VIEWPOINT

W. Nicholson Price II, JD, PhD
University of Michigan Law School, Ann Arbor.

Sara Gerke, Dipl-Jur Univ
Project on Precision Medicine, Artificial Intelligence, and the Law, Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics, Harvard Law School, Harvard University, Cambridge, Massachusetts.

I. Glenn Cohen, JD
Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics, Harvard Law School, Harvard University, Cambridge, Massachusetts.

Potential Liability for Physicians Using Artificial Intelligence

Artificial intelligence (AI) is quickly making inroads into medical practice, especially in forms that rely on machine learning, with a mix of hope and hype.[1] Multiple AI-based products have now been approved or cleared by the US Food and Drug Administration (FDA), and health systems and hospitals are increasingly deploying AI-based systems.[2] For example, medical AI can support clinical decisions, such as recommending drugs or dosages or interpreting radiological images.[2] One key difference from most traditional clinical decision support software is that some medical AI may communicate results or recommendations to the care team without being able to communicate the underlying reasons for those results.[3]

Medical AI may be trained in inappropriate environments, using imperfect techniques, or on incomplete data. Even when algorithms are trained as well as possible, they may, for example, miss a tumor in a radiological image or suggest the incorrect dose for a drug or an inappropriate drug. Sometimes, patients will be injured as a result. In this Viewpoint, we discuss when a physician could likely be held liable under current law when using medical AI.

Medical AI and Liability for Physicians

In general, to avoid medical malpractice liability, physicians must provide care at the level of a competent physician within the same specialty, taking into account available resources.[4] The situation becomes more complicated when an AI algorithmic recommendation becomes involved. In part because AI is so new to clini-

---

The challenge is that current law incentivizes physicians to minimize the potential value of AI.

---

cal practice, there is essentially no case law on liability involving medical AI. Nonetheless, it is possible to understand how current law may be likely to treat these situations from more general tort law principles.

The **Figure** presents potential outcomes for a simple interaction—for instance, when an AI recommends the drug and dosage for a patient with ovarian cancer. Assume the standard of care for this patient would be to administer 15 mg/kg every 3 weeks of the chemotherapeutic bevacizumab.

The first question (column 2) is what the AI recommends. To simplify, assume the AI makes 1 of 2 recommendations (there could be many more in clinical care): the standard-of-care dosage or a much higher dosage: 75 mg/kg every 3 weeks.

The AI could be correct or incorrect either way (column 3). Even though 15 mg/kg every 3 weeks is the standard-of-care dosage, perhaps for some reason the higher dosage is right for this particular patient.

Such a recommendation would be consistent with one of the goals of some AI, to personalize care.

Next, the physician could either follow or reject the AI recommendation (column 4). In this example, the physician retains this discretion, although in the future a health system or a payer may limit physician discretion.

Column 5 represents the patient outcome. If the physician follows a correct recommendation or rejects an incorrect recommendation, the outcome is good; if the opposite, the outcome is bad (in this very stylized example).

Eight possible scenarios result (column 6). The law treats them differently, and understanding these differences rests on some basic elements of US tort law. First, if there is no injury, there will be no liability (green boxes in the Figure); this is a good outcome, whether it happens because the physician accepts a correct recommendation (scenarios 1 and 5) or rejects an incorrect recommendation (scenarios 4 and 8).

Second, tort law typically privileges the standard of care, regardless of its effectiveness in a particular case—ie, whether providing that care leads to a good or bad outcome. When physicians follow the standard of care (eg, 15 mg/kg of bevacizumab every 3 weeks; scenarios 1, 3, 6, and 8), they will not generally be held liable for a bad outcome, even if a different course of action would have been better for a particular patient in a particular case (yellow boxes in the Figure).

Thus, under current law, a physician faces liability only when she or he does not follow the standard of care *and* an injury results (red boxes in the Figure).

This analysis suggests an important implication for physicians using medical AI to aid their clinical decisions: because current law shields physicians from liability as long as they follow the standard of care, the "safest" way to use medical AI from a liability perspective is as a confirmatory tool to support existing decision-making processes, rather than as a source of ways to improve care.

Although many physicians may be comfortable with this approach, the challenge is that current law incentivizes physicians to minimize the potential value of AI. If the medical AI performs a task better than the physician, such as recommending a higher dosage of a drug, it will provide some results different than the physician. The difference will increase if, in the future, some medical AIs perform better than even the best physicians, a goal for some algorithms. But because threat of liability encourages physicians to meet and follow the standard of care, they may reject such recommendations and thus fail to realize the full value of AI, in some cases to patients' detriment.

