Artificial Intelligence in Healthcare: Balancing Risks and Rewards

July 31, 2023

Introduction

Artificial intelligence ("AI") has the potential to radically transform the healthcare industry. While AI boasts a wide range of promising use cases, from practice management to clinical decision-making support, the risks associated with adopting AI are complex and ever-changing, encompassing, among other things, cybersecurity and privacy, data quality, professional ethics, fairness, fraud and intellectual property ("IP"). Many regulatory agencies may also assert jurisdiction, including the U.S. Food and Drug Administration ("FDA"), the Federal Trade Commission ("FTC"), the U.S. Department of Health and Human Services ("HHS") Office of Civil Rights ("OCR"), the U.S. Department of Justice ("DOJ"), state agencies and the European Union ("EU") and other foreign regulatory bodies. Naturally, this complicated amalgamation of regulatory authorities implicates a complex set of international, federal and state laws, some of which are listed in a chart below.

Healthcare companies must carefully assess potential risks to ensure practitioners and patients can fully realize the potential benefits of AI in the healthcare space; developers, in turn, must promote both innovation and safety to accelerate the effective adoption of these powerful technologies. In this Debevoise In Depth, we examine the legal and regulatory risks posed by AI adoption in healthcare and identify ways stakeholders can mitigate these risks to achieve their goals and better position themselves to respond nimbly to disruptive innovation.

AI-Specific Regulations and Guidance: Global and U.S. Developments

In light of the evolving regulatory approaches to AI, healthcare companies must carefully monitor piecemeal developments at multiple agencies to fully understand applicable opportunities and risks. The AI regulatory landscape falls into many high-level categories, including general AI regulations, sector-specific AI regulations and

nonbinding regulatory guidance. Below, we address each category and the resulting implications on the healthcare space.

General Global AI Regulations

The most notable of the general AI regulations currently under consideration is Europe's draft legislation of the Artificial Intelligence Act (the "EU AI Act"). The EU AI Act was introduced in a proposal by the European Commission in April 2021,[1] and has been followed by negotiating positions adopted by the Council of the European Union[2] and, most recently, by the European Parliament.[3] These three bodies will now negotiate their positions, with the goal of adopting a finalized form of the EU AI Act by the end of 2023. If it comes into effect, the EU AI Act will broadly govern the use of AI in the EU.

If enacted, the EU AI Act would place potentially onerous compliance obligations on a wide spectrum of companies using AI systems based on the potential risk posed by an AI system's intended use: AI systems that pose an "unacceptable risk" would be banned outright; those classified as "high risk" would be subject to stringent regulatory and disclosure requirements; and certain systems, including specific generative AI systems, would be subject to heightened transparency obligations.

For instance, AI systems used in medical devices by healthcare providers, as well as those used for remote biometric identification of natural persons and those used to gain "access to and enjoyment of essential private services and public services and benefits," will be classified as "high risk," subject to certain pre- and post-market assessments. Healthcare companies with products that incorporate AI, and healthcare providers whose uses of AI systems fall into these use cases, would thus be subject to stringent limitations and/or heightened requirements under the EU AI Act. As the EU AI Act undergoes extensive negotiations, the scope of requirements for healthcare companies may change. As such, healthcare companies should remain up-to-date with respect to relevant developments.

In addition to the EU, Brazil, China and Canada have all published general AI regulations that apply to all sectors of the economy.[4] Each of these regulations is also in draft form and is thus subject to amendment.

---

1 The proposal is *available* here.
2 The Council of the European Union's negotiating positions are *available* here.
3 The European Parliament's negotiating positions are *available* here.
4 Debevoise Data Blog: Overview of Global AI Regulatory Developments and Some Tips to Reduce Risk (May 3, 2023), *available* here.

U.S. Sector-Specific AI Regulations

The regulation of AI in the United States is largely carried out by various state and local agencies with specific authority over banking, insurance, securities markets, criminal justice, employment, etc. While there are currently no AI-specific regulations in the healthcare sector, healthcare companies and providers should remain aware of developments in other sectors of the economy, which may serve as templates for future legislation, both in terms of requirements and level of prescriptiveness.

