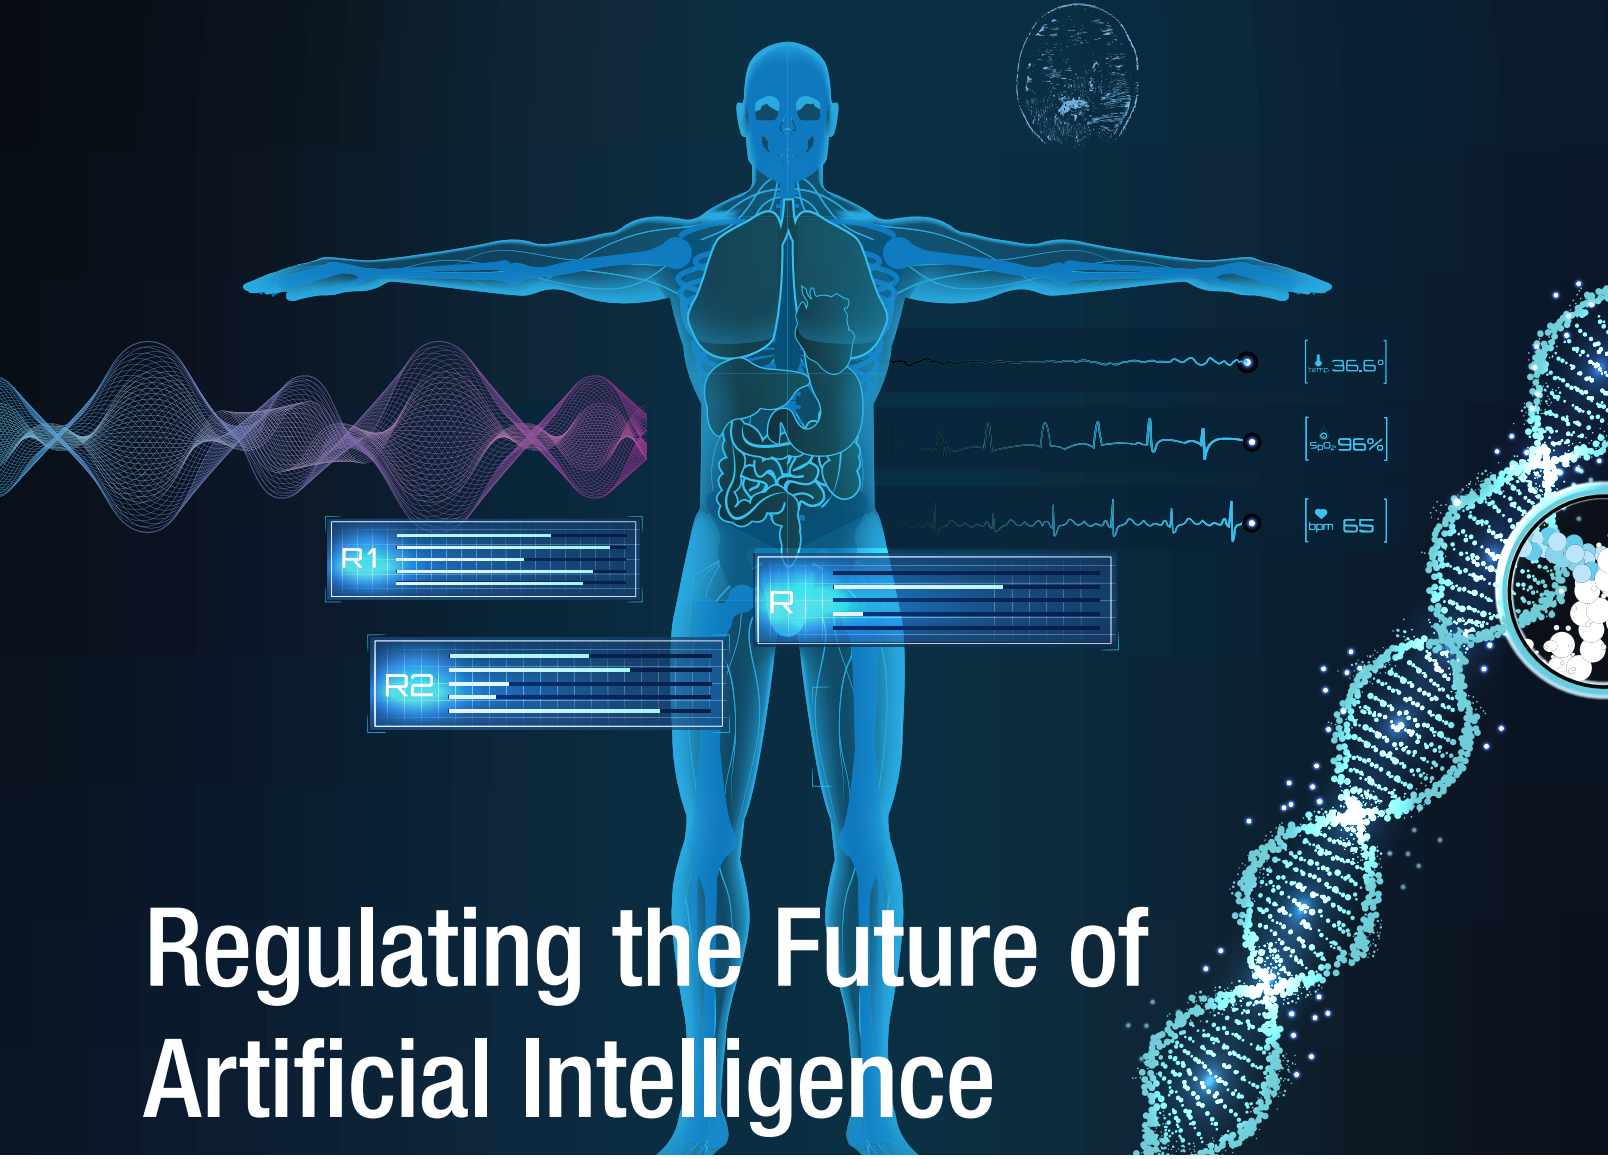

Regulating the Future of Artificial Intelligence

By Carolyn V. Metnick, McDermott Will & Emery LLP, and Dina B. Ross, Dina B. Ross Law Offices

Artificial intelligence (AI) is becoming a more frequent discussion item among health care providers, clinicians, manufacturers, and developers. With the rise and deployment of AI, we expect to see earlier disease detection, greater diagnostic accuracy, and more tailored treatment plans. The possibilities are beyond our imaginations, and we are only beginning to see the opportunities and potential. Although interest in AI is growing, and the use of AI is becoming more common, the regulatory framework in the United States remains in the nascent stages of development. This article provides an overview of the existing and evolving U.S. federal regulatory framework.

What Is Artificial Intelligence?

AI means different things to different stakeholders, but it is generally and broadly defined as computer engineering that results in the imitation of intelligent behavior or human learning and reasoning. There are different classes of AI. At one end of the spectrum we have impressive but basic reactive machines that store millions of bits of data and can weigh best options based on that data, but that cannot learn from past mistakes. Computer chess is an example of a reactive machine. At the other end of the spectrum are more sophisticated examples, including machine learning (ML), where a system uses human-programmed algorithms to review data over and over again until it "learns" to determine something or predict an event. A queue of recommended movies based on prior viewing habits is an example of ML. One particularly sophisticated form of ML is deep learning, where a machine "learns" similarly to a human's neural system. It is fed an enormous amount of data until it piles layers of data together, discovering patterns, pathways, interconnectivities, and relational concepts in the data far beyond what it was programmed to identify. Driverless cars are an example of deep learning.

AI applications in health care typically constitute software as a medical device[1] (SaMD), which is software for one or more medical purposes that performs without being part of a hardware medical device and may run on different operating systems or in virtual environments. The Food, Drug and Cosmetic Act considers medical purposes to be purposes that are intended to treat, diagnose, cure, mitigate, or prevent

disease or other conditions. An example of SaMD is software that allows a smartphone to view images from an MRI. Software that drives or controls the motors and pumping of medication in an infusion pump, however, is not SaMD, as the hardware medical device requires the software to perform its intended use. ML and deep learning AI differ from other SaMD in that each has the potential to adapt and improve quickly and in real time.