The legal standard of care is key to liability for medical AI, but it is not forever fixed. Over time, the standard of care may shift. What happens if medical practice reaches

Corresponding Author: W. Nicholson Price II, JD, PhD, University of Michigan Law School, 625 S State St, Ann Arbor, MI 48109 (wnp@umich.edu).

Figure. Examples of Potential Legal Outcomes Related to AI Use in Clinical Practice

Scenario	AI recommendation	AI accuracy	Physician action	Patient outcome	Legal outcome (probable)
1	Standard of care	Correct	Follows	Good	No injury and no liability
2			Rejects	Bad	Injury and liability
3		Incorrect (standard of care is incorrect)	Follows	Bad	Injury but no liability
4			Rejects	Good	No injury and no liability
5	Nonstandard care	Correct (standard of care is incorrect)	Follows	Good	No injury and no liability
6			Rejects	Bad	*Injury but no liability*
7		Incorrect	Follows	Bad	*Injury and liability*
8			Rejects	Good	No injury and no liability

AI indicates artificial intelligence.

a point where AI *becomes* part of the standard of care, the consensus view of good medical practice?[4] If and when that happens, scenarios 6 and 7 (italicized text in the Figure) may change substantially: physicians may incur liability for *rejecting* correct but nonstandard AI recommendations and may conversely avoid liability for injury if they were *following* incorrect AI recommendations. Because tort law is inherently conservative, the second alternative (scenario 7) is a more likely first step: reliance on medical AI to deviate from the otherwise known standard of care will likely be a defense to liability well before physicians are held liable for rejecting AI recommendations. But physicians should watch this space because it may change quickly.

What Should Physicians Do?

Physicians have a substantial role in shaping the liability issue. In their practices, physicians should learn how to better use and interpret AI algorithms, including in what situations an available medical AI should be applied and how much confidence should be placed in an algorithmic recommendation. This is a challenge, and evaluation tools are still very much under development.

Physicians should also encourage their professional organizations to take active steps to evaluate practice-specific algorithms. Review by the FDA will provide some quality assurance, but societies will be well placed to provide additional guidelines to evaluate AI products at implementation and to evaluate AI recommendations for individual patients. The analogy to practice guidelines is strong; much as societies guide the standard of care for specific interventions, they can guide practices for adopting and using medical AI reliably, safely, and effectively.

As part of care settings, such as in hospitals and health systems, physicians should also ensure that administrative efforts to develop and deploy algorithms reflect what is truly needed in clinical care. When external AI products are procured, physicians should advocate for safeguards to ensure that such products are rigorously vetted before procurement, just as with other novel medical devices.

In addition, physicians should check carefully with their malpractice insurer to determine how the insurer covers the use of medical AI in practice. Is care that relies on AI recommendations covered the same as care without such recommendations, or does the insurer treat such practices differently? Are practices different for more opaque algorithms that provide little or no reasoning? Collectively, physicians and their hospital systems may be able to make demands for changes in terms of insurance coverage to better accommodate the needs of a future of AI-enabled medicine.

Although current law around physician liability and medical AI is complex, the problem becomes far more complex with the recognition that physician liability is only one piece of a larger ecosystem of liability. Hospital systems that purchase and implement medical AI, makers of medical AI, and potentially even payers could all face liability.[5] The scenarios outlined in the Figure and the fundamental questions highlighted here recur and interact for each of these forms of liability. Moreover, the law may change; in addition to AI becoming the standard of care, which may happen through ordinary legal evolution, legislatures could impose very different rules, such as a no-fault system like the one that currently compensates individuals who have vaccine injuries.

As AI enters medical practice, physicians need to know how law will assign liability for injuries that arise from interaction between algorithms and practitioners. These issues are likely to arise sooner rather than later.

REFERENCES

1. Emanuel EJ, Wachter RM. Artificial intelligence in health care: will the value match the hype? *JAMA*. 2019;321(23):2281-2282. doi:10.1001/jama.2019.4914

2. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. doi:10.1038/s41591-018-0300-7

3. Burrell J. How the machine "thinks": understanding opacity in machine learning algorithms. *Big Data Soc*. 2016;3:2053951715622512. doi:10.1177/2053951715622512

4. Froomkin AM, Kerr I, Pineau J. When AIs outperform doctors: confronting the challenges of a tort-induced over-reliance on machine learning. *Ariz Law Rev*. 2019;61:33-99.

5. Price WN. Medical malpractice and black-box medicine. In: Cohen IG, Fernandez Lynch H, Vayena E, Gasser U, eds. *Big Data, Health Law, and Bioethics*. Cambridge, England: Cambridge University Press; 2018:295-306. doi:10.1017/9781108147972.027