At the state level, the Colorado Division of Insurance's ("CO DOI") risk-based draft Algorithm and Predictive Model Governance Regulation imposes significant compliance and operational obligations on regulated insurance entities, such as requiring insurance organizations to identify governance principles for AI, create oversight by senior management and the Board, formulate a cross-functional AI governance committee and develop a risk assessment rubric to assess and prioritize risks.[5] These obligations have evolved since the first draft of the regulation, and are likely to continue to evolve as regulators evaluate their approaches to governing AI. While this regulation does not directly apply to healthcare companies and providers, such governance principles may be applied to healthcare entities under a regulatory regime that mimics the CO DOI draft regulation. Accordingly, healthcare entities should continue to track these developments, especially as obligations continue to change.

At the local level, New York City's recently adopted Automated Employment Decision Tool Law ("NYC AEDT") requires, *inter alia*, covered employers to perform annual independent bias audits and to post public summaries of those results.[6] This regulation may apply to healthcare entities that meet the jurisdictional threshold of the law as a result of their operations as employers, regardless of the sectoral nature. Covered healthcare companies currently using automated employment decision tools to make certain employment decisions such as hiring, promotion, firing or even resume screening could be in violation of the NYC AEDT law, which is enforceable as of July 5, 2023.

U.S. Non-Binding Guidance

Various executive branch stakeholders and regulatory bodies have issued non-binding guidance on the regulation of AI. For example, the White House released its Blueprint for an AI Bill of Rights in October 2022, which provides a collection of principles in a

---

[5] Debevoise Data Blog: The Revised Colorado AI Insurance Regulations: What Was Fixed, and What Still May Need Fixing (May 31, 2023), *available* here.

[6] Debevoise Data Blog: NYC's AI Hiring Law Is Now Final and Effective July 5, 2023 (Apr. 12, 2023), *available* here.

rights-based approach to mitigating AI risk.[7] The principles include guidance regarding: safe and effective systems, discrimination protections, data privacy, notice and explanation capabilities, and human alternatives and fallbacks. The Blueprint notes that certain protections, such as data privacy and transparency, should be applied to automated systems with the potential to impact access to critical resources, such as healthcare. As recently as last week, President Biden announced a voluntary commitment between seven companies to commit to three fundamental principles of responsible innovation: safety, security and trust.[8]

The U.S. Department of Commerce's National Institute of Standards and Technology ("NIST") recently released its Artificial Intelligence Risk Management Framework 1.0 ("AI RMF"), which presents a non-binding, flexible framework designed to guide entities in their development and use of AI systems.[9] NIST's AI RMF describes the characteristics of trustworthy AI systems as those that are: reliable, valid, safe and secure, resilient, transparent, explainable and interpretable, and fair with respect to harmful bias management. NIST then outlines core tenets that it deems essential to developing and maintaining responsible AI systems. While this regulatory guidance is not directly applicable to healthcare companies, it provides helpful context for how regulatory bodies are thinking about AI and is thus important for healthcare companies to keep in mind.

The U.S. Data Privacy Landscape: Federal and State Developments

The U.S. data privacy landscape is comprised of numerous federal and state laws that address different aspects of data privacy. At the federal level, AI use by threat actors may expose covered entities and their business associates to liability under the Health Insurance Portability and Accountability Act of 1996 ("HIPAA"). Simultaneously, state data privacy laws and regulations implicate the use of consumers' health data. Below, we discuss each in turn.

Federal Data Privacy

HIPAA requires covered entities and their business associates to ensure the confidentiality and integrity of any electronic protected health information ("ePHI") utilized, accessed, disclosed or stored. On July 13, 2023, the HHS Health Sector

---

7 Debevoise Data Blog: The White House's Blueprint for an AI Bill of Rights: What It Gets Right and What It Gets Wrong About Artificial Intelligence Regulation (Oct. 26, 2022), *available* here.

8 President Biden's full statement is *available* here.

9 Debevoise Data Blog: Overview of Global AI Regulatory Developments and Some Tips to Reduce Risk (May 3, 2023), *available* here.

Cybersecurity Coordination Center ("HC3") published a brief on AI,[10] describing the threat AI-powered tools pose to the health sector and mitigation efforts healthcare entities should consider to better ensure their security strategies adequately address the evolving threats posed by AI. The HC3 warns that AI tools are being used by malicious actors to accelerate malware development and evade security solutions. Further, the HC3 brief describes the ease with which AI tools can be leveraged by malicious actors to create targeted phishing email templates with convincing lures to more easily trick recipients into opening malicious attachments or clicking malicious hyperlinks.