The U.S. Food and Drug Administration (FDA) is responsible for the regulatory framework for AI, yet it has encountered challenges and continues to wrestle with the framework for reasons that are as unique as AI itself.

> AI means different things to different stakeholders, but it is generally and broadly defined as computer engineering that results in the imitation of intelligent behavior or human learning and reasoning.

21st Century Cures Act

The 21st Century Cures Act (Act), which was enacted in December 2016, codified that certain digital technologies fall outside the FDA's regulation authority. The Act specifically removed certain clinical decision support software (CDS) from the definition of a medical device—and therefore from the FDA's authority—as long as the CDS does not:

> acquire, process, or analyze a medical image or signal from an in vitro diagnostic device, or a pattern or signal from a signal acquisition system for the purpose of—(i) displaying, analyzing, or printing medical information about a patient or other medical information (such as peer-reviewed clinical studies and clinical practice guidelines); (ii) *supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition*; and (iii) *enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any such recommendation to make a clinical diagnosis or treatment decision regarding an individual patient*.[2]

Although the Act excludes certain CDS from the definition of medical device, the Act specifically grants the FDA discretion to determine that a technology is a device if the Department of Health and Human Services Secretary makes a finding that the software function would be reasonably likely to have serious adverse health consequences, and the software function has been identified in a final order issued by the Secretary after considering: "(i) the likelihood and severity of patient harm if the software function were not to perform as intended; (ii) the extent to which the software function is intended to support the clinical judgment of a health care professional; (iii) whether there is a reasonable opportunity for a health care professional to review the basis of the information or treatment recommendation provided by the software function; and (iv) the intended user and user environment . . . ."[3]

AI that does not explain how input data is analyzed or otherwise "show its work" clearly requires FDA regulation, as it does not allow for the independent review of the basis for such recommendations, and a false diagnosis could be deadly. A black box algorithm would fall into this category. Software that allows the independent review of the basis for the recommendation may or may not be subject to FDA regulation depending on the Secretary's weighing of the four aforementioned variables. Of course, the FDA retains enforcement discretion. Assuming the software produces a diagnosis, and a provider relies on that diagnosis without the ability to question or review it, the software would clearly be subject to FDA regulation. The challenge that presents is how the FDA should regulate a device that learns and that is continuously and constantly evolving. That is the underlying difficulty in regulating a ML or deep learning modality. For example, the device may not be the same thing now that it was five minutes ago because during the intervening time, it has continued to obtain new information, adapted to that information for subsequent purposes, and improved its techniques.

There are other overarching challenges as well. Imagine, for instance, a virtual personal assistant providing a turnkey experience in managing your health care. It does not take much imagination to envision a voice-activated smart device having the ability to monitor the user's health vitals and to automatically schedule a doctor's appointment—at a time that is convenient with other scheduled priorities—based on the urgency of the perceived risk. The smart device could also track and schedule routine checkups, vaccines, and tests; reorder prescriptions; and notify emergency contacts when something is amiss. The smart device could also tell your self-driving car to fully take over if your blood pressure or blood sugar drops or spikes, and the smart device could be programmed to automatically send all results to your physician. For routine checkups, you may not even need to leave work or home, as your smart device could double as a telemedicine platform, taking your blood pressure, analyzing blood oxygen, and running other assessments of vital signs, all while you interface in real-time with your provider. AI can also help your doctor make sense of decades worth of electronic health data, identifying seemingly unrelated symptoms and ailments and identifying trends in laboratory results and clinical assessments that assist your provider in accurately diagnosing and prescribing a customized treatment based on your unique clinical history and physiology. Any prescriptions could be at your front door before you are home from work, and dosing reminders are automatically populated in your smart device, which will schedule a follow-up visit if your symptoms do not improve. Although this was a seemingly fantastical notion ten years ago, this possibility is on our horizon, and every single step poses a challenge in

today's regulatory framework. Existing laws and regulations will need to be overhauled, while new laws and regulations will need to be crafted to reflect our new AI-driven reality. The FDA is working to make this happen.

