling into mauve and violet. Shadows grow long and soft, sounds hush, and everything feels briefly suspended, as if the earth is holding its breath while the light slips away.”

⁶ See, <https://chatgpt.com>.

A careful observer will note the similarities between the three responses. The light “softens” or turns “liquid” or “spills.” The colors referenced are similar as well, with gold, amber, rose, bruised purple, or peach, or “embered orange.” The light may also “stretch” or “slip away” as shadows “lengthen” or “grow long.” The reason these items appear in each of the responses is that they likely appear in a statistically significant portion of the various descriptions of sunsets on which ChatGPT v5.2 has been “trained.” In other words, the ChatGPT software has reviewed scores of sunset descriptions, and seen references to these words and phrases or synonymous versions thereof. So, when prompted to “Describe a sunset,” the software presents the most statistically probable answer, which usually contains descriptions of amber/gold/embered orange light that softens, lengthens, and slips away. Put still another way, the software is simply picking the most commonly occurring phrases, and stringing them together in response to the prompt; pattern recognition, and pattern repetition at its core.

This functionality also helps explain what are referred to as AI “hallucinations.” These occur when the AI software recognizes a pattern that humans themselves do not recognize, but which likely does match the algorithm’s function. For example, if an AI chatbot invents caselaw that does not actually exist when prompted to provide cases regarding a given legal issue, this is fundamentally no different from the AI describing a sunset as in the three responses above. Rather than functioning as a caselaw search engine, the AI is simply assembling a “case” out of statistically likely arrangements of words and phrases based on the cases on which it has been “trained,” and presenting that collection of words and phrases to the user. This still qualifies as an “hallucination” by the AI, just as much as drawing a picture of someone with too many fingers would be, but it is also worth noting that the AI in this circumstance is actually doing what it was programmed to do, and the problem is more that the user is using the AI in an

inappropriate manner (i.e., using a chatbot to conduct legal research). It is likely safer to use purpose-built tools, although even in such circumstances, the risk of hallucinations and other technological failures still exists. This presents specific problems for health care providers, including physicians, depending on how they use AI.

B. AI in the Health Care Industry

Within the health care industry, AI uses can be broken into several different categories. These include: (1) clinical uses; (2) documentation- and administration-related uses; (3) claims-oriented uses; and (4) uses by payors. A clearer understanding of how AI is actually being used can help to anticipate the likely areas in which problems will arise.

1. Clinical Uses of AI

While adoption of AI in the clinical space has increased in recent years, full integration of AI into the clinical setting has yet to truly occur. However, there are examples of AI being gradually deployed through various demonstrations and pilot projects that raise questions about how AI may be used in a future that draws ever nearer. For example, in October, 2025, an AI program known as “Dr. CaBot” was developed at Harvard Medical School to serve as an educational tool. The AI was trained on roughly 100 years of clinicopathological conferences, wherein a patient’s doctor provides a detailed presentation of a patient’s case, and an expert who is not themselves involved in the care likewise offers a presentation and their own differential diagnosis. The doctors explain their reasoning throughout the process, and ultimately present what each believes to be the most likely diagnosis of the patient’s complaint.

In this instance, a physician presented alongside the Dr. CaBot AI on a case involving a 36-year-old man who was admitted to the hospital with abdominal pain, fever, hypoxemia, and

thrombocytopenia. Dr. CaBot presented alongside Dr. Gurpreet Dhaliwal, and the results were published in the New England Journal of Medicine.⁷ Dr. Dhaliwal presented first, and ultimately concluded that the cause was a perforation of the duodenum by a foreign body (i.e., a toothpick). Dr. CaBot presented second, ultimately arriving at the same conclusion, but through different reasoning.

Artificial intelligence has also been used in several instances in the realm of robotic surgery. In June, 2025, a robot performed a realistic surgery without the physical assistance of a human being. Instead, the human instructed the robot using voice commands, but did not otherwise physically intervene, much the way an attending physician would coach a student or resident in performing a surgery.⁸ A previous robot, known as the Smart Tissue Autonomous Robot (“STAR”) performed the first autonomous surgery on a pig, using laparoscopy. However, the STAR surgery required that the pig have specially marked tissue, and the robot followed a pre-determined surgical plan. In the 2025 case, however, the newer robot adapted in real-time and made “decisions” in the moment, rather than simply following a programmed path, to perform a gallbladder removal procedure, requiring 17 specific tasks to be performed, including the identification of various ducts and arteries, carefully grasping them, placing clips, and then severing portions with a pair of scissors. To do this, the robot used machine learning software similar to ChatGPT, whereby the robot was trained on videos of Johns Hopkins surgeons performing the same surgery on pig cadavers. The robot also received reinforcement training by

⁷ See, Dhaliwal, G., Hood, Manrai, Buckley, Asombang, and Hohmann, “Cas 28-2025: A 36-year-old Man with Abdominal Pain, Fever, and Hypoxemia,” New England Journal of Medicine, 393; 14, pp. 1421-1434 (2025), available at, <https://www.nejm.org/doi/pdf/10.1056/NEJMcp2412539>.

⁸ See, Rosen, Jill, “Robot Performs First Realistic Surgery Without Human Help,” Johns Hopkins University press release, July 9, 2025, available at, <https://hub.jhu.edu/2025/07/09/robot-performs-first-realistic-surgery-without-human-help/>.

human-added captions describing the tasks. In the final surgery, the robot performed with 100% accuracy, although the process took longer than would be expected by a human surgeon. Still, the results were compared to that of an “expert” surgeon.

Relatedly, a surgeon in Chile performed an AI-guided camera-assisted gallbladder surgery by himself in September, 2025.⁹ The surgery itself used magnetic surgical instruments, alongside software that independently directed the surgical camera without human control, with the camera programmed to follow along with the surgeon’s tools and hands. Similarly, a robotic surgery was performed in China at the Chinese University of Hong Kong, involving an AI-directed surgical arm that acted as a “third arm” to assist surgeon-controlled robotic arms. Similar to prior surgeries, the Hong Kong surgery was to remove a pig’s stomach. The AI surgical arm assisted in tasks like retracting tissue, picking up gauze, and clipping blood vessels, with four out of six clippings performed successfully.

Still more robotic surgeries were reviewed as part of a survey paper on the subject of robotic and AI-assisted surgery, looking at twenty five different peer-reviewed sources to examine both the role and effectiveness of AI-assisted robotic surgery.¹⁰ The article ultimately concluded that there is currently a “paradigm shift” underway in the precision, decision-making, patient safety, and efficiency of robotic surgeries, and that AI integration seems to have enhanced performance and outcomes. One method the article referenced was “digital twin” surgeries, where a “virtual patient” is created to allow for pre-operative rehearsal and risk assessments.

⁹ See, Blum, Karen, “AI is enabling robots to assist in surgery. What to know,” Association of Health Care Journalists, September 17, 2025, available at, <https://healthjournalism.org/blog/2025/09/ai-is-enabling-robots-to-assist-in-surgery-what-to-know/>.

¹⁰ See, Wah, Jack Ng Kok, “The Rise of Robotics and AI-Assisted Surgery in Modern Healthcare,” Journal of Robotic Surgery, 19: 311, June 20, 2025, available at, <https://doi.org/10.1007/s11701-025-02485-0>.

The article also discussed how AI could assist in automating routine surgical maneuvers that require consistency, which could in turn lead to fewer intraoperative complications.

Another area of AI development and integration with clinical practice is the increase in AI driven clinical decisions support systems (“CDSS”). One description of CDSS tools defines them as systems “intended to improve health care delivery by enhancing medical decisions with targeted clinical knowledge, patient information, and other health information.”¹¹ In broad strokes, the technology takes individual patient information and compares that data against computerized clinical knowledge databases. Based on the comparison, the software provides recommendations to the physician, and the physician may choose to follow them or not. The software is not intended to substitute physician information or knowledge, but rather to supplement it.

The use of CDSS generally is nothing new, and was incentivized as part of the Health Information Technology for Economic and Clinical Health Act of 2009 (“HITECH”) that promoted the adoption of certified EHRs more broadly.¹² Federal regulations governing the certification of electronic health records technology include the capacity to provide clinical decision support as a requirement for certification.¹³ As a result, certified EHRs already incorporate CDSS. However, the introduction of AI into the CDSS process is a new development. According to Andrew Taylor, an associate professor of emergency medicine, director of emergency department clinical informatics, and associate director of informatics and

¹¹ Sutton, Reed, David Pincock, Daniel C. Baumgart, Daniel C. Sadowski, Richard N. Fedorak, Karen I Kroeker, “An Overview of Clinical Decision Support Systems: Benefits, Risks, and Strategies for Success,” NPJ Digital Medicine, 2020, vol. 17, available at, <https://www.nature.com/articles/s41746-020-0221-y>.