Although the HHS OCR—the federal agency responsible for HIPAA enforcement—has yet to issue formal guidance about the use of AI as it relates to HIPAA, OCR has stated that "HIPAA regulated entities should determine the potential risks and vulnerabilities to ePHI before adding any new technology into their organization."[11] HIPAA-regulated entities should conduct an AI-focused risk assessment and utilize the findings to update their risk management plans. The HC3 recommends healthcare entities consider the following AI risk mitigation strategies:

- review the "Artificial Intelligence Risk Management Framework" from NIST;
- review the "MITRE Atlas" knowledge base of adversary tactics, techniques and case studies for machine learning ("ML") systems;
- adopt AI-based tools for defense, including penetration testing, threat detection, threat analysis and incident response; and
- provide AI training for cybersecurity personnel.

State Data Privacy Laws

Generally applicable U.S. state privacy laws may implicate healthcare providers' use of AI. We previously provided guidance to companies on how to prepare for compliance with new state privacy laws,[12] and below we address a few of these laws and their applicability to healthcare companies implementing AI.

California's Consumer Privacy Act (as amended by the California Privacy Rights Act that went into effect on January 1, 2023, ("CCPA")) established a new California Privacy Protection Agency ("CPPA").[13] The CPPA was charged with adopting regulations

---

[10] The HC3 brief is *available* here.
[11] More on OCR's statement is *available* here.
[12] Debevoise Data Blog: Getting Ready for 2023: What Companies Can Do Now to Prepare for New Privacy Laws (Dec. 16, 2021), *available* here.
[13] The text of the CCPA is *available* here.

"governing access and opt-out rights with respect to businesses' use of automated decision-making technology," including providing meaningful information about the logic of the decision and the likely outcome with respect to the consumer. Notably, the CPPA's mandate to issue automated decision-making regulations is not currently limited to "solely" automated decisions or those with legal effects. Healthcare companies using or considering implementing AI solutions should monitor developments in this space that may require the implementation of opt-out rights or other guardrails.

Illinois has also established a robust data privacy regulatory scheme that could apply to healthcare companies' use of AI. Illinois' Biometric Information Protection Act ("BIPA") protects consumers from the collection, use and sharing of their biometric information without prior consent.[14] We have previously described the legal risks of voice analytics and the potential for lucrative class actions under BIPA given that it includes a right of private action of $1,000 – $5,000 per violation.[15] Whereas BIPA specifically covers biometric identifiers that could be traced back to an individual, such as fingerprints and retinal scans, Illinois' forthcoming Data Privacy and Protection Act ("DPPA") aims to cover a broader scope.[16] If it passes, the DPPA would limit companies from collecting, processing or transferring "sensitive covered data," which includes information regarding an individual's past, present or future physical or mental health, without consent.

Similarly, Washington state recently enacted the My Health My Data Act ("MHMD"), which regulates businesses that process "consumer health data," which is personal information linked to a consumer that can identify their past, present or future physical or mental health.[17] MHMD requires businesses to obtain consent from customers before processing their health data, unless it is necessary to provide a product or service that the consumer has requested. The law also requires companies to obtain separate consent to sell or share consumer health data. Both the forthcoming DPPA and Washington's MHMD regulations could present limitations on healthcare companies seeking to use AI tools.

BIPA, MHMD and other state privacy laws have tangible implications in the healthcare space. This is particularly evident when considering the legal implications of wearable technology (i.e., wearables). Wearables typically use biosensors to monitor the health and wellness of the individual wearing them. Some common examples of wearables include smartwatches, fitness trackers and continuous glucose monitoring devices.

---

[14] The text of Illinois' BIPA is *available* here.
[15] Debevoise Data Blog: Legal Risks of Using AI Voice Analytics for Customer Service (Jan. 10, 2023), *available* here.
[16] The text of Illinois' DPPA is *available* here.
[17] The text of Washington's MHMD is *available* here.

These products may not necessarily track biometric identifiers that would allow the identification of the individual wearing them, such that they would fall under the purview of BIPA. However, these devices do generally track information related to the individual's past and present physical health. Entities that use wearables to collect or access individuals' data should thus consider whether Washington's MHMD or Illinois' DPPA, if it passes, would apply to them. Companies creating such wearables could consider disclosing their use of health data and obtaining appropriate consent from individuals such that they would be in compliance with existing data privacy regulations.

Federal Agency Oversight: FTC, FDA and DOJ Developments

Healthcare companies that seek to incorporate AI must also keep aware of the various federal agencies that exert regulatory authority over different aspects of the technology: the FTC regulates deceptive and misleading advertising claims, including those that pertain to AI; FDA regulates the use of AI in medical devices, as well as drug and biological product development; and the DOJ enforces fraud, waste and abuse laws which may cover false claims submitted as a result of using AI for administrative purposes. Below, we discuss how these federal agencies are adapting their regulatory and enforcement efforts to incorporate AI considerations and use.