FDA's Consideration of Total Product Lifecycle-Based Framework

The FDA already has processes through which it reviews new medical devices, including 510(k) notification, de novo, or the premarket approval application. As a brief background, the 510(k) process arises when a manufacturer brings to market a moderate-risk device substantially equivalent to a device that has already been approved by the FDA. The FDA de novo process is reserved for low- or moderate-risk devices with no substantially equivalent device on the market, and the premarket approval process is meant for high-risk devices requiring demonstrated safety effectiveness before the FDA issues its approval.

The FDA has guidance to assist manufacturers in the decision whether to seek authorization to market a software change to a previously approved medical device. Generally, the FDA considers the 510(k) pathway to be potentially appropriate for software changes to a previously approved device, depending on the outcome of a risk-based flowchart provided by FDA.[4] Modifications to previously approved 510(k) devices that may require subsequent premarket submission include changes that introduce a new risk or modify an existing risk that could result in significant harm, a change to risk controls to prevent significant harm, and a change that significantly affects clinical functionality or performance.

Acknowledging that the existing framework does not accommodate a device that changes quickly and in real time, as an AI device does, the FDA is exploring a total product lifecycle approach for AI and ML technologies. In April 2019, the FDA published a paper entitled *Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)—Based Software as a Medical Device (SaMD)* that lays out the framework the FDA is considering (FDA Paper). The FDA Paper proposes a framework based on the International Medical Device Regulators Forum (IMDRF) risk categorization principles, the FDA's benefit-risk framework and risk management principles in the software modification guidance, and the organization-based total product life cycle approach envisioned in the Digital Health Software Precertification Program.

In conjunction with the FDA Paper, the FDA Commissioner released a statement promoting the FDA's commitment to applying "current authorities in new ways to keep up with the rapid pace of innovation and ensure the safety of [artificial intelligence] devices."[5] The Commissioner noted that devices with AI capabilities that have been approved to date were limited to "locked" algorithms or reactive machines that do not continually update and improve and acknowledged the need for a framework that would allow for modifications to be made to algorithms from real-world learning and adaptation while ensuring safety.

> The challenge that presents is how the FDA should regulate a device that learns and that is continuously and constantly evolving.

Risk Categorization

The framework proposed in the FDA Paper evokes the IMDRF risk categorization framework but does not explain how it would be applied. The IMDRF risk categorization framework examines the significance of the information provided by software and the state of the health care situation or condition. The significance of the information pertains to the intended use of the information provided to treat or diagnose and inform clinical management. The state of the health care condition relates to the seriousness of the disease or condition. Together these two factors can categorize the software from low to high risk. For example, if the software is only being used to inform clinical management (as opposed to diagnosing), and the health care condition is non-serious, the risk classification would be low. On the other hand, if the software is being relied upon to diagnose a patient, and the patient has a critical health condition, then the risk would be very high. The FDA Paper does not address how the risk classification will direct the review. Presumably, the higher the risk classification, the greater the scrutiny to be provided.

The FDA assumes that most software modifications to existing, approved medical devices will involve changes to coding and programming that result in different decision matrices and retraining existing algorithms with new data sets, which would be subject to premarket review. The FDA contemplates three categories of changes: performance, input, and intended use. The performance category includes modifications related to performance with no change to the intended use or new input type. The inputs category contemplates changes to the inputs used by the algorithm and their clinical association to the SaMD output. Modifications related to intended use include those that result in a change in significance of information provided by the software. The FDA notes that changes may not be mutually exclusive, and that while many changes would be subject to premarket review, some modifications may not require such review. Although the FDA contemplates different categories of modifications, the FDA Paper fails to describe how different modification categories will be treated.

Pre-Certification

The FDA Paper contemplates a pre-certification process whereby a device manufacturer is pre-certified by the FDA to qualify for a more streamlined premarket review based upon meeting certain criteria and demonstrating a history of and commitment to quality, excellence, and safety.