¹² P.L. 111-5. The HITECH Act was itself part of the American Recovery and Reinvestment Act of 2009.

¹³ 45 CFR § 170.102. The definitions of both “Base EHR” and “Qualified EHR” include this language.

data science research at Yale University School of Medicine, "It's about cultivating AI tools that support clinicians rather than replace them, insuring technology is a means to enhance the human-centric care that lies at the heart of medicine."¹⁴ Taylor explains that, within the emergency room context, AI-enhanced CDSS could improve the precision and speed of triage, and rapidly assess how severe a patient's condition actually is.

AI-enhanced CDSS tools can be provided as standalone products, or integrated into larger EHR software suites. One standalone product is UpToDate Expert AI, which is a platform developed by Wolters Kluwer. The platform provides AI CDSS for professionals and small groups, and offers an enterprise-level version of the same software. It describes itself as an "evidence-based GenAI in clinical decision support," and promotes itself as offering transparent reasoning with "traceable answers."¹⁵ By contrast, Oracle Health provides an EHR suite that includes an "embedded" AI which, according to Oracle, is "built in—not bolted on," and provides "clinical intelligence" offering AI-driven suggestions about patients. These suggestions include potential risks and gaps in care, as well as financial concerns for patients, and their risk of readmission.¹⁶

¹⁴ Siwicki, Bill, "How Integrating AI and Clinical Decision Support Systems Can Help in the ER," HealthIT News, January 12, 2024, available at, <https://www.healthcareitnews.com/news/how-integrating-ai-and-clinical-decision-support-systems-can-help-er>.

¹⁵ For more on UpToDate Expert AI, see, <https://www.wolterskluwer.com/en/solutions/uptodate/ai-clinical-decision-support?compositeLink=%7BF33D03DC-5B75-46BD-8A8F-A8E4BBB9D1FB%7D>. Note that this article does not take a position on the efficacy or appropriateness of any products discussed herein; rather these products are presented merely as examples of currently available AI-driven CDSS.

¹⁶ See, "Oracle Health EHR" website, specifically the "Clinical intelligence" subsection, available at, https://www.oracle.com/health/clinical-suite/electronic-health-record/?elqak=8AF5A6FB517A9D3DD13D760375C6B97F21ACF5D75730F5D3F6319915949606BBB727&source=:ad:pas:go:dg:a_nas:l5:RC_GOOG240501P00011C00560:MainAd+:ow:lp:cpo::&gad_source=1&gad_campaignid=21316894519&gbraid=0AAAAADtVWWwgjjBKX_-l6LHbX5Rg6ncUs&gclid=Cj0KCQiAnJHMBhDAARIsABr7b85zfOsuydmgvTn8_G-NIQ3kn3AII1f5s6_6Yz2JqVBftlmkurHsdrYaAsnVEALw_wcB#clinical-intelligence.

2. Documentation-Related Uses of AI

Adoption of AI in documentation services has been somewhat more rapid, and presents a more attractive use case for many physicians. Physicians have likewise been much more comfortable with the idea of using AI for these purposes. In a study published in 2025, the American Medical Association conducted a survey of 1,183 Physicians regarding their views around artificial intelligence in medicine.¹⁷ The survey found a noticeable increase in enthusiasm regarding use of AI for documentation purposes, with the top three anticipated documentation uses being: (1) documentation of billing codes, medical charts, or visit notes; (2) creation of discharge instructions, care plans, and/or progress notes; and (3) generation of chart summaries. Enthusiasm for these use cases grew by an average of slightly more than 6% for each use from 2023 to 2024, with the creation of discharge instructions, care plans, and/or progress notes increasing by 9% from year to year.

Another study, also published in 2025, was conducted by the American Academy of Family Physicians and Rock Health, and found that half of the respondents had already used AI for at least one use case at work, with 32% having used AI for clerical support such as drafting messages, or performing “ambient documentation” (i.e., where the software records the visit, and generates a transcript automatically).¹⁸ More than 1,200 physicians participated in the survey, consisting of primary care and other family medicine clinicians, to ask about their perceptions

¹⁷ See, “AMA Augmented Intelligence Research: Physician Sentiments Around the Use of AI in Health Care: Motivations, Opportunities, Risks, and Use Cases – Shifts from 2023 to 2024,” published February, 2025, available at, <https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf>.

¹⁸ See, Ryan, Bob, Kavya Tallapaka, Chris Lew, “AI Is in the Doctor’s Bag – And Primary Care is Ready to Use It,” Rock Health, February 26, 2025, available at, <https://rockhealth.com/insights/ai-is-in-the-doctors-bag-and-primary-care-is-ready-to-use-it/>.

and use of AI in health care. Roughly 70% of the respondents expressed that they expected AI to improve their own wellbeing, and 66% of respondents expected improvements to their work loads.

Physicians in Philadelphia, interviewed in late 2025, also reported using AI for auto-scribing purposes.¹⁹ One physician reported that they had found AI auto-scribing to be a “huge time saver,” after having used it for a year, and stated that it allowed a greater degree of direct interaction with patients. Rather than staring at a computer while taking notes, the physician was able to interact directly with the patient, and be sure that the AI would effectively transcribe the exchange. The same article reported that more than 80% of family medicine physicians at Temple Health, and over 70% of all physicians at Penn Medicine are already using generative AI tools in their practices. Jefferson Health estimated that AI saves each of its physician an hour of time per day, while WellSpan Health claimed that, since adopting AI technology, physicians were able to save enough time to see between two and five additional patients each day.

Epic released its own “AI Charting” software module very recently, as part of its “Art” AI software suite for clinicians, which Epic incorporates into its EHR platform.²⁰ The software records patient visits, drafts a clinical note, and then prepares orders based on the conversation between the physician and the patient, all in real time. Clinicians can also modify the formatting of notes through the use of voice commands, so that future AI-generated notes are more personalized to the physician’s preferred style. Epic notes that its Art AI platform has grown to a

¹⁹ See, Yu, Alan, “Why Many Philadelphia Doctors Now Use AI to Record Patient Visits,” WHYY, December 3, 2025, available at, <https://whyy.org/articles/artificial-intelligence-doctors-patient-visits/>.

²⁰ See, “Epic AI Charting Rolls Out Alongside Expanding Set of Built-in AI Capabilities,” Epic news release, February 4, 2026, available at, <https://www.epic.com/epic/post/epic-ai-charting-rolls-out-alongside-an-expanding-set-of-built-in-ai-capabilities/>.

point where it is used “over 16 million times each month a nearly 3x increase in usage from November 2025.”²¹

Some physicians are also using an AI product to “chat” with their EHR software, appropriately named “ChatEHR.”²² The program was developed in 2023, and applies large language models (similar to how ChatGPT works). This allows the physician to use natural language prompts to interact with a patient’s electronic medical records. The data for the software was gathered from patient medical records, allowing the software to summarize patient charts, and allowing the physician to ask questions about the chart. This is, in essence, a faster way of “searching” the chart, rather than trying to think of the ideal search terms. The software, however, is only in a testing phase, and is deployed at Stanford University among only 33 clinicians, comprised of a mix of physicians, physician assistants, nurses, and nurse practitioners. During testing, users monitor the software’s performance, and interact with it to enhance its accuracy and usefulness, serving as an example of human-led machine learning. The developers noted that ChatEHR is not meant to provide medical advice, but instead is designed to help speed information gathering. The software is also being modified to add functionality to do things like review a patient’s history and record to determine if transfer to another hospital is appropriate, or to determine patient eligibility for hospice care, or recommendations for heightened supervision of post-surgical care, but these are still in development.

²¹ “Epic AI Charting Rolls Out Alongside Expanding Set of Built-in AI Capabilities,” Epic news release, February 4, 2026, available at, <https://www.epic.com/epic/post/epic-ai-charting-rolls-out-alongside-an-expanding-set-of-built-in-ai-capabilities/>.