FTC Regulation of AI-Related Claims in Advertising

Due to the proliferation of AI products and the related marketing of such products in recent years, the FTC has repeatedly stated its intention to crack down on deceptive and misleading AI-related advertising claims.[18] The FTC has regulatory authority over most advertising claims made for AI products, and has indicated that it plans to focus on the following areas when evaluating whether AI-related claims are deceptive: (1) exaggerations as to what an AI product can actually do; (2) promises that an AI product provides superior performance to a non-AI product; and (3) whether a product actually utilizes AI at all.[19] In particular, the FTC has warned against "overusing and abusing" the term AI as a marketing tool. As AI continues to be integrated into the life sciences and healthcare fields, companies and investors should carefully evaluate AI-related product claims prior to dissemination, as it is clear the FTC will enforce against violative claims.

Although the FTC has not yet announced any significant enforcement actions for deceptive AI-related advertising, companies engaging in deceptive advertising are being

---

[18] *Keep Your AI Claims in Check*, FED. TRADE COMM'N (Feb. 27, 2023), *available* here.

[19] *See* Debevoise In Depth: Risks of Overselling Your AI: The FTC is Watching (Mar. 6, 2023), *available* here.

held accountable by other entities; these actions may presage the types of claims the FTC will bring under its own enforcement authority, and thus healthcare companies should monitor such developments. For instance, Engineer.ai, a startup that claimed to have built an AI-assisted app development platform, was sued by its chief business officer who claimed the company was exaggerating its use of AI technology.[20] It was unveiled that the company was largely using human engineers to build apps for customers while leveraging the popularity of terms like "AI" to drive business and raise additional funds to further develop its AI capabilities.

Certain class action lawsuits also provide an apt model for the types of claims the FTC may seek to bring against healthcare companies utilizing AI. Zillow, for example, was sued in a class action lawsuit for representations related to its program Zillow Offers, an instant buying organization that used data matching algorithms to provide cash offers on for-sale homes.[21] The suit alleged that Zillow was aware that its algorithm failed to accurately predict future home prices, yet company leadership continued to promote its success to both shareholders and the public. Similarly, several actions were brought against DoNotPay, a company that claimed to have produced the first AI "robot lawyer."[22] One class action alleged that the company advertised that its AI lawyer could perform on par with traditional legal services provided by a lawyer, but, in reality, its performance was allegedly inadequate. A separate petition for pre-action discovery alleged that DoNotPay did not actually employ AI at all.

As healthcare and life sciences companies continue to adopt AI tools, they should be mindful of FTC's regulatory authority and take concrete steps to mitigate potential liability. Companies should ensure the appropriate disclosure of AI tools, as the failure to mention that a company is using AI for patient care (e.g., diagnostics, treatment recommendations) may constitute a material omission that could invite scrutiny. When a company is making direct claims related to its AI use, such claims should be truthful and non-misleading. Equally important are internal education efforts for both marketing and compliance teams on FTC guidance regarding deceptive advertising. Finally, companies should implement a thorough review process for all current and proposed claims to evaluate claim substantiation and verify that the claim does not exaggerate the capabilities of any AI product.

---

[20] *See This AI Startup Claims to Automate App Making But Actually Just Uses Humans* (Aug. 14, 2019), *available* here.

[21] *See Seattle-based Zillow Faces Another Shareholder Suit Over Failed House-Flipping Business* (July 25, 2022), *available* here.

[22] *See Analysis: DoNotPay Lawsuits: A Setback for Justice Initiatives?* (Mar. 28, 2023), *available* here.

FDA Regulation of AI and ML Medical Devices

Over the past several years, FDA has been developing a framework to regulate the use of AI in FDA-regulated products.[23] Whereas FDA's traditional medical device regulatory framework is predicated on approval or clearance of static, unchanging devices, the dynamic nature of AI and ML (collectively, "AI/ML") requires flexibility to accommodate changes based on the addition of new data.[24] In particular, medical devices designed to continually update and adapt in real time based on new input data challenge FDA's longstanding regime, necessitating the development and implementation of an updated regulatory process.

