Featured Article

AI Health Care Liability: From Research Trials to Court Trials

Gary E. Marchant and Lucille M. Tournas

What is the issue? Artificial intelligence is about to disrupt clinical medicine—how will this new technology impact medical liability?

What is at stake? Historically, new medical technologies generate new waves of medical liability. Artificial intelligence is likely to have a similar effect in creating new liability risks and exposures for a variety of health care actors.

What do you need to know? Machine learning has given artificial intelligence systems new types of data-based capabilities that will greatly improve medical diagnosis, treatment, and administration, but will also raise difficult issues in attributing and evaluating responsibility for any resulting injuries.

Gary E. Marchant & Lucille M. Tournas, *AI Health Care Liability: From Research Trials to Court Trials*, J. Health & Life Sci. L., Feb. 2019, at 23.

Marchant and Tournas: AI Health Care Liability

CONTENTS

Introduction

There are two key points to understand about the role of artificial intelligence (AI) in health care. First, AI is a powerful transformative technology that is disrupting many industries, and it is already starting to have a significant impact on the practice of medicine. Second, and not notwithstanding point one, there is enormous hype and exaggeration about the actual extent and pace of AI's demonstrated value in changing health care practice at this time. AI is not going to cure cancer or displace doctors anytime soon, but it is showing promising results in the diagnosis, prognosis, and treatment of dozens of health conditions. Both of these points, especially when considered together, create a significant and growing liability exposure for health care providers, product manufacturers, and health care institutions, given the uncertainties created about whether, how, and when AI should be measured, if at all, by traditional notions of the standard of care.

This article addresses the potential liability implications of AI in the practice of medicine, and begins with a brief section on what AI is, provides some examples of how it is being applied in experimental trials, and then addresses the critical need and role of data in health care AI. The second half of this article discusses the liability implications of using AI, including the difficulties created by the rapid technological pace of AI, the question of who will be liable for injuries resulting from the inappropriate use or failure to use AI, and finally the implications of AI in litigating the standard of care.

Two key themes, which remain somewhat unresolved, emerge from our analysis. The first is that the regulatory status of AI in health care will have significant implications for liability risks and exposure. AI in health care has the potential to impact a large range of different tools and applications in the delivery of health care, ranging from medical procedures and services that are not regulated by the Food and Drug Administration (FDA), to products that are regulated by the FDA such as medical devices or software. In addition, the approach that the FDA is applying to AI and other advanced medical technologies continues to rapidly evolve and still remains somewhat uncertain, which in turn creates uncertainty about how possible regulatory approaches will affect AI liability risks and preemption of state malpractice liability.

Second, medical AI is unique in that, for the first time, decisions made by a machine could displace the professional judgment of a trained health care provider in certain AI applications. This raises the question of whether the potential liability of such an AI system should be evaluated under the traditional product liability standards for manufactured products, or whether professional malpractice liability based on negligence should be the applicable doctrine—just as it is for the physicians whose judgment may be replaced by the AI. No courts have yet had the opportunity to address this question, but it is likely to be central to the future liability questions in the era of AI in health care, and we analyze this issue below drawing on the closest available precedents.

What is Artificial Intelligence?

Artificial intelligence can be broadly defined as the capability of a machine to perform cognitive functions that previously could only be performed by a human brain. Within this broad definition are several different kinds of AI. Much of the public discussion has focused on the possibility of artificial general intelligence (AGI) that resembles or surpasses the human brain in its entirety, but AGI is at least several decades away. Rather what we have for the immediate future is more narrow AI systems that are capable of only performing a specific task or set of tasks.