²² Armitage, Hanae, “Clinicians Can ‘Chat’ with Medical Records Through New AI Software, ChatEHR,” Stanford Medicine News Center, June 5, 2025, available at, <https://med.stanford.edu/news/all-news/2025/06/chatehr.html>.

3. Claims-Oriented Uses of AI

Artificial intelligence can also be used to both automate and improve the efficiency of a variety of claims-oriented activities, including both claims preparation and submission, and auditing and internal oversight and compliance efforts. The American Academy of Professional Coders has identified several activities that AI can perform in the realm of billing and claims.²³ These include automated suggestions, based on natural language processing applied to information contained in electronic medical records and other related data, which can be used to determine appropriate ICD-10-CM, CPT, and HCPCS codes; real-time coding edits that flag incomplete or non-compliant documentation that is likely to result in a claim denial; predictive denial analysis using machine learning to review prior denied claims, the reasons for the denials, and then use that information to flag high-risk claims prior to submission; and, AI can be used to assist in audit preparation to ensure that audit trails are maintained effectively, thereby making it easier and faster to respond to payor audits (including government audits), and to demonstrate coding logic and decisions through reproduceable and transparent mechanisms.

Usage of AI in billing and coding can also reduce billing staff burdens by allocating certain repetitive tasks, such as verifying patient eligibility, to AI automation. For example, coding and billing AI can be programmed to prompt users with different types of alerts to help improve coding outcomes. When new coding requirements are published, this information can be added automatically to the software, and the AI can then prompt coders to take the new data

²³ See, Balaji, Vanmathi, "AI and Medical Coding," AAPC Knowledge Center, November 1, 2025, available at, <https://www.aapc.com/blog/93444-ai-and-medical-coding/?srsltid=AfmBOooCObE3QpLF7kY5dBtqNAKehDsU84FGwDaQUSStWFqwMELXuPZ>.

into account. Likewise, coders can be notified by the AI about patient charts that require additional review.

Some EHR platforms are already incorporating this technology into their software suites. For example, Athenahealth’s athenaOne EHR has incorporated AI across multiple different facets of its software, and advertises that “every payer update, every clean claim, and every denial resolution strengthens the platform for all customers in real time – helping all parties move closer to desired outcomes throughout each phase of revenue cycle management.”²⁴ Athenahealth further claims that its AI-powered software can improve the speed at which claims are created through automatic claim creation, where a claim is automatically generated after each patient encounter. The software identifies possible claims problems in real time, based on data analyzed from 160,000 providers in Athenahealth’s network, which in turn allows users to identify and fix more rapidly potential problems that might result in a claim denial. Athenahealth also advertises that the AI can help streamline the prior authorization process, and boasts a greater-than-98% success rate for authorizations.

Similarly, Epic already includes “Penny,” its AI for operational workflows, as part of its hospital-facing software suite.²⁵ Penny, Epic notes, is capable of automating billing and coding. Epic has stated that the software has reduced coding-related denials by 20%, and is capable of drafting letters to challenge medical necessity denials 23% faster than without using the software.

²⁴ “Top Ways AI-Native athenaOne EHR lifts RCM Capabilities,” athenahealth.com blog entry, March 11, 2025, available at, <https://www.athenahealth.com/resources/blog/reduce-administrative-burden-with-athenaone-ai-tools>.

²⁵ See, “Epic AI Charting Rolls Out Alongside Expanding Set of Built-in AI Capabilities,” Epic news release, February 4, 2026, available at, <https://www.epic.com/epic/post/epic-ai-charting-rolls-out-alongside-an-expanding-set-of-built-in-ai-capabilities/>.

4. Payor Usage of AI

Physicians and physician practices are, of course, not the only ones adopting AI; payors have been fast to implement the technology in several controversial ways, primarily focused around claims processing (especially denials), and prior authorizations. This use of AI in the realm of claims submissions and denials, as well as with prior authorizations, has been described as part of an “arms race” between payors and providers. Aron Klein, vice president of finance operations and supply chain at Carle Health in Urbana, IL has specifically described it as such, and noted “payers are definitely ahead.”²⁶

UnitedHealth Group, for example, has reported that it has 1,000 uses for AI technology currently in production across all of its various businesses, including insurance, care delivery, and pharmacy services.²⁷ These uses are reported to be a mix of generative and traditional AI, including uses in transcribing clinician visits, help with claims processing, and support chatbots. The 1,000 uses figure represents a marked increase in AI use by the company; in 2024, the company reported only 500 AI applications then in use. UnitedHealth Group has also used AI in direct consumer calls through the deployment of AI phone agents.²⁸

²⁶ See, Condon, Alan and Madeline Ashley, “The ‘AI Arms Race’ Reshaping Healthcare,” Becker’s Hospital Review, November 26, 2025, available at, <https://www.beckershospitalreview.com/finance/the-ai-arms-race-reshaping-healthcare/>. Others have likewise described the AI “arms race” between payors and providers. See, Mello, Michelle M., Artem A. Trotsyuk, Abdoul Jalil Dijberou Mahamadou, and Danton Char, “The AI Arms Race in Health Insurance Utilization Review: Promises of Efficiency and Risks of Supercharged Flaws,” *Health Affairs*, 45 No. 1, pp. 6-13, January, 2026, available at, <https://www.healthaffairs.org/doi/10.1377/hlthaff.2025.00897>.

²⁷ See, Emerson, Jakob, “UnitedHealth Has 1,000 AI Use Cases: Report,” Becker’s Payer Issues, May 6, 2025, available at, <https://www.beckerspayer.com/payer/unitedhealth-has-1000-ai-use-cases-report/>.

²⁸ The author has been unable to find reporting on the number of consumers reported to have shouted “Live agent!” to circumvent such systems, but speculates that it would be a statistically significant figure.

Other health plans are using AI in their own call centers, as well as to increase the speed of claims processing, and to “manage utilization.”²⁹ Conner Coles, the director of Public Affairs for AHIP has stated,

“Health plans recognize that AI has the potential to drive better health care outcomes – enhancing patient experience, closing gaps in care, accelerating innovation, and reducing administrative burden to improve the focus on patient care...As adoption grows, strong industry standards and balanced national policies are critical to foster innovation while ensuring safety and trust for patients, providers and other health care stakeholders.”³⁰

Payors have reported using AI in other areas, including predictive risk modeling, administrative simplification, and payment integrity and fraud detection. AHIP has argued that although AI may accelerate prior authorization reviews and approvals, it is not used to automatically deny requests based only on automated decision-making; instead, AHIP claims that licensed clinicians review the requests.

Nevertheless, physicians are concerned about the use of AI by payors. The AMA, in a recent survey, found that 61% of respondents were concerned that health plans’ use of AI has increased the rate of prior authorization denials, worsening patient harms that would otherwise be avoidable, and generating waste both in the present and the future.³¹ The same press release observed that, in spite of a 2018 agreement with the insurance industry to address the burdens of

²⁹ See, Condon, Alan, “How Payers are ‘Embracing AI’s Benefits’ – and What’s Next: AHIP,” Becker’s Payer Issues, December 15, 2025, available at, <https://www.beckerspayer.com/payer/how-payers-are-embracing-ais-benefits-and-whats-next-ahip/>.

³⁰ Condon, Alan, “How Payers are ‘Embracing AI’s Benefits’ – and What’s Next: AHIP,” Becker’s Payer Issues, December 15, 2025, available at, <https://www.beckerspayer.com/payer/how-payers-are-embracing-ais-benefits-and-whats-next-ahip/>.

³¹ “Physicians Concerned AI Increases Prior Authorization Denials,” AMA Press Release, February 25, 2025, available at, <https://www.ama-assn.org/press-center/ama-press-releases/physicians-concerned-ai-increases-prior-authorization-denials>. The press release references the AMA Augmented Intelligence survey (discussed above), from which the statistics from the press release are derived.

prior authorizations, the payors' follow-through has been "ineffectual." The press release further noted that only 16% of physicians who work with UnitedHealthcare, and only 16% who work with Cigna reported that prior authorizations had reduced in number. In terms of high numbers of prior authorizations, the worst offenders, according to the press release, were UnitedHealthcare, Humana, Anthem/Elevance, Aetna, Cigna, and Blue Cross/Blue Shield.