To address the recent dramatic increase in the utilization of AI/ML in medical devices, FDA released its AI/ML-Based Software as a Medical Device ("SaMD") Action Plan in January 2021,[25] followed by draft guidance in April 2023, formalizing its intended approach to regulation of AI/ML-based medical devices.[26] By establishing clear expectations, including the concept of a predetermined change control plan ("PCCP"), this guidance is expected to spur innovation and increase the number of approvals and clearances of AI medical devices while providing a reasonable assurance of safety and efficacy. Specifically, under this guidance, FDA would allow medical device manufacturers to include a PCCP for a device as part of the device marketing application (e.g., 510(k) application) and, if successful, would permit the manufacturer to make modifications anticipated by the PCCP without additional marketing submissions to FDA.

As AI capabilities continue to develop and companies increasingly incorporate AI/ML into their products, medical device companies and investors must be cognizant of anticipated changes to the current regulatory regime to ensure compliance with relevant laws and regulations,[27] including the broad state privacy laws discussed above,

---

[23] In addition to medical devices, discussed herein, AI is also implicated in drug and biological product development. In May 2023, FDA released a discussion paper addressing the use of AI/ML in the development of drug and biological products. FDA, Using Artificial Intelligence & Machine Learning in the Development of Drug & Biological Products, Discussion Paper and Request for Feedback (2023), *available* here.

[24] ML algorithms are data-driven AI systems that "learn" from examples in large datasets (i.e., training sets) without being explicitly programmed to reach a particular answer or conclusion. These algorithms can learn to decipher data patterns at scales unattainable by humans to identify relationships between the input (e.g., radiologic images) and output (e.g., diagnoses/clinical decision support).

[25] FDA, Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD) Action Plan (Jan. 2021), *available* here.

[26] FDA, Draft Guidance: Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions (Apr. 3, 2023) ("PCCP Guidance"), *available* here.

[27] For more information, *see* Debevoise In Depth, Artificial Intelligence and the Life Sciences Industry: FDA and FTC Regulatory Update (May 16, 2023), *available* here.

which add regulatory complexity and may impose additional requirements on healthcare companies.

Companies seeking to incorporate AI/ML into their medical devices should review FDA's PCCP Guidance to understand how FDA intends to regulate changes to medical devices. To date, the algorithms used in most authorized AI/ML devices have remained static because any change would require an additional submission to FDA. An exception came on February 7, 2020, when FDA announced the marketing authorization, through the De Novo pathway, of Caption Guidance software, the first cardiac ultrasound software using AI to help users capture images of a patient's heart for diagnostic purposes. The Caption Guidance software incorporates a locked algorithm—i.e., the algorithm is not adaptive, and changes must be implemented manually in accordance with the PCCP. To date, no products have been approved with PCCPs that cover automatic updates (whereby the software automatically changes based on continuous learning or adaptive models).

Accordingly, when developing a medical device that incorporates AI/ML, companies should carefully consider how the device will be updated (i.e., manually or automatically). FDA has indicated that its review of automatically implemented modifications will be comparatively more complex and incorporate a risk-benefit assessment. To facilitate the regulatory process, companies should initiate a dialogue with FDA and begin drafting the PCCP early in the development process. Pursuant to the PCCP Guidance, FDA recommends that a PCCP include a limited number of modifications that can be verified and validated to allow for efficient review.[28]

DOJ Enforcement of Fraud, Waste and Abuse Laws

Although efforts to combat healthcare-related fraud, waste and abuse are myriad, the primary tool used by the DOJ is the False Claims Act ("FCA"). The FCA imposes civil penalties for knowingly (or with reckless disregard) submitting or causing others to submit false records, statements or claims for payment to the federal government, as well as for wrongfully concealing or failing to return an overpayment. FCA actions may be brought by the DOJ or by private individuals with nonpublic information about

[28] FDA recommends the PCCP include: (1) a detailed description of each planned modification to the device, including the rationale for the modification and anticipated changes to the device characteristics and performance; (2) a modification protocol describing the methods that will be followed when developing, validating and implementing modifications, including how information on the modifications will be communicated to users and how real-world data on the impact of the modifications will be monitored after implementation; and (3) an assessment of the benefits and risks of implementing a PCCP.

alleged misconduct (i.e., "relators").[29] Further, the majority of states have state-based FCA analogues that allow relators to bring claims on behalf of the state.