Narrowly focused, AI has achieved tremendous progress in the past few years, including in the medical realm, due to a relatively new approach known as machine learning. Earlier AI systems consisted of pre-programmed responses to specific events—a human programmer created a decision tree based on various contingencies or inputs and instructed the machine on how to respond to each possible input. Such an approach is referred to as a rules-based approach. Simply put, the machine applies a set of rules written by a human programmer in response to inputs. Machine learning implements a data-based rather than rules-based approach. A machine learning system teaches itself how to respond based on patterns detected in large data sets. No human tells the machine what to do, but in most cases a human supervises the training of the machine learning system by telling the machine whether it made a right or

wrong decision in response to a particular input in the training period of the system. A more advanced form of machine learning is known as deep learning, which uses more complex "neural networks" that try to emulate the multi-step analytic pattern of the human brain to learn from and apply data.[1]

To take a simple example, a machine learning system may be given a large set of chest x-rays and asked to determine which of the scans indicate the presence of a lung tumor. An expert radiologist may tell the machine whether it gives a right or wrong answer to each scan. At first, the machine learning system will likely perform at or near the level of chance, but will gradually improve and eventually, with enough examples, become as proficient as the best radiologists. At some point, the machine learning system may even surpass human capabilities and correctly predict the presence of tumors—as revealed by subsequent patient outcomes—that even the expert radiologist missed.

Machine learning is now being applied in hundreds of applications in the health care industry and other industry sectors. Two significant implications of machine learning AI are important to note. Since machine learning systems are data-based, they require large sets of data from which to learn. This is particularly challenging in the health sector, as elaborated below. The second notable implication is that how an AI system integrates data to identify patterns and make predictions is often impervious to human understanding.[2] In the health care context, this results in something known as "black-box medicine," where machine learning systems can be applied to provide valuable information and diagnoses without human understanding of the inner workings of the system.[3] In response to this dilemma, AI developers are trying to improve the transparency of machine learning systems. For example, Google recently published a study in which its AI system pinpoints the specific

1 Geoffrey Hinton, *Deep Learning—A Technology with the Potential to Transform Health Care*, 320 J. Am. Med. Ass'n 1101 (2018).
2 Will Knight, *The Dark Secret at the Heart of AI*, MIT Tech. Rev., Apr. 11, 2017, *available at* www.technologyreview.com/s/604087/the-dark-secret-at-the-heart-of-ai/.
3 W. Nicholson Price II, *Black-Box Medicine*, 28 Harv. J. L. & Tech. 419 (2015).

portions of CT scans it relied on to diagnose various eye diseases.[4] However, for the immediate future, such efforts to penetrate the black box represent the exception rather than the rule, and most initial clinical applications of AI systems will result in such black boxes in which the system's decision-making methodology will not be auditable by humans.

Even when acting as a black box, AI systems will have the potential to provide tremendous value to the health care system in terms of more accurate and quicker diagnoses, treatment recommendations that are better informed by all relevant information, and cost savings associated with these improvements in disease diagnosis, prognosis, and treatment.[5] Yet, these AI black boxes are prone to unrecognized biases in the underlying data and reasoning errors that a human would quickly detect. For example, one hospital system used an AI machine learning system to determine which pneumonia patients could best be treated at home rather than admitted to the hospital.[6] The AI recommended that pneumonia patients with asthma were one group of patients that could be safely treated at home, based on the good average outcome of such patients treated at the hospital. However, the better than average outcome for these patients was an artifact of the fact that such pneumonia patients with asthma are immediately admitted to acute care units because they are at greater rather than less risk than most pneumonia patients. A human doctor would know this immediately, whereas the AI system lacked the common sense and tacit knowledge of a human physician. This type of error demonstrates the need for physicians to oversee the application of AI systems in clinical practice involving real patients.[7] For the moment, AI is not ready to displace physicians or "practice medicine" as that term is defined by the 50 states.

4 Jeffrey De Fauw et al., *Clinically Applicable Deep Learning for Diagnosis and Referral in Retinal Disease*, 24 Nature Med. 1342 (2018).

5 Monique Brouillette, *Deep Learning Is a Black Box, but Health Care Won't Mind*, MIT Tech. Rev., Apr. 27, 2017, *available at* www.technologyreview.com/s/604271/deep-learning-is-a-black-box-but-health-care-wont-mind/.

6 Rich Caruana et al., *Intelligible Models for Healthcare: Predicting Pneumonia Risk and Hospital 30-Day Readmission*, *in* KD '15 Proceedings of the 21st ACM SIGKDD International Conference on Knowledge Discovery and Data Mining 1721 (2015), *available at* http://people.dbmi.columbia.edu/noemie/papers/15kdd.pdf.