Meanwhile, insurers have argued that AI has accelerated the processing of prior authorization decisions from several days on average, to under a minute. This, however, has raised concerns among health care providers regarding just how integrated human decision-making has been in the prior authorization process.³²

The use of AI has not been limited to private payors, either; public payors are also using AI. Acentra Health, a cloud-based claims processing and care management service provider for 45 different state Medicaid programs, along with 25 federal agencies across 12 departments, offers several AI services.³³ These include AI-powered automated determination letters for payors. Since January, 2024, Acentra has estimated that its AI correspondence engine has drafted sections of 65,000 determination letters to providers and beneficiaries. Additionally, it has reduced the average time for a nurse to draft a complete document from 6 minutes and 35 seconds, down to 3 minutes and 28 seconds.

³² Mello, Michelle M., Artem A. Trotsyuk, Abdoul Jalil Dijberou Mahamadou, and Danton Char, "The AI Arms Race in Health Insurance Utilization Review: Promises of Efficiency and Risks of Supercharged Flaws," *Health Affairs*, 45 No. 1, p. 9, January, 2026, available at, <https://www.healthaffairs.org/doi/10.1377/hlthaff.2025.00897>.

³³ See, Harrison, Sean, Melissa Leigh, Daniel Hallenbeck, "Applying AI to Care Management and Claims Processing," *North Carolina Medical Journal*, vol. 85, no. 4, pp. 260-262, available at, <https://ncmedicaljournal.com/article/120569-applying-ai-to-care-management-and-claims-processing>.

III. Legal Risks & Considerations

Physicians and physician practices electing to employ artificial intelligence in their work face several potential risk areas. These risks primarily stem from the potential for the AI to malfunction (e.g., to “hallucinate”), but can also arise from technical legal compliance introduced by how AI functions. Hallucinations or other technical malfunctions can raise risks especially in the realm of malpractice, but also can expose practices to false claims liability. Depending on how the AI uses and secures protected health information (“PHI”), compliance with the Health Insurance Portability and Accountability Act of 1996 (“HIPAA”) and its regulations can also be a concern. And, of course, AI must be considered as a factor when dealing with payors regarding claims submission and denials, as well as prior authorizations.

A. Malpractice Risk

Depending on how AI is actually used, it can introduce an element of medical malpractice risk. Malpractice risk can take two forms, however: (1) treatment malpractice, and (2) informed consent malpractice. With treatment malpractice, the risk is fairly obvious: the physician relies too heavily on the AI, and the AI fails to perform effectively (e.g., providing incorrect clinical decision support services as a result of an hallucination), and the patient is harmed. With informed consent malpractice, the central issue is whether the patient consented to the use of AI in the course of a procedure. For example, as AI-assisted (or even AI-performed) surgery becomes more common, a physician who fails to inform a patient that AI tools will be used in the course of the procedure may expose themselves to malpractice risk if the patient did not consent to the use of AI, such as where use of the AI is still considered experimental. This risk does not manifest if AI is only being used in ways that do not impact the patient, such as in billing or coding. The consent malpractice risk likewise does not really affect documentation

(although consent regarding AI documentation may be an issue in a HIPAA context, discussed further below). However, if the doctor is using AI-assisted or AI-operated robotic arms, or other surgical devices in the midst of a procedure, that might alter the patient’s consent, depending on how involved the AI is in the actual service itself.³⁴

At its core, medical malpractice involves a breach of a duty of care that violates the standard of care. Assuming one can establish that a duty of care exists, the next question is whether the breach violated the standard of care. As AI becomes integrated into medical practice, however, this may change what constitutes the standard of care. How much reliance upon AI will be considered “standard”? At what point will over- or under-reliance deviate from that standard of care? As use of AI for clinical purposes becomes more common, these questions will settle themselves, but already there are signs that the public expects at least some use of AI by physicians. A recent study examined the attitudes of potential jurors when analyzing a hypothetical case involving a radiologist’s diagnosis (or failure to diagnose) a brain bleed.³⁵ The study found that, when AI found the presence of a brain bleed in the image, but the radiologist disagreed, the jurors found in favor of the plaintiff 56% of the time. Thus, the radiologist ignoring or countermanding the AI actually resulted in the jurors “punishing” the radiologist more often than not, and at the very least cut against the radiologist in the verdict.

³⁴ For more on the issues of informed consent malpractice and AI, see, Duffourc, Mindy N. and Ronan A. Cahill, “Legal Implications of AI Standard of Care Integration on Patients’ Informed Consent: Lessons from Surgery,” *NPJ Digital Medicine*, vol. 9, article no. 4, pp. 1-5, December 13, 2025, available at <https://www.nature.com/articles/s41746-025-02157-1>.

³⁵ Garris, Jonathan, “Use of AI Complicates Legal Liabilities for Radiologists, Study Finds,” blog post at Warren Alpert Medical School, Brown University, June 28, 2025, available at, <https://medical.brown.edu/news/2025-07-28/radiology-artificial-intelligence-malpractice-study>.

By contrast, doctors' overreliance upon AI could result in their skills degrading over time. A study of gastroenterologists in Poland found that, after three months of using AI-assisted imaging systems, physicians were 20% worse at spotting the presence of polyps and other abnormalities in colonoscopy imaging.³⁶ The underlying issue appears to be that the physicians gradually lose competence because they rely upon the AI too much; once the functionality of the AI was disabled, the physicians became less likely to detect abnormalities. Some dispute the study, however, arguing that it is unlikely that physicians would lose in months skills that it took them decades to obtain and hone, but one concern may be that the standard of care could shift to a lower standard, depending on how often physicians rely upon AI in their practice.

The Federation of State Medical Boards has waded into the matter, taking the position that use of AI must ultimately be balanced effectively.³⁷ To this end, the FSMB has stated, "Underappreciation of the ability of AI to improve health care delivery may restrict physicians from practicing to the top of their license and may result in a physician not taking full advantage of the tools that can improve patient outcomes. At the same time, over-reliance on AI can lead to real harms in independent clinical thinking and critical decision making such as misdiagnosis, medical errors, dependence, and skill degradation. This risk of harm increases in situations of algorithmic bias or where misinformation is present."³⁸

³⁶ Brumfiel, Geoff, "Research Suggests Doctors Might Quickly Become Dependent on AI," Shots – NPR Health News, August 19, 2025, available at, <https://www.npr.org/sections/shots-health-news/2025/08/19/nx-s1-5506292/doctors-ai-artificial-intelligence-dependent-colonoscopy>.

³⁷ See, "Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice," adopted by FSMB House of Delegates, April, 2024, available at, <https://www.fsmb.org/advocacy/news-releases/fsmb-releases-recommendations-on-the-responsible-and-ethical-incorporation-of-ai-into-clinical-practice/>.

³⁸ "Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice," adopted by FSMB House of Delegates, April, 2024, p. 5.

The FSMB has further taken the position that physicians are ultimately responsible for all aspects of care, even in the case of software failures. “Consistent with the prevailing standards for any tool used in the delivery of health care, the physician is ultimately responsible for the use of AI and should be held accountable for any harms that occur.”³⁹ While the document published by the FSMB acknowledges that the degree of accountability depends upon the nature of the AI use, the more involved in patient care the AI is, and the closer to actual practice of medicine the tasks it undertakes are, the more responsible the physician ultimately is. Clinical decision support services require the greatest level of scrutiny by physicians, while administrative tasks require the least because they are the least like practicing medicine. While the FSMB does not dictate policy for all medical licensure boards across the country, taking this position strongly suggests that the consensus among board regulators is that physicians still hold the ultimate responsibility to practice medicine responsibly, regardless of AI malfunctions.

B. False Claims Act Risks

The Federal False Claims Act⁴⁰ (“FCA”) can also be a source of risk for physicians and practices employing AI, with such risk mainly arising from AI usage in documentation, and in claims preparation and submission. For example, with AI-scribed visit notes or AI-penned records could be vulnerable to errors that, in addition to possibly creating malpractice risk, could generate FCA risk. For example, a physician might dictate a note, but the AI records it incorrectly, thus a “0.5mg” becomes “5mg”, or the AI could hallucinate what was actually said and transcribes something completely incoherent. If the physician or their practice does not

³⁹ “Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice,” adopted by FSMB House of Delegates, April, 2024, p. 5.