In addition to clinical applications of AI/ML-based software,[30] described below, the ability to leverage these technologies to streamline administrative tasks, including billing and coding, can result in substantial efficiencies. By automating such processes, however, there is a risk that software may generate inaccurate or upcoded claims, thereby creating the potential for federal and state FCA liability. For example, autonomous coding solutions have been developed to lessen the volume of manual coding, reduce costs and alleviate administrative burden. To do so, these platforms use AI/ML to translate clinical evidence in electronic health records ("EHR") into billing codes. Liability arises for upcoding if such codes are submitted and deemed to represent more expensive diagnoses or services than the provider actually diagnosed or performed. Relatedly, by using algorithms to automate claims management, there is also risk of liability for improper claim denial.[31]

In accordance with traditional billing and coding assessments, it is crucial to ensure that AI applications are trained to engage in appropriate downcoding to mitigate the risk of fraud. Additionally, to counterbalance the use of such AI/ML software, periodic manual compliance reviews should be implemented as standard practice.

Ethics, Licensure and the Corporate Practice of Medicine ("CPOM")

While proponents of healthcare AI laud its potential to revolutionize care quality and delivery, they must also grapple with a host of complex ethical considerations. For instance, AI poses a well-documented but poorly understood "transparency problem." The term "transparency" in the context of AI carries multiple meanings—not only the plain language meaning, which queries whether people know that AI is involved in making a particular prediction, recommendation or decision, but also the concept of "explainability"—that is, "do we know how an AI tool is making the decisions it is making?" and, ancillary to that, "do we know how to get a better result?" A physician using an AI-powered clinical decision-making tool may, for instance, be unable to

[29] Relators are incentivized to bring such lawsuits on behalf of the government (i.e., "qui tam" actions) because they are entitled to receive a significant share of any monetary recovery. Lawsuits brought by relators remain under seal while the government investigates and decides whether to join. If successful, the government can potentially recover treble damages and statutory fines of $25,000 per violation.

[30] One relevant consideration is whether AI-assisted clinical services should be billed under alternate billing codes or for fewer units of time.

[31] For example, health insurers that use AI to deny claims for lack of medical necessity have been subject to class action lawsuits.

explain to a patient why the tool's algorithm generated a particular diagnosis. Further, because algorithms must be trained on high-quality datasets, flaws in data may lead to algorithmic bias: given the number of independent actors (and potentially the "independence" of the AI algorithm itself), accountability gaps may arise that stall remediation of such flaws.

The CPOM doctrine aims to protect patients by prohibiting nonlicensed individuals and entities (other than professional corporations) from practicing medicine or employing a physician to provide professional medical services. Although regulated at the state level and, therefore, enforced to varying degrees, CPOM regulations are intended to promote the physician-patient relationship by limiting clinical decision-making to duly licensed healthcare professionals.

Within the realm of AI, CPOM raises interesting questions with respect to clinical decision support ("CDS") tools, particularly those that are designed to dynamically update based on the continuous adaptation of new inputs. CDS tools can quickly analyze large datasets, identify patterns and provide, among other things, alerts, clinical guidelines, diagnostic support and contextually relevant reference information.[32] Although such applications can enable healthcare practitioners to make more informed clinical decisions, AI/ML-powered CDS software cannot be independently licensed and, as a result, could be deemed to exert undue control over clinical decision-making in contravention of CPOM laws. Further, because healthcare providers are liable for the practice of medicine, where use of AI-informed CDS tools is permitted, practitioners cannot wholly rely on such outputs and must retain their own professional judgment.[33]

Advocates of AI describe useful applications for both imaging and diagnostic purposes. For example, AI algorithms can scan chest X-rays to simultaneously evaluate a patient for a multiplicity of conditions.[34] These tools can empower physicians to make diagnostic decisions more quickly and accurately. In areas like radiation therapy, AI is being used to produce a tailored treatment plan based on the patient's medical records, calculating and recalculating radiation dosages based on changes in the patient's anatomy.[35] AI incorporating CT or MRI imaging can create three-dimensional, interactive anatomy models that provide information on tumor size, location and vascular structure for surgical planning.[36] By expediting therapeutic treatment and

---

[32] FDA, Clinical Decision Support Software, Guidance for Industry and Food and Drug Administration staff (Sept. 2022), *available* here.

[33] *See id.* describing when FDA considers a CDS tool to be an FDA-regulated medical device.

[34] *See Artificial Intelligence Rivals Radiologists in Screening X-Rays For Certain Diseases* (Nov. 20, 2018), *available* here.

[35] *See AI Can Jump-Start Radiation Therapy for Cancer Patients* (Jan. 27, 2020), *available* here.