7 *See generally* Kate Crawford & Ryan Calo, *There is a Blind Spot in AI Research*, Nature, Oct. 13, 2016, at 311.

Examples of AI in health care

Health care is one of the industries in which AI will have the most significant and rapid impact. The growing reliance on big data, the current inefficiencies built into the system, the complexity of health problems, and the potential for major improvements in clinical outcomes all make health care prone to improvement and disruption by AI. Major technology companies, such as IBM, Microsoft, Apple, Amazon, and Google, as well as over 100 AI start-ups, are already racing towards applying AI to the health care sector. Many of the initial applications of AI in health care do not directly involve clinical care directly, but rather involve functions such as scheduling appointments, billing, and medical record administration, which are unlikely to raise significant liability issues. For this reason, these applications of AI are not further addressed in this article. AI is also being used in drug discovery and in improving clinical trials to look for subtle patterns in large data sets, among other non-clinical uses.

The introduction of AI into clinical care has begun and, based on early successes, will only continue to increase. Radiology is one of the most noteworthy applications so far. To start, there is high demand on radiologists—in some regions a clinician would need to read an image every 3–4 seconds in an eight-hour shift to meet demand.[8] AI can help alleviate such burdens. Additionally, AI is particularly skilled in this specialty as it involves narrow, task-specific diagnoses. For example, researchers at Stanford developed an algorithm that is able to detect pneumonia from a chest X-ray at a level that exceeds practicing radiologists.[9] A research team led by Google was able to detect diabetic retinopathy from retinal photographs using machine learning at a sensitivity equal to or greater than ophthalmologists.[10] Additionally, there is also the opportunity for interpretation tasks that radiologists do not cur-

8 Ahmed Hosny et al., *Artificial Intelligence in Radiology*, 18 Nature Reviews Cancer 500 (2018).
9 Pranav Rajpurkar et al., CheXNet: Radiologist-level Pneumonia Detection on Chest X-Rays with Deep Learning (2017), *available at* https://arxiv.org/pdf/1711.05225.pdf.
10 Varun Gulshan et al., *Development and Validation of a Deep Learning Algorithm for Detection of Diabetic Retinopathy in Retinal Fundus Photographs*, 316 J. Am. Med. Ass'n 2402 (2016).

rently perform, such as radiogenomics, which focuses on the relationship between imaging features of tumors and their genomic features.[11] This could prove to be groundbreaking work towards directing oncology treatment plans.

AI applications such as these have led many in the field to believe AI may eventually displace many of the tasks currently performed by radiologists. Radiologist Robert Schier, MD, wrote, "The advent of computers that can accurately interpret diagnostic imaging studies will upend the practice of radiology…[t]he two currently unanswered questions are just how much upending there will be and how long it will take to happen."[12] While some tasks currently performed by radiologists can be automated, others cannot. Therefore, an important distinction is to be made between those radiographic analyses that can be automated and those that cannot, all while realizing that this line may shift with time and further technological advances.[13]

IBM Watson Oncology offers clinicians AI processing software that is equipped with guidelines, best practices, medical journals, and textbooks.[14] It can then assess information from medical records and weigh medical evidence in order to indicate potential treatment options for clinicians complete with a confidence score and supporting evidence. As Watson needs large amounts of patient data to learn and improve, its implementation has been slower than originally projected and it is currently being used only to supplement physician knowledge rather than replace it.[15]

A research group at Mount Sinai Hospital in New York applied machine learning to the hospital's massive database of patient records. This program, named Deep Patient, was trained using data from about 700,000 patients, and

11 Maciej A. Mazurowski et al., Deep Learning in Radiology: An Overview of the Concepts and a Survey of the State of the Art, (2018), *available at* https://arxiv.org/pdf/1802.08717.pdf.

12 *Quoted in* Michael Walter, *If You Think AI Will Never Replace Radiologists—You May Want to Think Again,* Radiology Bus., May 14, 2018, *available at* www.radiologybusiness.com/topics/artificial-intelligence/if-you-think-ai-will-never-replace-radiologists-you-may-want-think.

13 Saurabh Jha & Eric J. Topol, *Adapting to Artificial Intelligence: Radiologists and Pathologists as Information Specialists*, 316 J. Am. Med. Ass'n 2353 (2016).

