

## Medical AI and tort liability

I. Glenn Cohen<sup>a,b</sup>, Andrew Slottje<sup>c</sup>, and Sara Gerke<sup>d</sup>

<sup>a</sup> James A. Attwood and Leslie Williams Professor of Law, Harvard Law School, Cambridge, MA, United States

<sup>b</sup> Petrie-Flom Center for Health Law Policy, Biotechnology & Bioethics, Harvard Law School, Cambridge, MA, United States

<sup>c</sup> J.D., Class of 2023, Harvard Law School, Cambridge, MA, United States

<sup>d</sup> Penn State Dickinson Law, Carlisle, PA, United States

### Introduction

This chapter examines the role of liability in shaping the use of artificial intelligence (AI) in medicine. It begins with a hypothetical vignette, then examines various forms of liability—physician medical malpractice, lack of informed consent, corporate liability for hospitals, and developer liability. Finally, it turns to the preemption of liability and regulation. While many of the points raised apply across legal systems, this chapter focuses primarily on US law, with some reference to EU law.

#### A starting vignette

Imagine an oncologist providing cancer treatment to a patient in a hospital somewhere in the United States. Imagine that all else equal, the physician in question would administer paclitaxel to the patient at an established dose of 175 mg/m<sup>2</sup> body surface area. Now, imagine instead that our physician has access to an AI system developed by the hospital with a software company. Recently cleared by the Food and Drug Administration (FDA), the system recommends a dose that varies over time, based in part on the patient's changing biomarker levels, as measured by hospital employees. While the hospital (and indeed the insurer's guidelines) promotes the use of the AI system, there is no requirement to use it. Indeed, even when using it, the final decision on the correct dosage remains the physician's to make. The physician ultimately chooses to use the medical AI system in setting the dose for this patient. The AI system recommends a dose higher than the established standard of care, and the physician follows the recommendation. The patient experiences side effects that we have reason to believe are the result of the elevated dose, causing physical, emotional, and financial injury [1].

Under the current law, against whom, if anyone, may the patient successfully bring a lawsuit? The physician, the hospital, or the AI developer? If we were system designers able to freely alter the prevailing rules to govern this kind of vignette, what liability rules would be optimal?

While this vignette is hypothetical, AI systems like it have already entered clinical use [2]. Many of these systems use machine learning algorithms that model relationships in data to make predictions. Such systems can help interpret X-rays [3] or determine insulin dosage [4]. Nowadays, these systems are usually “locked” by their developers, but they may also be “adaptive,” adjusting

to data in the field [5]. Many apply algorithms like neural networks—so-called “deep learning,” a subset of machine learning. Such algorithms are considered “black boxes” because they are exceedingly difficult for humans to understand [6]. Thus far, the vast majority of AI systems are intended to support rather than fully determine a clinical decision, though some like Digital Diagnostics’s LumineticsCore (formerly called “IDx-DR”) are meant to be “autonomous” in the sense of providing, for instance, a complete analysis to a primary care physician as to whether to recommend referral to an ophthalmologist or rescreening in 12 months [7].

The major reason for adopting AI in medicine is its potential to produce better results than those a physician would reach without its assistance. However, demonstrating that a particular AI will achieve this is not easy. It is one thing to show success in development and testing, but quite another to prove that this improvement will be achieved reliably in the real world across different health-care systems—especially because so much of AI’s performance may depend on the larger context in which it is embedded [8]. For example, how do physicians react to AI recommendations? Do they over- or under-correct for what they perceive as its shortcomings? What role is played by insurer decisions that limit reimbursement to AI-recommended treatments?

But even when medical AI demonstrably improves outcomes *overall for a patient population*, as to a particular patient, it may produce a worse outcome than if it had not been used. In a particular case, algorithmic predictions may be inaccurate when based on erroneous data, biased data, or data otherwise poorly matched to the clinical application, when a model’s specifications make it too specific or too general, or when a system is administered incorrectly. And the picture is further complicated by system integrity issues, like adversarial attacks and data breaches [9], that are beyond the scope of this chapter.

When such adverse events occur, healthcare professionals, health systems, and product developers can face liability claims by the patient who is injured. Understanding the scope of such liability is crucial because it will shape the design and adoption of AI in health care.

### **Liability for medical AI**

A patient injured by medical AI may sue responsible parties in court to recover damages. Liability in such cases is determined according to the principles of tort law. Although case law directly addressing medical AI liability is missing, medical professionals, health systems, and medical device manufacturers may all be held liable when applying general tort law principles. Lawsuits, contracts, and regulations may affect how they share this liability.

Regulation, enforced by administrative liability, may also govern AI design and use. For instance, in the European Union, explainability obligations under the General Data Protection Regulation may inform the use of black-box algorithms in medical AI systems [10]. AI development is further shaped by the US and EU privacy laws.

### **Tort liability for medical AI: An overview**

Medical professionals may be held liable for negligent treatment. In its most basic form, such a claim requires demonstrating that a health professional breached a duty that caused an injury. In medical malpractice, such a breach occurs if a physician provides medical care that deviates from the standard of care. These principles will apply when a physician uses an AI system. Thus, a physician can be liable for injuries from the use of AI technology in diagnosis or treatment when deviating from the standard of care.

Health-care professionals may also be liable for failing to obtain a patient’s informed consent

for treatment. Again, to simplify, if a (reasonably prudent) patient would have chosen a different treatment with proper disclosure, a physician may be liable for injuries from the treatment [11]. Could the failure to disclose that AI was involved in a physician's recommended course of action lead to liability? As far as we are aware, there have been no reported cases directly on the topic, but under current case law established on informed consent more generally, liability would be most likely when AI decision-making is treated as determinative rather than as an input into the physician's decision, and when the physician's control over a procedure, or level of experience with that procedure, seems relevant to the patient's choice. One example of possible informed consent liability might be that of a physician who fails to disclose AI use in AI-driven surgery.