[36] *See The Potential for Machine Learning Algorithms to Improve and Reduce the Cost of 3-Dimensional Printing for Surgical Planning* (May 2018), *available* here.

continuously tailoring treatment as the patient progresses, AI can improve the effectiveness of existing therapies and streamline the physician's care management process. As AI tools receive access to greater volumes of specialized training data, accuracy and usefulness are likely to increase.

On the other hand, some thought leaders in the healthcare industry have voiced concerns regarding the efficacy of AI and remain skeptical of its ability to perform in real-life clinical contexts, which are immensely more complex than the textbook-style cases that AI may be trained on. Certain case studies have demonstrated that initial AI systems that appear to boast high rates of diagnostic accuracy can exhibit errors or markedly lower accuracy rates once tested in clinical practice. For example, Google Health's AI tool, designed to scan images of diabetic patients' eyes for diabetic retinopathy, initially demonstrated up to 90 percent accuracy in internal tests.[37] During clinical testing, however, unforeseen environmental variables, such as lower-quality retinal scans submitted by clinic nurses, resulted in the AI tool rejecting over 20 percent of the images, causing substantial delays and frustration; the high number of rejected images led to unnecessary follow-up appointments for images the nurses believed "showed no signs of disease." While Google Health's AI was trained on high-quality scans to ensure accuracy, this counterintuitively hampered the tool's real-world usability. In a recent interview, the president of the American Medical Association ("AMA") expressed his belief that the types of AI algorithms currently in use are too narrow and cannot fully grasp the nuanced and complex considerations (e.g., medical, social and psychiatric backgrounds) that accompany patients, nor can current AI adapt to the different goals that patients may have with respect to their care.[38] In essence, the concern is that patients often present with complicated conditions, but AI is designed to produce generic answers; in such situations, a clinician's experience and judgment may be invaluable.

To address both the ethical and legal risks that arise from the use of AI in healthcare, companies should strive to create structured environments that give physicians meaningful control over AI while still allowing them to extract the enormous benefits that the technology can provide. Creating such a space for AI requires advanced knowledge of the types of training data being used. Disparate datasets for training AI can cause the same algorithm to reach starkly different conclusions; vendors and healthcare professionals should know where the data used to train a specific AI tool has deficiencies to allow for appropriate testing and remediation before it is used with patients. Structured environments also require knowledge of the expected scenarios an AI tool may face, including whether that number is indeterminate, as this will inform

---

[37] *See Google's Medical AI Was Super Accurate In a Lab. Real Life Was a Different Story* (Apr. 27, 2020), *available* here.

[38] The full interview with AMA president Jesse Ehrenfeld is *available* here.

what information should serve as inputs; improper inputs may lead to inappropriate outputs. And finally, the intended use of the AI tool and desired outputs should be clearly defined, because without this information the effectiveness of the tool cannot be measured. Healthcare companies that are equipped with this knowledge can begin to peel back the opaque lid of the so-called "black box" to increase both accuracy and accountability. Such knowledge is not a panacea for the ethical and legal concerns that accompany AI, but it represents a significant step towards meaningfully addressing them.

Intellectual Property

As companies adopt the use of AI technologies into their business operations, they will want to assess the IP issues implicated by such use. The input data used to train, develop or operate an AI model may be subject to third-party IP or other proprietary rights or license terms that limit, restrict or condition the use of such data. As healthcare companies source various types of input data to train, develop and use with varying types of AI models, it is important to consider the potential ownership claims and consent obligations that may arise, including those from patients and researchers involved in the creation of such inputs. Further, where companies use AI models to generate new content, they may encounter barriers to claiming IP rights and associated protections on such outputs (e.g., if use of the input data was infringing or if the AI model simply regenerates unaltered portions of a copyrighted work as an output).

An example of a commercial arrangement that underscores many of the IP issues present in this intersection between healthcare and AI is the exclusive licensing deal that one healthcare organization entered into with a technology start-up engaged in the development of AI applications focused on clinical diagnostics. The healthcare organization granted the start-up the exclusive rights to a large volume of biomedical and research data. Although the biomedical dataset would have been de-identified, there were potential questions as to whether the necessary consents were obtained, or notices given, to allow for the commercialization and use of any underlying patient data. Additionally, the licensed research data, in some cases, resulted from federally funded research, which raised concerns about making such research exclusively available to a for-profit entity.