⁴⁰ 31 U.S.C.A. § 3729.

engage in careful oversight of AI-generated documentation, the record might no longer be effective for supporting the billed claim, or could affect a properly-functioning automated AI-driven billing system. In other words, the AI scribing aspect of the software could malfunction in a way that would lead to a properly-working AI billing aspect to bill for a more expensive service, effectively “upcoding” based on incorrect records.

Similarly, AI in a CDSS role could provide information to a physician about a course of care for a patient. If the physician follows the recommendation, but the CDSS tool has not been sufficiently updated to take into account new federal billing guidance, the recommendation – while not necessarily deviating from the standard of care – could result in an improperly billed claim that is, again, not supported effectively by clinical documentation. Alternatively, while the claim itself might be reimbursable, the medical records might only support a service that is reimbursed at a lower rate. Worse, the AI could hallucinate in its guidance even with the correct programming, and recommend care that deviates from coverage requirements, even though the AI has the most current coverage information programmed into it. Additionally, if the AI-powered CDSS tool offers incorrect recommendations and the physician follows the recommendations, FCA liability could attach if the patient is harmed and the services are deemed “worthless services.”

There are similar risks with respect to AI usage in billing and coding, especially when those processes become automated. For example, an AI-driven billing system could recommend, or automatically fill in a claims form with, an incorrect CPT code. The incorrect code could be based on hallucinations regarding the requirements to bill for the service that are programmed into the AI. Alternatively, if the AI bases its recommendations on documentation, it could

hallucinate in its interpretation of the documentation and apply an incorrect, higher-reimbursed code based on its misreading of the record.

Thus far, there has been no FCA enforcement based on these scenarios, but the risk of such incidents increases the more that AI is integrated into medical practices. Moreover, FCA liability always rests with the health care provider; it cannot be shifted to a third party, even if the third party's actions gave rise to the FCA liability in the first place. At best, one can rely upon indemnification clauses, or attempt to bring suit against the third party (in this case, software vendors). For this to change, the FCA itself would need to be modified, or a separate law would have to be passed that would create direct statutory liability for software vendors in FCA violations arising from malfunctioning software. While this may seem unlikely, to some extent there exists precedent for such an outcome. Prior to 2009, business associates under HIPAA only faced contractual liability for their own HIPAA violations; only the covered entity faced liability, and would have to sue the business associate under breach of contract, or hope that indemnification language would at least make the covered entity whole. However, with the passage of the HITECH Act, business associates now faced direct, statutory liability.⁴¹ Regardless, short of Congressional intervention and modification of the laws, liability will remain with health care providers for the foreseeable future.

C. HIPAA Risks

At its core, AI relies upon vast quantities of aggregated data to function effectively. The requirements of HIPAA, however, place particular and complicated burdens on how, what kind,

⁴¹ The specific provisions that imposed liability are P.L. 111-5 §§ 13401 and 13404. For more on business associate direct liability, see "Direct Liability of Business Associates," OCR information page, available at, <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/business-associates/factsheet/index.html>.

and how much data is used by AI, and for what purposes. “Protected health information” is defined as individually identifiable health information that is transmitted or maintained in electronic media, or in any other form or medium.⁴² Information that is de-identified is not considered PHI.⁴³ Depending on what kind of information software developers use when employing AI, and how that information is used, there may be different HIPAA obligations for both business associates and physician and physician practice covered entities.

For example, if AI developers only use de-identified information to train AI, then HIPAA ceases to be a concern with respect to how the AI trains. If, however, the AI works off of the raw data contained in patient and billing records, without regard to whether that information is PHI, its use could create risk under HIPAA. Data aggregation performed by business associates on behalf of covered entities is permitted under HIPAA without a patient authorization.⁴⁴ The definition of “data aggregation” also permits this to occur on behalf of other covered entities with which the business associate contracts. Put another way, it is permissible for the business associate to aggregate data from Covered Entity A and Covered Entity B, so as to assist both covered entities with their respective health care operations. Note, however, that “health care operations” is not a blanket permission to use any/all PHI aggregated by a business associate for any purpose. The definition of “health care operations” is limited, and pertains to certain quality assessment and improvement activities, evaluating practitioner performance, providing training

⁴² 45 C.F.R. § 160.103.

⁴³ For more on de-identification of PHI, see 45 C.F.R. § 164.514.

⁴⁴ 45 C.F.R. § 164.504(e)(2)(i)(B). The term “data aggregation” is further defined at 45 C.F.R. § 164.501.

programs, arranging for medical and legal review services as well as auditing functions and fraud and abuse detection and compliance programs, and certain other specific activities.⁴⁵

Absent from this definition, however, are the two other main reasons why PHI is used by covered entities and business associates: for treatment purposes, or for payment purposes. Both treatment and payment are distinct uses, separate from the definition of “health care operations,” each having their own specific definitions.⁴⁶ Many of the activities for which AI is deployed fall into “treatment” such as with CDSS, and “payment” such as with AI-powered billing and claims activities. These uses may raise separate issues under HIPAA, when dealing with data aggregated across multiple covered entities.

Consider how Athenahealth advertises its AI products, which claims that it can improve payment recovery by “analyzing data from over 160,000 providers on the [Athenahealth] network...[to identify] potential claim issues in real time.”⁴⁷ If Athenahealth, functioning as the business associate of Provider A collects data from Provider B, which it integrates into its AI claims software to allow Provider A to better prepare and submit its own claims, has Athenahealth improperly used and/or disclosed PHI? Has Provider B improperly disclosed PHI to Athenahealth when it knows that Athenahealth is using that information to help other covered entities? Some of these issues can be mitigated through the use of health information exchanges, which allow covered entities to share PHI with each other when treating the same patients. But that does not necessarily address circumstances where PHI is being shared and/or used between

⁴⁵ The definition of “health care operations” can be found at 45 C.F.R. § 164.501.

⁴⁶ “Payment” and “treatment” are each defined at 45 C.F.R. § 164.501.

⁴⁷ “Top Ways AI-Native athenaOne EHR lifts RCM Capabilities,” athenahealth.com blog entry, March 11, 2025, available at, <https://www.athenahealth.com/resources/blog/reduce-administrative-burden-with-athenaone-ai-tools>.

covered entities where no such common treatment relationship exists, and where the use does not fit into the definition of “health care operations.” While data aggregation services are permissible for the performance of “health care operations” across multiple covered entities, the definition does not account for payment or treatment activities across multiple covered entities. Thus, without obtaining an authorization from patients to use and disclose their PHI for these purposes, such uses and disclosures may violate HIPAA.

To date, there has been no enforcement of HIPAA for this type of use, but such uses are relatively new, and the Department of Health and Human Services Office for Civil Rights (“OCR”) may have other priorities at the moment. Additionally, a business associate or covered entity might be able to argue that the data aggregated for use in AI-powered CDSS tools or billing and claims services is de-identified, which would sidestep HIPAA, but the devil will lie in the details of how the AI in such circumstances actually functions, along with how the OCR perceives that functionality.

The OCR has not provided any real guidance with respect to how AI interacts with HIPAA, with the sole exception being a “Dear Colleagues” guidance letter published on January 10, 2025, by the Biden-administration-era OCR, titled “Ensuring Nondiscrimination Through the Use of Artificial Intelligence and Other Emerging Technologies.” This document appears to have been removed from the OCR’s website and is only accessible through an archival link.⁴⁸ Guidance regarding apps and cloud computing might offer some insight as to the OCR’s perspective regarding software⁴⁹, but such guidance generally reiterates the existing obligations

⁴⁸ Available at, <https://www.hhs.gov/sites/default/files/ocr-dcl-section-1557-artificial-intelligence.pdf>.

⁴⁹ See, “The Access Right, Health Apps, and APIs,” available at, <https://www.hhs.gov/hipaa/for-professionals/privacy/guidance/access-right-health-apps-apis/index.html>.

of business associates and contextualizes them in the realm of mobile apps, with the central focus being upon whose behalf uses and disclosures of PHI by the app are being taken. For example, in response to a question as to whether a covered entity is liable for releasing information to an app at an individual's request, the guidance explains that the covered entity could be liable if the app receives, transmits, or maintains electronic PHI on behalf of the covered entity. In such circumstances, rules regarding further improper disclosure would apply. This is helpful in recognizing that the general rules of HIPAA continue to apply even as the context of their use shifts to platforms like apps, but it does not address the peculiar and specific issues presented by AI in the HIPAA context. Time will tell whether and how the OCR regulates these matters, but in the meantime, the risk of HIPAA violations remains for improper use and disclosure of PHI for payment and treatment purposes, even while other activities may be permitted by data aggregation services provided for the purpose of health care operations.