14 IBM, *IBM Watson for Oncology*, www.ibm.com/us-en/marketplace/ibm-watson-for-oncology.

15 Casey Ross & Ike Swetlitz, *IBM Pitched its Watson Supercomputer as a Revolution in Cancer Care. It's Nowhere Close*, STAT, Sept. 5, 2017, *available at* www.statnews.com/2017/09/05/watson-ibm-cancer/.

when tested on new records was particularly capable at predicting disease.[16] It discovered patterns hidden in hospital data, including patterns of liver cancer as well as schizophrenia—which is particularly difficult for physicians to detect. Google also used an AI system to mine electronic health records and developed an AI algorithm that is able to predict a variety of clinical events—including a patient's death—better than traditional, clinically-used predictive models.[17]

Mobile health is another area in which AI is already starting to make large impacts. For example, ResApp Health released its ResAppDx application, which has shown the potential to accurately diagnosis croup, pneumonia, and other respiratory conditions based on the audio of an individual's cough.[18] This has proven beneficial for doctors to use in telemedicine, it allows consumers access to diagnostic tools before going to their physician, and it can provide access to respiratory diagnoses in the developing world where access to physicians is poor.

Intel's IQ also released an application, named Doctor Hazel, which uses a picture of a mole taken by a user to sort through a database of images and classify said mole as to its risk, letting the patient know if the mole looks benign or may be potentially cancerous.[19] If the application predicts a risk of cancer, it advises the patient to follow up with a dermatologist. Skin cancer is the most common cancer in the U.S. and melanoma is particularly fatal but has a high cure rate if caught early. Initially, the company sees opportunity for use by primary care physicians but aims eventually to put it in the hands of consumers.

While most of the AI applications are still in the experimental stage, these and numerous other advances in using AI to diagnose and treat patients signal that the evolution of AI in health care is likely at a tipping point and may soon be readily available to clinicians.

16 Riccardo Miotto et al., *Deep Patient: An Unsupervised Representation to Predict the Future of Patients From the Electronic Health Records,* 6 Sci. Rep. 26094 (2016).

17 Alvin Rajkomar et al., *Scalable and Accurate Deep Learning with Electronic Health Records*, 1 NPJ Digital Med., no. 18, 2018.

18 *Respiratory Disease Diagnosis Using Only the Sound of a Patient's Cough*, ResApp Health, *available at* www.resapphealth.com.au/technology/.

19 Joyce Riha Linik, *Skin Cancer Detection Using Artificial Intelligence*, Intel IQ, Jan. 31, 2018, *available at* https://iq.intel.com/skin-cancer-detection-using-artificial-intelligence/.

The regulatory status of AI in health care

The regulatory status of AI in health care will be an important determinant of the liability doctrines that apply to AI systems if harms occur, as well as the commercialization pathways and speed with which AI enters our health care system. However, the regulatory status of medical AI is uncertain at this time. A key issue is whether an AI system is a medical device or a medical service or procedure. The traditional distinction is that the FDA regulates products such as medical devices, while states regulate the "practice of medicine."[20] There has been some blurring of this demarcation over time.[21] Moreover, this distinction is anything but clear for AI systems. The FDA has already approved several AI diagnostic devices as medical devices. For example, the agency announced in April 2018 that it had approved the machine-learning device described above for detecting diabetic retinopathy from retinal photographs without an eye specialist to validate the AI's scan interpretation.[22] On the other hand, other AI systems, such as the IBM Watson Health also described above are currently being used for clinical use without any FDA approval or oversight, apparently accepted as medical service or procedure not requiring FDA approval. The precise line between product and medical practice by AI systems remains to be defined.

Additional recent and pending change further complicate AI health regulation. In December 2017, the FDA issued a series of draft guidances and requested public comment on the regulation of medical software, in part to implement the 21st Century Cures Act.[23] These draft guidances raise significant questions about how and whether the FDA will regulate clinical support

20 United States v. Evers, 643 F.2d 1043, 1048 (5th Cir. 1981) (FDA "was obviously intended to control the availability of drugs for prescribing by physicians," but it "was not intended to regulate the practice of medicine ….").