The Software Pre-Cert Program is an existing voluntary FDA program through which FDA evaluates a device manufacturer's

"culture of quality and organizational excellence" to determine whether the organization has demonstrated its ability to "build, test, monitor and proactively maintain and improve the safety, efficacy, performance and security of their medical device software products, so that they meet or exceed existing FDA standards of safety and effectiveness."[6] This Program is being used for approval of electronic medical record systems. FDA's approach in this regard is based on several principles: (1) establishing clear expectations on quality systems and good ML practices; (2) conducting premarket review for SaMD that requires premarket submission to demonstrate reasonable assurance of safety and effectiveness and establishes clear expectations for manufacturers; (3) expecting manufacturers to monitor the device and incorporate a risk management approach in development, validation, and execution of algorithm changes; and (4) enabling increased transparency to users and FDA using postmarket real-world performance reporting.[7]

With respect to the first principle, the FDA describes good ML practices as those akin to good software engineer practices or quality system practices, and notes the following considerations for SaMD:

- Relevance of available data to the clinical problem and current clinical practice;
- Data acquired in a consistent, clinically relevant, and generalizable manner;
- Appropriate separation between training, tuning, and test datasets; and
- Appropriate level of transparency of the output and the algorithm aimed at users.[8]

With respect to the second principle of premarket assurance of safety and effectiveness, the FDA proposes a "predetermined change control plan" that includes the types of anticipated modifications, which it places in two categories: SaMD Pre-Specifications (SPS) and Algorithm Change Protocol (ACP). SPS include modifications and changes to performance or inputs related to the intended use of the AI (i.e., what the manufacturer intends the algorithm to become as it learns). The ACP specifies the methods in place to control the risks of anticipated modifications.

The third principle calls for manufacturers to evaluate modifications to SaMD based on risks to patients. The software modifications guidance expects manufacturers to perform a risk assessment and evaluate whether the risks are reasonably mitigated, and then either submit a new 510(k) for premarket review or document the modification and analysis in the risk management and 510(k) files.

The final principle involves transparency about the device's intended function and any associated modifications in furtherance of safety. Manufacturers would commit to transparency and real-world performance monitoring and provide periodic reports to the FDA on updates that were implemented as part of the approved SPS and ACP as well as performance metrics. The FDA notes that transparency may include providing updates to the FDA, device companies, collaborators, and the public. The FDA suggests that manufacturers explore different mechanisms for being transparent, including email, letters, and software notifications, and consider the amount of information to provide. Real-world performance monitoring may be achieved through reports to the FDA, and reporting may depend on the device, number of modifications, and maturity of the algorithm.

The proposed framework outlined in the FDA Paper describes a predetermined change control plan in premarket submissions that includes types of anticipated modifications and the associated methodology to implement those changes. The approach requires a commitment from manufacturers on transparency and real-world performance. This framework would allow the FDA to oversee AI that is constantly evolving and improving in a safe and efficient manner. The FDA Paper sought feedback through June 3, 2019. No date has been announced for further steps by the FDA.

Digital Health Guidance Documents

In October 2019, the FDA released new guidance that reflects small changes to prior guidance and provides clarity on the FDA's thought process with respect to SaMD. While the guidance does not alter the FDA's prior positions, it provides a more detailed analysis of risk factors that FDA will use to determine whether AI-driven CDS is a device subject to FDA regulations. Specifically, the Clinical Decision Support Software—Draft Guidance for Industry and Food and Drug Administration Staff (Draft Guidance) distinguishes software functions as either "Device CDS" or "Non-Device CDS" based on certain criteria. CDS that permits a health care provider to independently review the basis of the output recommended by the CDS and also meets the criteria of Section 520(o)(1)(E) of the Federal Food, Drug, and Cosmetic Act will be considered a "Non-Device CDS," and in most cases either will be not subject to FDA regulatory oversight or not subject to enforcement at this time. The Draft Guidance also provides that, currently, certain Device CDS also will not be subject to enforcement.[9] The determination of whether something is subject to regulatory enforcement, as described in the Draft Guidance, depends on whether it is intended to be used by a health care provider or by a patient/caregiver, whether the health care provider can independently review the conclusion or recommendation of the CDS, and if not, whether the condition or circumstances being evaluated by the CDS is significantly serious. The Draft Guidance included the helpful table (located on the following page):

> When the FDA finalizes its regulatory approach, stakeholders will be able to more easily identify and resolve challenges and share the benefits of AI with health care providers and consumers.