D. Reimbursement and Claims Denial Risks

Separate from federal statutory and regulatory liability risks, health care providers remain locked in the aforementioned “arms race” with payors surrounding AI-powered documentation, and claims preparation and submission on the one hand, and automated claims denial and prior authorizations by payors on the other. The “risk” here is less about legal exposure, and more about the risk from payor reliance upon AI-powered prior authorizations, and an increase in denials of prior authorizations due to automated processes.

In the face of this, some patients have attempted to push back through lawsuits, although such cases are thus far relatively uncommon. One case, however, has proceeded and survived a

motion to dismiss. In *Estate of Barrows v. Humana, Inc.*,⁵⁰ the Western District for the District of Kentucky addressed circumstances in which the estate of a deceased Medicare Advantage plan member sued the plan for denials of post-acute care. The plaintiffs in the case alleged that Humana had used AI-powered software to deny coverage of post-acute care services. Interestingly, the suit did not argue that the denials themselves were improper. Instead, the plaintiffs sidestepped the issue of the propriety of the denials, and instead focused on a breach of contract theory, breach of the implied covenant of good faith and fair dealing, unjust enrichment, and various claims grounded in state law. The plaintiffs claimed damages were the insurance premiums paid to Humana for services that were allegedly not delivered. Humana, by contrast, filed a motion to dismiss, and argued that the plaintiffs had failed to use the mandatory federal administrative process that is required prior to bringing such lawsuits. Humana also argued that the case was preempted under provisions of the Medicare Act.

With respect to the administrative exhaustion requirement, Humana argued that the Medicare Act⁵¹ required the plaintiffs to exhaust administrative remedies before pursuing judicial remedies. The plaintiffs, however, argued that their case was not brought under the Medicare Act and was not inextricably entwined with it. The court first analyzed the issue of whether the plaintiffs' claims for breach of contract even arose under the Medicare Act. The court explained that claims are only brought under the Medicare Act when the standing and substantive basis for the presentation of the claims is the Medicare Act itself, or where the claim is inextricably intertwined with a claim for Medicare benefits. In this case, however, the court found that most

⁵⁰ 2025 WL 2375645 (W.D. Ky., 2025).

⁵¹ Specifically, 42 U.S.C.A. §§ 405(g) and (h).

of the plaintiffs' claims did not fit into either category. The plaintiffs were claiming breach of contract not because of a failure to provide benefits, but rather because of Humana's use of an AI product to make coverage determinations and a failure to disclose that process to the plaintiffs. Thus, the plaintiffs argued that they had not received the actual services paid for (i.e., claims review by humans), and therefore Humana breached its contract with the plaintiffs.

With respect to the administrative exhaustion requirement, the court noted that there were several exceptions to this, including where the claimants would be irreparably harmed if the exhaustion requirement was applied against them, and where exhaustion of such remedies would be futile. The plaintiffs had already suffered irreparable harm, the court determined, by having experienced medical setbacks while engaging in the administrative remedy process, where such setbacks could not be remedied by a later payment for services. In specific, that plaintiffs had foregone medically necessary care, which in turn resulted in them being admitted to the hospital or requiring additional treatment. Retroactive payment, therefore, would not remedy the harm the plaintiffs had already suffered.

With respect to the issue of futility, taking the plaintiffs' assertions about the administrative process as accurate, the court noted that the plaintiffs' complaint described that, when a denial was reversed on appeal, Humana simply issued a new denial, thereby forcing the beneficiary to restart the appeals process. In one instance, a beneficiary received seven denials for the same care within 30 days, each after successful appeals and each lacking any factual basis for the denial. When the beneficiary in this case was later hospitalized, the plaintiffs noted that Humana issued five more denials for post-acute care, each also after successful appeals and again without factual bases. If a claim made it to an administrative law judge, Humana would only then pay the claim to terminate the appeals process, and avoid (according to the plaintiffs)

judicial review. Taking these allegations as true, the court found that the plaintiffs had demonstrated futility in the administrative appeals process. The court ultimately dismissed several state-law claims, but permitted the plaintiffs' remaining claims to proceed beyond a motion to dismiss. As of this writing, the case had yet to reach a final disposition.

Bear in mind that the underlying problem identified by the plaintiffs was that all of these denials were occurring because of Humana's use of automated AI-driven prior authorization systems. Moreover, the plaintiffs were able to proceed with their suit and survive a motion to dismiss primarily because their allegations were sufficient to meet exceptions to otherwise disqualifying aspects of the case (such as the requirement to exhaust administrative remedies before bringing suit), and because the claims were couched primarily in breach of contract surrounding the use of the AI, rather than an improper denials of benefits. While the denial of benefits did harm the patients, the damages sought were not consequential damages, but rather a return of premiums paid.

Other cases against payors, however, have been less successful. For example, in *Carlin v. United HealthCare Ins. Co. of New York, Inc., et al.*,⁵² two 80-year-old Medicare beneficiaries, who also had a supplemental insurance plan, sued Novitas Solutions and National Government Services, Inc., in their capacities as Medicare Administrative Contractors ("MACs"), as well as United HealthCare in its role as a private insurer. The plaintiffs filed several Medicare claims personally for a range of services, some of which were unpaid, and some of which were only partially paid. Several claims did not receive an initial determination due to insufficient information in the claim itself (e.g., no rendering practitioner named) or because the claim was

⁵² 2025 WL 2549106 (S.D.N.Y. 2025).

inappropriate (e.g., a claim for services under Medicare Part-A, which only health care providers may submit). In none of the claims did the plaintiffs attempt to file appeals. The MACs moved to dismiss for lack of subject matter jurisdiction, again relying on the fact that the plaintiffs had not exhausted (indeed, had not even begun) administrative remedies. The court, unlike in the *Estate of Barrows* case discussed above, noted that, at base, the plaintiffs' claims were a demand for benefits, and thus subjected them to analysis under the Medicare Act. Because the plaintiffs had not even had initial determinations in some instances, let alone appealed any denied claims, the court agreed with the MACs, and accordingly granted the MACs' motion to dismiss.

The plaintiffs' claims against United HealthCare, however, were more complicated, and turned on civil RICO claims. These included, among other things, that United HealthCare had implemented practices and procedures designed to ensure that most claims would be denied, including the use of AI algorithms that were programmed to refuse, delay, or deny most claims. However, the court found that this claim did not rise to the level of fraud or deception. Specifically, the court noted "Practices or procedures that favor denial of policyholders' claims might be self-interested, harmful, or unfair, but it is well-established that 'conduct is not deceptive merely because it is wrongful or because it harms another.'"⁵³ Because the plaintiffs' complaint did not state the necessary elements of a civil RICO claim, including that the claims surrounding AI usage did not amount to fraud, the court granted United HealthCare's motion to dismiss.

⁵³ Carlin, at *7.

IV. Risk Mitigation Strategies in the Face of AI

Health care providers must contend with a variety of risks when adopting AI.

Nevertheless, they do have options to mitigate those risks; several strategies can help, especially when taken in conjunction. Efforts can also be made to try to shift these risks so that they are shared more with the AI and software industry. Changes in contractual clauses, development of internal compliance plans and auditing procedures, and the political process itself may prove useful as methods for risk mitigation.

A. Contractual Clauses and Legal Challenges

Electronic health records software is a fact of life in the modern era of medicine. As such, most physician practices either have signed, or will sign, a license agreement for EHR software.⁵⁴ It is always advisable for physician practices to have legal counsel review such license agreements, but these documents should carefully consider how the software will be used, and how AI will be integrated into the software. Several provisions will be relevant to this effect.

One key provision that protects software developers and vendors from liability is the disclaimer of liability. This type of clause appears in nearly every software license agreement, and typically disclaims a range of otherwise implied warranties and other liability. The clause usually limits damages available under the contract to only the fees paid under the contract. As

⁵⁴ For more on EHR end user 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.

an example, consider the following contract language, taken directly from a software vendor's license agreement provided to a small physician practice:

“Except as expressly provided herein, [Vendor] disclaims all representations and warranties of any kind or nature, express or implied (either in fact or by operation of law), with respect to any service or item provided hereunder. Except as expressly provided herein, [Vendor Software] is provided ‘as is.’ [Vendor] does not warrant that [Vendor Software] will be error-free or will be provided (or available) without interruption or meet Client’s business or operational needs.”