21 Patricia J. Zettler, *Toward Coherent Federal Oversight of Medicine*, 52 San Diego L. Rev. 427 (2015).

22 Press Release, U.S. Food & Drug Admin., FDA Permits Marketing of Artificial Intelligence-Based Device to Detect Certain Diabetes-Related Eye Problems (Apr. 11, 2018), *available at* www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm604357.

23 *See, e.g.,* U.S. Food & Drug Admin., Clinical and Patient Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff (2017), *available at* www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/UCM587819.pdf.

software that relies on complex machine learning functionality. The FDA has also created a Digital Health Software Precertification (Pre-Cert) Program which seeks to regulate digital companies as a whole rather than specific products.[24] There are many uncertainties about what this program will require, how many AI companies it will apply to, and the FDA's authority to regulate companies in addition to devices.

In short, there are many moving parts, uncertainties, and complexities about how the FDA will regulate AI health systems. How these questions are resolved will have significant implications for potential AI liability. If an AI system is determined to be a medical device by the FDA and requires a pre-market authorization (PMA), the FDA's regulatory approval will generally preempt state product liability claims. On the other hand, if the AI system is a medical device but is regulated by some other mechanism such as Section 501(k) clearance, then no preemption will apply but the regulatory classification will suggest that product liability is the appropriate tort theory if the system causes injury. Finally, if the FDA determines that the AI system is not a medical device and is not subject to FDA regulation because it's use amounts to medical practice, such an outcome might suggest that medical malpractice rather than product liability is the appropriate tort remedy.

Liability Risks for AI in Health Care

AI will present some unique liability issues. The key dimensions include the rapid pace of AI development and deployment and questions about who the responsible party is, and by what standard of care they should be held responsible. In particular, we address whether an AI system that results in an injury to a patient should be treated as a product and subjected to the product liability doctrine, or as the practice of medicine subject to medical malpractice liability rather than product liability.

24 *Digital Health Software Precertification (Pre-Cert) Program*, U.S. Food & Drug Admin., (updated Sept. 27, 2018) *available at* www.fda.gov/MedicalDevices/DigitalHealth/DigitalHealthPreCertProgram/default.htm.

The rapid pace of technology development

Historical analyses of medical liability demonstrate that change in technology is one of the primary drivers of liability in the health care system.[25] The rapid pace at which AI health care applications develop will present difficult challenges for health care actors. The new emerging technology will quickly eclipse the training of most practitioners, resulting in many providers not using the best available technology or making mistakes in applying such AI systems for which they have not been properly trained. In addition, clinical guidelines and clinical support systems cannot keep pace with the rapidly evolving technologies, leaving providers uncertain on whether and how to apply newly available AI tools. Of course, when there is a new technology, some providers will utilize the technology while others will not, further exacerbating uncertainty and controversy over the applicable standard of care.[26]

The rapid development of AI medical technology creates a challenge for providers that we call the "tipping point" dilemma. As technologies change so rapidly, what might be malpractice if relied on today may be negligent to not use tomorrow. For example, as discussed above, there are several diagnostic modalities where preliminary studies have shown the AI system appears to out-perform experienced physicians. It could constitute malpractice for a physician to defer to such an unproven AI test at this time, as these applications are still experimental. But what type of validation is required for such systems before they are ready for prime time in clinical care? There is no clear agreement on this. Many researchers and policymakers would prefer to see multiple prospective controlled studies before adopting AI systems—or any medical technology—into mainstream clinical care. Yet, requiring such studies is not always followed for other medical technologies and procedures. Further, the tremendous benefits shown by some AI programs in initial studies can place pressure on clinicians to use these tools sooner rather than later. As these AI systems become more developed and start being applied in the clinic,

25 William M. Sage, *Medical Liability and Patient Safety*, Health Aff., July 2003, at 26,28; Kenneth De Ville, *Medical Malpractice in Twentieth Century United States. The Interaction of Technology, Law and Culture,* 14 Int'l J. Tech. Assessment Health Care 196, 200 (1998).

26 Michael D. Greenberg, *Medical Malpractice and New Devices: Defining an Elusive Standard of Care*, 19 Health Matrix 423 (2009).

society may reach a point where a jury could find that it was negligent for a physician not to use or rely on AI technology. When does this tipping point occur? There is no official notice or declaration to providers when this tipping point is reached, it will only be determined *ex post* by a jury. This puts providers in a very difficult situation.