Conducting diligence on the source of any input data, including by undertaking efforts to understand how the data was obtained, identify whether there are any contractual conditions or other limits on the scope of use for such data, and obtain any necessary consents, is an important risk mitigation measure in the context of developing and using AI tools in the healthcare space. It is also important to consider the relevant

stakeholders who may have rights to the input data or who may seek to exercise rights or attribution over the outputs of the AI models, which could include researchers and other contributors to the materials, patients and individuals whose data is used, and individuals or organizations involved in the funding of such research. When entering into contracts for the development and use of AI tools by third parties, companies should ensure there are clear contractual terms with respect to the use of the input data that align with the permitted scope of use for the sourced data. Companies should also obtain appropriate reps and warranties and indemnifications that provide that their use of the AI tools will not infringe a third party's rights.

* * *

We will continue to monitor the legal landscape governing the use of AI in healthcare. Please do not hesitate to contact us with any questions.

Regulatory Authority	Applicable Law(s)/Proposal(s)	Scope	Jurisdiction
European Commission, Council of the European Union, European Parliament	EU Artificial Intelligence Act (proposed)	AI systems broadly	European Union
United States Executive Branch	Blueprint for an AI Bill of Rights	AI systems broadly	United States (non-binding guidance)
U.S. Department of Commerce, National Institute of Standards and Technology	Artificial Intelligence Risk Management Framework 1.0	AI systems broadly	United States (non-binding guidance)
U.S. Department of Health and Human Services, Office of Civil Rights	Health Insurance Portability and Accountability Act of 1996	Governs HIPAA covered entities (i.e., payers, providers, health care clearinghouses) and business associates that electronically transmit protected health information	United States
U.S. Department of Justice	False Claims Act (31 U.S.C. §§ 3729-3733)	Records, statements or claims submitted to the U.S. federal government	United States
U.S. Food and Drug Administration	Food, Drug, and Cosmetic Act; AI/ML-Based Software as a Medical Device Action Plan	AI/ML used in FDA-regulated products	United States

Regulatory Authority	Applicable Law(s)/Proposal(s)	Scope	Jurisdiction
U.S. Federal Trade Commission	Federal Trade Commission Act Section 5, Fair Credit Reporting Act, Equal Credit Opportunity Act	Most advertising claims	United States
California Attorney General; California Privacy Protection Agency	California Consumer Privacy Act, as amended by the California Privacy Rights Act	Governs certain covered entities that collect, use, share or sell California consumers' "personal information"	California
Colorado Division of Insurance	Algorithm and Predictive Model Governance Regulation	Governs insurers authorized to do business in the state of Colorado that use external consumer data and information sources ("ECDIS"), algorithms and predictive models that use ECDIS	Colorado
Civil Lawsuit (Statutory Private Right of Action)	Biometric Information Protection Act	Governs private entities that operate or do business in Illinois in possession of Illinois residents' biometric identifiers or biometric information	Illinois
Illinois Attorney General, Civil Lawsuit (Statutory Private Right of Action)	Data Privacy and Protection Act (proposed)	Governs covered entities' collection, processing or transferring of Illinois residents' "sensitive covered data"	Illinois
New York City Department of Consumer and Worker Protection	Automated Employment Decision Tools Law	Governs the use of "Automated Employment Decision Tools" by Covered Entities	New York City
Washington Attorney General	My Health My Data Act	Governs data controllers that do business or provide goods and services to Washington residents that collect, process, share or sell consumer health data	Washington State
Various State Agencies	Corporate Practice of Medicine Doctrine	Prohibits the ownership of a medical practice or employment of medical professionals by non-licensed individuals or companies	Majority of states

Regulatory Authority	Applicable Law(s)/Proposal(s)	Scope	Jurisdiction
Private Parties	Intellectual Property Rights	Patents, copyrights, trademarks and trade secrets	United States

NEW YORK

Andrew L. Bab
albab@debevoise.com

Jennifer L. Chu
jlchu@debevoise.com

Avi Gesser
agesser@debevoise.com

Kevin Rinker
karinker@debevoise.com

Caroline P. Geiger
cpgeiger@debevoise.com

Hannah R. Levine
hrlevine@debevoise.com

Michael L. Cederblom
mlcederblom@debevoise.com

Lex Gaillard
adgaillard@debevoise.com

Clara Geffroy
cgeffroy@debevoise.com

WASHINGTON, D.C.

Paul D. Rubin
pdrubin@debevoise.com

Tigist Kassahun
tkassahun@debevoise.com

Melissa Runsten
mrunsten@debevoise.com

SAN FRANCISCO

Kim T. Le
kle@debevoise.com

Mengyi Xu
mxu@debevoise.com