A further disclaimer in the same agreement included limitations on damages up to only 1 year's worth of service fees, and explicitly disclaimed liability for any consequential damages, lost profits or business opportunities, loss of data, or costs associated with procuring replacement items or services.

Such provisions have proven effective at court. Consider the case of *Wired Informatics, LLC v. OmniMD, Inc.*⁵⁵ The case involved a dispute between two software vendors. OmniMD, a developer of EHR software, sued Wired Informatics for failures of Wired's software, Invenio: an AI software designed to read and parse language in medical transcripts, among other functions. The software would transfer information from medical transcripts to specific locations in standardized electronic health records for patients of OmniMD's customers. The software, however, did not function effectively, and Wired did not provide sufficient support, according to OmniMD.

OmniMD brought several claims against Wired, including fraud, and breaches of implied warranties of fitness and merchantability. With respect to the fraud claim, the court granted Wired's motion to dismiss, finding that OmniMD's mere statements that Wired had represented

⁵⁵ 2019 WL 3429486 (D.Mass., 2019).

the software was “ready to use” when it was not, that the software “was barely working for only a small portion of functions” it was supposed to perform, and that the nine-month timeframe to fix these problems would be sufficient when it was not, were not allegations made with sufficient particularity to meet the requirements under Federal Rule of Civil Procedure 9(b). Instead, the court found “These allegations do not describe with particularity the alleged misleading representations by Wired and do not identify the ‘who, what, where, and when’ of the misrepresentations.”⁵⁶ The court did allow the pleading to be amended, however.

With respect to the breaches of implied warranties of fitness and merchantability, because the license agreement between OmniMD and Wired included disclaimers and limitations, which explicitly stated that the software would not perform error-free, and further disclaimed all warranties of merchantability or fitness for a particular purpose, the court dismissed this count as well. The court was quick to note that the section of the agreement⁵⁷ was explicit and clear, and was even written in all capital letters and bolded, clearly setting it apart from other portions of the agreement, thereby supporting a conclusion that the disclaimer was conspicuous enough that a reasonable party would have noticed the language. If, the court noted, the contract had been for services only, rather than goods, and the Uniform Commercial Code had not applied, then OmniMD would not have had a cause of action for the breach of such warranties. However, because the language of the contract was clear and OmniMD had waived such warranties, it could not then sue for breaches of them. Accordingly, the court dismissed these claims without leave to amend the complaint.

⁵⁶ Wired Informatics, Inc., at *3.

⁵⁷ The full text of the section is quoted in the opinion, at *4, and is worth reviewing as another example of a typical disclaimer clause in software license agreements.

Presented with such disclaimers, physician practices can try to argue that either the disclaimers should be removed, or that the vendor or developer should stand good for their work by including indemnification language for certain specific malfunctions of the software (e.g., surrounding AI hallucinations, instead of all software malfunctions across the board). Such indemnification language could look something like the following:

“Indemnification. [Vendor] shall indemnify and hold the other harmless [Practice], including its respective directors, officers, employees and agents, from and against any and all costs, losses, damages, expenses and liabilities (including without limitation reasonable attorneys’ fees) caused by or incurred as a result of any hallucination or other malfunction of any artificial intelligence functions of [Vendor’s Software].”

The language presented here is relatively simple, and only meant to provide a general sense of what an indemnification clause might look like in a software license agreement. Such language would need to be carefully tailored to address the specific uses and functionality of the AI in question.

It is likely, however, that vendors and developers will reject such changes for several reasons. First, because such vendors and developers usually are providing standardized contracts to a broad range and high number of health care providers, they are unlikely to be comfortable modifying a specific contract just to suit a single physician practice. Allowing for such changes could increase administrative burdens for the vendor or developer. They also would likely claim that, were such a change to be made, they would have to make it for all health care providers with which they contract, since doing otherwise would require tracking unique provisions across any health care provider with a “customized” contract, and either approach would be infeasible.

Second, doing so would expose them to legal risk and liability that the vendors and developers do not wish to assume, even on the limited basis of AI hallucinations. Given the potential harm that AI malfunctions could cause, and the range of penalties or other damages that

could result from such malfunctions, vendors and developers are unlikely to indemnify physician practices without significant pressure to do so coming from across the health care sector.

Nevertheless, if more physician practices and other health care providers demand that software developers stand good for their AI-powered products by indemnifying their customers for software malfunctions like hallucinations, the software industry may begin to provide them.

Physician practices should also carefully consider the functionality of AI as it is integrated into their software, and how that integration will use the data they input into the software. As addressed above in the discussion on HIPAA risks, PHI entered into an AI that is used for purposes other than “health care operations” by parties other than the covered entity legally cannot be the subject of data aggregation services. In other words, if a vendor purports to use data entered by one physician practice to assist with claims preparation and submission (i.e., “payment” under HIPAA, rather than “health care operations”), the license agreement should address this at a minimum in sections pertaining to data usage.

Additionally, the physician practice should consider whether it will obtain authorizations from all of its patients regarding such usage, and should seriously consider doing so (or updating existing authorizations to include such uses). Alternatively, the practice should obtain written assurances from the software developer or vendor that the software’s AI will only use de-identified aggregated data, and not PHI for any uses that do not fit within the definition of “health care operations”; such assurances should be added to the license agreement itself. Similar considerations will apply for, for example, auto-scribing AI, if the improvements provided by the recordings of interactions with patients are used for purposes other than “health

care operations,” the practice will either need to obtain individual patient authorizations⁵⁸ for such use, or will need written assurances from the vendor or developer that such data will be de-identified. However, similar to demands for indemnification language, physician practices should be prepared for software developers and vendors to refuse to include in the license agreement any guarantees that their AI does not use identified data and only uses de-identified aggregated data.

B. Compliance Plans and Internal Audits

As a general matter, compliance plans form the foundation of any health care provider’s efforts to reign in fraud and abuse, and maintain compliance with the myriad regulatory schema that providers face. When health care providers introduce AI tools into their practices, such compliance plans must be updated to take into account their impact. The HIPAA Security Rule⁵⁹ requires that covered entities revisit their security risk analysis, and update their policies and procedures to reflect the updated risk analysis, when introducing new technology.⁶⁰ The OCR has further issued guidance that states that the risk analysis should be “ongoing” although it does not establish any requirements regarding the frequency of such analyses. Instead, the guidance states,

“The frequency of performance will vary among covered entities...A truly integrated risk analysis and management process is performed as new technologies and business operations are planned, thus reducing the effort required to address risks identified after implementation. For example, if the covered entity...is planning to incorporate new

⁵⁸ As a separate matter, in states with anti-wiretap laws requiring the consent of both parties to being recorded, physician practices will already need to obtain consent for the recording. For example, Pennsylvania’s own anti-wiretap law prohibits recording of parties, unless consent has been given prior to the recording. 18 Pa. C.S.A. §§ 5703 for the general prohibition, and 5704(4) for the exception where consent is given.

⁵⁹ 45 C.F.R. §§ 164.302-164.318.

⁶⁰ 45 C.F.R. § 164.308(a)(8).

technology to make operations more efficient, the potential risk should be analyzed to ensure the e-PHI is reasonably and appropriately protected. If it is determined that existing security measures are not sufficient to protect against the risks associated with...the introduction of new technology, then the entity must determine if additional security measures are needed.”⁶¹

Once AI is introduced into a covered entity’s electronic infrastructure and security environment, that alone should prompt a HIPAA risk analysis of the new AI functionality with respect to electronic PHI, which in turn will necessitate the revision of policies and procedures designed to address and mitigate security risks. Depending on how pervasive the AI will be within the electronic infrastructure, the breadth and scope of the analysis may vary.

Separate from the requirements under HIPAA, however, it will be wise to address the introduction of AI in other compliance programs, especially those pertaining to fraud and abuse, to assess and address the risks posed by the new technology. This may change how often internal audits of billing, coding, and related records occur, as well as the tools with which such audits are performed. For example, physician practices may consider using AI in the performance of internal billing audits, although this may introduce additional elements of risk if the software fails to detect problems that a human might find. Certainly, any use of AI in the conduct of internal audits must itself be supported by human decision-making and judgment.