There has been very little precedent on how the medical liability system deals with rapidly changing technologies and standard of care. The paradigmatic medical malpractice case involving a changing standard of care is *Burton v. Brooklyn Doctors Hospital*,[27] where the court held that the prevailing standard of care of giving pure oxygen to premature babies was no longer consistent with new evidence. The court stated that:

> [a]lthough the conventional medical wisdom at the time believed that increased oxygen was essential to the survival of premature babies, the hospital and Dr. Engle cannot avail themselves of the shield of acceptable medical practice when a number of studies, including their own, had already indicated that increased oxygen was both unnecessary and dangerous …[28]

With the unprecedented pace of technological change that AI is advancing, health care actors will not be able to rely on "accepted medical practice" but will need to continually monitor the latest developments to determine the most effective and safest treatment option. As one expert recently noted, "[t]he question for hospitals isn't 'who will be liable if we use AI' but rather 'how soon will we be liable if we don't?'"[29]

Who is the liable party?

AI will also introduce new complexity to the question of who is liable if the failure to use, or the inappropriate use of, an AI system results in harm to a

27 Burton v. Brooklyn Doctors. Hosp., 452 N.Y.S.2d 875 (N.Y. 1982).
28 *Id.* at 879–80.
29 Sally Daub, *Defensive Diagnostics: The Legal Implications of AI in Radiology,* AIMed, Aug. 2, 2018, *available at* http://ai-med.io/defensive-diagnostics-the-legal-implications-of-ai-in-radiology/.

patient. The potentially liable party could be one or more of the following: (i) the treating physician who recommended and supervised the AI system; (ii) the technician who administered or interpreted the AI system; (iii) the manufacturer who developed and marketed the AI system; and/or (iv) the health care institution in which the patient received care. Different criteria and standards of care will apply depending on who is the potentially liable party or parties.

If AI is deemed a medical product, commercial manufacturers of products such as AI systems will be subject to product liability for harm to persons or property caused by a defect associated with their products.[30] Product liability will be especially enticing to plaintiff attorneys, both because product manufacturers usually have deeper pockets than physicians, and also because product liability provides more favorable doctrines for plaintiffs, including strict liability and punitive damages. In contrast, medical malpractice lawsuits against physicians are adjudicated under a negligence standard, and damages are frequently limited.[31]

This dynamic can be seen in litigation involving robotic surgery. For example, *Taylor v. Intuitive Surgical* was a case involving patients that were injured during the course of procedures using Intuitive Surgical's da Vinci robotic surgical system.[32] In that case, according to both the trial and appellate courts, the issue was whether the surgeon operating the machine or the machine itself is responsible for the injury. Given the advantages for plaintiffs of a product liability suit against the manufacturer, the plaintiff attorneys tried to structure the lawsuit as a product liability suit against the manufacturer, even when it appears that the surgeon made an error. The facts of the case revealed that the surgeon was trained and credentialed to use the da Vinci device through the hospital. The da Vinci System provided a manual to the physician stating the maximum recommended BMI for robotic surgery is 30. Yet, the patient presented to the physician with a recorded BMI of 39. The surgeon acknowledged

30 Restatement (Third) of Torts: Product Liability §1 (Am. Law Inst. 1998).
31 Christopher L. Thompson, *Imposing Strict Products Liability on Medical Care Providers*, 60 Mo. L. Rev. 711, 715–21 (1995).
32 Taylor v. Intuitive Surgical, Inc., 355 P.3d 309 (Wash. Ct. App. 2015), *aff'g* Case No. 09-2-03136-5 (Wash. Superior Ct. June 7, 2013).

that the patient was not an ideal candidate for this type of surgery using a robotics system but nevertheless proceeded. The surgery was marred with complications, and the patient was left with a poor quality of life including respiratory issues requiring ventilator support, neuromuscular damage that left the patient unable to walk unaided, and incontinence. He passed away from complications four years after surgery. Despite these facts, the plaintiff still argued that the manufacturer was ultimately to blame, usually for failing to adequately train or warn the physician. The jury found for Intuitive Surgical, determining that the company had a duty to warn the physician regarding the nature of the robotics system and that the company did provide an adequate warning to the physician. The court of appeals affirmed, holding that Intuitive Surgical did not have a separate duty to warn the hospital.[33]