Summary of Regulatory Policy for CDS Software Functions			
		Intended User Is Health Care Practitioner	Intended User Is Patient or Caregiver
IMDRF Risk Categorization	Can the User Independent Review the Basis?	FDA Regulation	FDA Regulation
Inform X Critical	Yes	Not a Device	Not a Device
	No	Oversight Focus	Oversight Focus
Inform X Serious	Yes	Not a Device	Oversight Focus
	No	Oversight Focus	Oversight Focus
Inform X Non-Serious	Yes	Not a Device	Enforcement Discretion**
	No	Enforcement Discretion**	Oversight Focus

**"Enforcement Discretion" indicates that, based on its current understanding of the risks of these devices, FDA does not intend at this time to enforce compliance with applicable device requirements.

The Draft Guidance was open for comment through December 26, 2019.

Conclusion

The FDA's proposed framework, while vague in some areas, reflects a predetermined control plan in premarket submissions that would offer manufacturers of medical devices with AI capabilities a more flexible and practical approach for complying with FDA approval. The plan would require manufactures to describe anticipated modifications and the methodology for implementing them in a controlled manner that manages risk. The approach would also require manufacturers to commit to transparency, reporting, and real-world monitoring. When the FDA finalizes its regulatory approach, stakeholders will be able to more easily identify and resolve challenges and share the benefits of AI with health care providers and consumers.

Any views and opinions expressed in this article are those of the authors alone and should not be attributed to AHLA.

Carolyn V. Metnick is a Partner at McDermott Will & Emery LLP in the firm's Chicago office. She focuses her practice on health information privacy and security. She helps covered entities and business associates develop effective policies, procedures, and workforce training to comply with HIPAA and the HITECH Act. She also counsels businesses in data breach investigations and compliance with other federal and state breach notification laws, and assists businesses in responding to cyberattacks and other security breaches. Carolyn also advises health care providers on transactions and business issues. She guides clients through joint ventures, mergers and acquisitions, and counsels on governance matters as well as regulatory issues. She advises on compliance with the federal fraud and abuse laws, including the Stark Law and the Anti-Kickback Statute. Carolyn's background as a former litigator helps inform her transactional work.

Dina B. Ross is the sole proprietor of Dina B. Ross Law Offices. She focuses her practice on advising and assisting clients in meeting their technology legal needs in a variety of industries, including health care. She is experienced in drafting and negotiating commercial technology agreements in both established and emerging technologies, for both established and emerging companies. Ms. Ross also assists clients in navigating the technology legal issues during mergers, acquisitions, divestitures and other corporate restructurings with a particular emphasis in helping clients negotiate transition services agreements.

The authors thank Gregory Fosheim, Associate, McDermott Will & Emery LLP, for his insightful comments to and editing of this article.

Endnotes

1 *See* International Med. Device Regulators Forum (IMDRF), Software as a Medical Device (SaMD): Key Definitions (Dec. 2013), http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf.

2 Section 3060(a) of the 21st Century Cures Act, Pub. L. No. 114-255 (emphasis added).

3 *Id.*

4 *See* Food and Drug Admin. (FDA), Deciding When to Submit a 510(k) for a Software Change to an Existing Device, https://www.fda.gov/media/99785/download.

5 FDA Statement, April 2, 2019.

6 Developing a Software Precertification Program: A Working Model, v1.0—January 2019.

7 Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning Based Software as Medical Device.

8 *Id.*

9 The Draft Guidance proposes that certain Device CDS would not be subject to enforcement of regulatory compliance at the current time, based on the FDA's understanding of the risks of such Device CDS. For example, software intended to "inform clinical management" for non-serious conditions, such as a Device CDS that provides a health care provider with options for different diagnoses (e.g., seasonal allergic rhinitis versus common cold) would not be subject to enforcement currently.

For upcoming webinars on AI and information about the AHLA Convener on AI and Health Law, please see the inside back cover of this issue.