Even if AI is not used, however, humans will need to be vigilant with respect to examining how AI coding and billing, as well as documentation software is functioning, and may need to periodically adjust how the AI functions to ensure greater accuracy and consistency. On an individual level, physicians using AI scribing software will need to review their own AI-generated documentation to be certain it is accurate, and must use their medical judgment in the

⁶¹ “Guidance on Risk Analysis,” Office for Civil Rights, Department of Health and Human Services, available at, <https://www.hhs.gov/hipaa/for-professionals/security/guidance/guidance-risk-analysis/index.html>.

course of such reviews. Compliance officers should also recognize that the act of reviewing AI-recorded documentation is fundamentally different from entering the information oneself.

Whereas a physician might very carefully enter documentation into a chart or progress note themselves, when reviewing an AI-generated note, there may be a temptation to “skim” the documentation rather than “flyspeck” it to ensure that the details are accurately captured.

Periodic billing and coding audits should be performed to guard against FCA liability, and to confirm that appropriate codes are being billed by the AI (if it is being used in such a manner). From a “defensive medicine” perspective, such audits may also serve two additional functions. First, they may ensure that AI-generated documentation can support the claims submitted for such services. Second, they may help protect against potential malpractice lawsuits, by guaranteeing that the documentation reflects that the appropriate standard of care was followed.

All of these audits should further include records demonstrating where, when, and why a clinician followed or disagreed with any AI diagnostic or CDSS tools, as well as any billing and coding tools. This helps to create a digital record that demonstrates the physician practice’s ongoing compliance efforts and supervision of the AI. This may help insulate from potential fraud claims, catch overpayments before they create FCA liability, and address HIPAA privacy and security incidents. Of course, all of this requires that physician practices actively engage with their compliance programs, and do not simply allow them to collect dust on a shelf. Unused compliance programs and policies not only do not help protect against compliance problems, they can also serve as proof of a physician practice’s negligence by demonstrating that the practice knew better and chose to ignore its own policies and procedures.

C. The Political Process and State Law/Regulation

As a risk mitigation strategy, entering into the political process may seem like a course of last resort. Unfortunately, it may simply be what is necessary to come to grips with the potential harms that AI may cause both to patients and to their clinicians, and the massive imbalance in power between such clinicians and the companies that create and distribute AI-powered software. To date, Congress has done relatively little to regulate AI. Moreover, as addressed above, lawsuits often provide weak or uncertain solutions to malfunctioning AI, or third party payors who use AI to deny claims. Of course, there is also always the incentive for a payor to settle a losing case before it goes to a verdict and establishes legal precedent with which they must then comply (or face further lawsuits and expense).

Health care providers find themselves in the worst possible position with respect to AI. They bear all of the risks posed by the AI malfunctioning and injuring their patients, all of the burdens imposed by the AI malfunctioning and overbilling or creating garbled documentation that cannot support claims for services, and increasing burdens placed upon them by payors which are themselves utilizing AI to assist in claims review and denials in the midst of a so-called “arms race.” Health care providers by themselves lack the power to force software companies to adjust their contracts, and thus may need to find ways to level the playing field. This may require coordinated efforts through specialty societies and political advocacy.

Some states have taken initial efforts to address the use of AI through the law, although these efforts are scattered and, in some cases, snarled by disputes between political parties in states with divided governments. For example, Pennsylvania’s House Democratic Caucus

introduced a bill to modify the Commonwealth's Health & Safety and Insurance laws to impose a range of requirements on both health care practitioners and insurance providers.⁶²

If passed, the law would impose requirements on health care providers that include notifying patients about using AI for clinical decision making, and requirements to conduct internal audits to ensure the accuracy of AI used by the provider. The law would also regulate developers and vendors of AI software, although details are somewhat sparse, and the Department of Health is directed to issue regulations to implement the statute; this would provide greater detail. The law would not apply to AI usage relating to administrative, scheduling, or scribing activities.

As applied to insurers, the law would exempt most Federal payors, as well as Medicare supplemental policies, and TRICARE policies, workers compensation, and automotive insurance. On other insurers, the law would impose duties to disclose the use of AI in utilization reviews to both network providers and covered individuals, and would impose "responsible use" requirements where the AI would be required to make decisions based on the medical and clinical history of a covered person, other non-clinical information as provided by their health care providers, and other relevant and nonclinical information contained in the medical record. Moreover, the AI's algorithms must not be based solely on a group data set, nor supersede the decision making of a health care provider conducting utilization review, and any data sets used would be prohibited from directly or indirectly discriminating against covered persons. The law in its current form is likely to face gridlock in the Republican-controlled Pennsylvania State

⁶² See, House Bill No. 1925, available at, <https://www.palegis.us/legislation/bills/2025/hb1925>.

Senate, but does at least stand as evidence that lawmakers are contemplating the need to regulate AI specifically within the health care arena.

Other states have also passed various AI-related legislation. California recently passed a sweeping AI bill that regulates a broad scope of activities in different economic sectors.⁶³ Illinois has passed a law restricting the use of AI in behavioral health.⁶⁴ Maryland has passed a law regulating the use of AI in health insurance.⁶⁵ In one cautionary tale, Colorado passed a consumer protection oriented law regulating AI in 2024, representing the first state to engage in such regulation.⁶⁶ The law, however, was subsequently delayed in its implementation until June, 2026,⁶⁷ having proven difficult to implement due to disagreements among various stakeholders on how exactly to put the law into practice. Perhaps unsurprisingly, allegations are that the AI industry pushed back against the bill due to concerns with liability. Given that violations of the law, which requires notifying regulators and consumers when AI is used for consequential decisions (e.g., employment, loans, housing, etc.), and penalties are up to \$20,000 per violation, industry concern is understandable. Still, the experience in Colorado demonstrates the potential difficulties as states move forward with AI regulation, while Congress seems unable to make any

⁶³ See, Cal. Senate Bill 53, available at, <https://legiscan.com/CA/text/SB53/id/3270002>.

⁶⁴ See, Ill. House Bill 1806, available at, <https://ilga.gov/Legislation/BillStatus/FullText?GAID=18&DocNum=1806&DocTypeID=HB&LegId=159219&SessionID=114>.

⁶⁵ See, Maryland House Bill 820, available at, <https://legiscan.com/MD/text/HB820/id/3251347>.

⁶⁶ See, Colorado SB24-205, available at, <https://leg.colorado.gov/bills/sb24-205>.

⁶⁷ See, Paul, Jesse and Taylor Dolven, “Colorado Lawmakers Abandon Special Session Effort to Tweak AI Law, Will Push Back Start Date to June, 2026,” *Colorado Sun*, August 25, 2025, available at, <https://coloradosun.com/2025/08/25/colorado-ai-law-tweak-dies/>.

headway on the matter. For now, though, this approach may be the best way for physicians to influence their own stake in the complicated realm of AI.

V. Conclusion

At present, there is no reason to expect a sudden collapse of AI and its influence. Far too much has been invested in the technology at this point for it to simply disappear overnight. With this in mind, it is incumbent upon physicians and physician practices, as well as their legal counsel, to remain informed about developments in AI function and regulation alike, and to think carefully about legal and regulatory obligations and how they will interact with AI functionality as it becomes more integrated into the practice of medicine and health care more broadly.

Artificial intelligence will continue to be used by insurance companies in both prior authorizations and claims review processes, and likely will be used in auditing practices as well. Electronic health records software platforms will see further integration of AI tools designed to automate various processes to enhance efficiency. These will, in some respects, make life easier for physicians and physician practices, but will also introduce liability risk from various sources. To the extent possible, physicians and other clinicians should attempt to minimize some of those risks by demanding license agreement changes, as well as by banding together through specialty societies that are better positioned to influence the political process. As discussed above, this may be the best, indeed the only way to rein in the influence of third party payors' uses of AI, and of AI software developers and vendors themselves, and bring the individual clinician back into some semblance of parity at the negotiating table.

For the time being, Congress does not seem likely to step in with sweeping regulation to govern the AI industry. Instead, it appears more likely that this will occur at the state level,

which in turn will create a patchwork of varying legal requirements for users and creators of AI.

This will necessitate health care counsel closely monitoring developments in the states where their clients work, to ensure continued compliance. At the same time, those legal frameworks will doubtless face pushback from the AI industry itself, as has occurred within Colorado.

Regardless, the current environment requires careful management of compliance obligations, and consideration of how regulatory frameworks are being applied to AI.