However, on further appeal the Washington State Supreme Court extended the company's duty to warn past the surgeon who was using the device, to the hospital.[34] The Washington State Supreme Court found that even though the company had trained and warned the physician, manufacturers also have a duty to warn to hospitals, as purchasing agents, about the dangers of their products. As such, the court vacated the lower courts' decisions in favor of the manufacturer and remanded for further proceedings.

As *Taylor v. Intuitive Surgical, Inc.* demonstrates, hospitals and other institutions offering medical care may also face increased liability risk with AI-enabled medicine. An individual physician practicing in a hospital—whether as an employee or an independent contractor—will rarely be in the position to purchase and use his or her own AI system. Rather, it will be the institution who selects, installs, trains, and operates an AI system that its physicians may utilize, and thus it is the institution that may be directly or vicariously liable for any faults, including a decision not to make the latest AI capabilities available, deficiencies in the installation or access to the AI system, failure to properly train staff on the AI system, and failure to provide informed consent to patients treated using the AI system. Licensing agreements between

33 *Id.*
34 Taylor v. Intuitive Surgical, Inc., 389 P.3d 517 (Wash. 2017).

AI vendors and health care organizations will play a critical role in assigning liability and providing indemnification as this new world of AI-assisted health care develops.[35]In the case of AI with machine learning systems, such systems require massive amounts of data, and storage to hold such data, in order for the system to learn and perfect itself in the long term. If this data is collected and used by only a very few large technology companies such as Google, Apple, or IBM, these will be very tempting liability targets for plaintiffs' attorneys if a patient is injured. To protect themselves against liability, these companies will need to carefully implement and document effective training and disclosure programs to hospitals, clinics, technicians, and physicians who will implement these AI technologies within patient care. Such liability exposure and safeguards slow down the clinical implementation of AI systems.

Standard of care

Medical malpractice has traditionally been based on negligence, which uses the reasonable physician standard. Simply put, the court looks to how other physicians, of similar specialty, would act in the same situation. This is often established through the testimony of expert witnesses. Unsupported deviation from these standards typically results in a conclusion that he or she has failed to provide the required level of care. In contrast, if a medical device or equipment is to blame, product liability—with its strict liability standard for manufacturing defects, design defects, and failure to warn defects—will apply.

AI disrupts this well-established order because AI systems may be deemed to replace doctors themselves.[36] If an AI machine is now displacing the professional judgment of a human physician, should it be judged by the same legal standard that the human physician is subjected to? In other words, should the AI machine be evaluated under the traditional medical malpractice negligence standard rather than product liability doctrine? If the machine is evaluated under a different standard than the human doctor who it replaces in perform-

35 Scott Bennett & Leeann Habte, *Artificial Intelligence in Health Care: Welcome to the Machine*, AHLA Connections, Dec. 2018, at 16, 18–19.

36 Jane R. Bambauer, *Dr. Robot*, 51 U.C. Davis L. Rev. 383 (2017).

ing a specific task, this discrepancy may bias the outcome of a head-to-head competition between human and machine. Imposing a higher standard on the AI machine may deprive patients of better care and could deprive the health system of potential cost savings provided by an AI system. On the other hand, imposing a lower standard of are on the AI machine may encourage the offering of substandard care.

In solving these problems, it will be important to determine whether the AI system is displacing a physician function or is simply assisting the physician in the care of the patient. Like a physician, when the machine errors in its learned judgment based on machine learning principles, it should be held to medical malpractice liability standards and potentially even regulated by state Boards of Medicine. But when some other defect in the manufacture, design, or warnings of the machine are at fault, traditional product liability should apply. For example, if IBM Watson diagnoses a type of cancer and establishes a treatment plan that turns out to be wrong, and the patient is harmed as a result, those decisions should be analyzed under negligence. In this instance, the machine isn't liable for every injury, but only for those in which the standard of care falls below that of a human physician. This allows for patients to be appropriately compensated when the standard of care is not met and regardless of whether a human or machine made the professional judgement about the patient's care. On the other hand, if IBM Watson has a poorly constructed part that causes injury, IBM would be evaluated under a product liability standard and therefore be strictly liable.

Applying a medical malpractice standard to machine learning AI systems would raise a number of subsidiary issues. Most jurisdictions apply a standard of care in which a physician's actions are compared to specialist in the same field. What would be the applicable comparison for an AI system? Given that medical malpractice applies because the AI is substituting for a physician, it seems obvious that the AI system should be compared to human physicians in the same specialty and jurisdiction when appropriate. As AI systems continue to develop, they may far surpass human capabilities in certain tasks. Should AI capabilities be considered in determining whether a particular AI system or

physician performance fell below the standard of care? *Brooklyn Doctors Hospital* suggests it should. If the standard of care for an AI system is based on the performance of both human and AI systems, so too should the medical malpractice standards for physicians. This means as AI systems exceed human capabilities in specific tasks, physicians will risk malpractice liability if they continue to try to undertake those tasks themselves. In this way, over time, more and more medical tasks are likely to become completely taken over by medical AI systems, displacing physicians from such tasks.

Another issue is whether AI systems will influence the choice between a community-based versus national standard of care for medical malpractice. In recent years, with the implementation of national standards for medical education and the universal availability of online guidance and evidence, there has been an argument and growing support for a shift to a national standard of care.[37] AI systems will likely be expensive and, thus, limited in access at least initially, as high-end academic medical centers and elite clinics lead the development and uptake of AI in clinical care. At this stage, it may be unfair to apply the same standard to providers at these elite institutions as to providers in institutions without access to the cutting-edge AI systems. However, as the price of these AI systems decline and their availability expands, the case for applying a national standard of care will be stronger than ever.

Finally, when a physician's error harms a patient, that usually only involves that specific incident, and affects only the one patient. While it is possible that the same physician could make a similar mistake with another patient in the future, patient care and negligence associated with such care, generally occurs on a patient-by-patient basis. In contrast, a faulty algorithm in a machine learning system could potentially affect many patients in short period of time.[38] This might suggest a higher standard of care of machine learning systems, given the increased harm that can result from an error.

37 Michelle Huckaby Lewis et al., *The Locality Rule and the Physician's Dilemma,* 297 J. Am. Med. Ass'n 2633, 2634 (2007).

38 Nilay D. Shah et al., *Big Data and Predictive Analytics: Recalibrating Expectations,* 320 J. Am. Med. Ass'n 27, 27 (2018).

Conclusion

Artificial Intelligence will rapidly disrupt the practice of medicine. Even if AI improves both the quality and cost of health care as expected, it will create new liability issues, just as previous technologies have done. Law has very little experience in evaluating machines that express some form of intelligence that previously could only be exhibited by humans. It remains to be seen how courts will treat the mistakes that AI systems will make, and whether liability will promote or impede improved patient safety. J

Gary E. Marchant received his Ph.D. in genetics from the U. of British Columbia, his J.D. from Harvard Law School (1990), and his M.P.P. from the Kennedy School of Government, Harvard University (1990). He serves as the Regents' Professor and Lincoln Professor of Emerging Technologies, Law & Ethics, and Faculty Director of the Center for Law, Science & Innovation, at the Sandra Day O'Connor College of Law, Arizona State University (ASU). Professor Marchant's research and teaching interests include the governance of emerging technologies, legal aspects of personalized medicine, artificial intelligence and the law, blockchain and the law, use of genetic information in the legal system, legal aspects of risk assessment and risk management, and the application of science and technology in the legal system. Contact him via email at gary.marchant@asu.edu.

Lucille M. Tournas is a research fellow in the School for the Future of Innovation in Society at Arizona State University, as well as research fellow in the Center for Law, Science and Innovation at Arizona State University's College of Law. Tournas earned her J.D. (2018) from Arizona State, and her B.A. (Philosophy) (2004) from Pepperdine University. Tournas' current research interests include the regulation of emerging technologies, specifically the role of artificial intelligence and big data in health care. Contact her via email at lucy.tournas@gmail.com.