Health systems, physician groups, and physician-employers may also be held vicariously liable for professional negligence. Employers, as principals, may be vicariously liable for the negligence of their employees in some instances while vicarious liability for independent contractors is possible, but more challenging, to establish. This distinction according to employment status matters, especially in the United States, where nurses are often employees while physicians are more traditionally independent contractors. A court might still impute vicarious liability to a health system for a physician's use of AI on theories such as apparent agency or agency by estoppel—for example, when an injured patient has relied on representations of safety made by a hospital advertising its provision of AI-driven medicine [12].

Health systems can also be directly liable for injuries. A health system has its own traditional duties of care to patients such as providing safe facilities and equipment. The emergence of corporate negligence doctrine has enabled some patients to bring claims for negligence in supervising medical care or in enforcing hospital rules about patient care [13]. For example, a negligent credentialing claim alleges that a hospital should have known a physician at its premises was not qualified to practice. It is possible that we will see theories of corporate negligence brought to courts in relation to AI adoption. Negligent credentialing or supervision theories might be extended to treat medical AI like an additional physician: to put it pithily, we could think of hospital systems as “hiring” not “buying” an AI [14]. A health system might thus be liable for making AI technology available to practitioners without proper vetting or validation. As with medical providers, the direct liability of hospital systems will evolve as standards for AI use change [15].

Turning to the development of AI, manufacturers can be liable for injuries from product defects. Liability for selling a defective product is strict, with no need to show negligence [16]. While the risks intrinsic to medicine need not render a medical product defective, sellers can be liable for manufacturing defects, errors in product design, and failures to provide adequate warnings. Showing a design defect may entail showing that a product is unreasonably dangerous under the “risk-utility” test, which may require showing that the product could have feasibly been designed in a better way.

To an extent, current case law may inform such liability. Some types of defects in medical AI, like mechanical flaws in a product's hardware, resemble traditional sources of product liability. Health systems and medical professionals traditionally do not face liability for product defects as sellers [17], so strict liability claims may be more likely against developers. On the other hand, a health system may itself participate in product development—for example, by sharing data with a software developer.

Claims about defective AI software may also face several particular barriers [18]. It is not clear that software is a product, as opposed to a service [19]. Increased system autonomy might impede attribution of responsibility to the product's manufacturer. Because AI can have multiple

developers, it may be unclear who the manufacturer is in any case. Although errors in dataset choice or algorithmic specification are plausibly design defects, courts, mindful of the needs of innovation, seem generally hesitant to review medical product designs. Algorithmic opacity may make it difficult to perform the cost-benefit analysis necessary under the risk-utility test. And the learned intermediary doctrine limits failure-to-warn liability for medical products on the premise of the physician's intervening judgment, a rationale that may come into question if AI distances physicians from treatment decisions.

### **Apportioning liability**

Now that we have a basic sense of the overall medical AI liability landscape, we can examine how the various liable parties may interact to reapportion that liability. A tort plaintiff may recover from multiple responsible parties. Historically, manufacturers and health-care providers might have shared fully in responsibility for an "indivisible injury" [20], putting the burden on defendants to apportion damages in actions for contribution. Changes in US tort law have increased up-front apportionment [21]. Thus, if the patient in the vignette sues the treating physician, the hospital, and the AI developer, tort liability could be apportioned among the parties according to their level of fault as determined in court.

Some of this liability may be shifted by contract. For example, a health system may contract to indemnify a practitioner's malpractice liability, or a developer may put hold harmless or arbitration clauses in a contract with a health system [22]. Practitioners and health systems may purchase insurance, shifting liability risk to insurers. Indemnity and insurance could expand the medical AI market by diminishing physician liability, by clarifying its uncertainty, or by establishing rules or incentives for the use of medical AI [23]. Developers could also indemnify the health systems that purchase and implement their products—the LumineticsCore has apparently done exactly this [24], though its eagerness to do so may also stem from the fact that it was first-in-class and needed to mollify more concerns to build the market than will successor products.

Patients are less likely to contract away their rights to sue. Insurers have viewed arbitration agreements as economically unfavorable, and some courts have taken a dim view of contracts that compel injured patients to arbitrate [25]. But after an injury, health systems may offer payment in exchange for a patient agreeing not to sue [26]. Uncertainty about liability for medical AI use could increase the risks of litigation, affecting the incentives for health systems to enable, and for patients to enter, such pre-litigation settlement agreements.

Regulation may preempt product liability or malpractice claims, shifting a controlled level of risk onto the public. Programs like the FDA approvals for medical devices may substitute ex ante regulatory requirements for ex post liability. Administrative remedies can shift a controlled level of liability onto the government or mandatory insurance programs. We discuss regulatory preemption further below.

There are interesting divergences between how tort law assigns responsibility and the perspectives of the public on who is responsible. One study indicates that the public might generally impute greater fault for medical AI errors to medical professionals and health systems than developers [27]. Previous decisions in other settings have similarly conveyed hesitance to excuse physician errors in administering defective products [28]. The relatively heavier share of liability that juries might apportion to physicians and health systems could make malpractice liability particularly important in shaping the adoption of medical AI, and we turn to it next.

## Malpractice liability and the standard of care

In this section, we take a deeper dive into what malpractice actions against physicians and other medical professionals may look like when AI is involved. Malpractice actions require the breach of a duty of care, but jury members may have little experience or even intuition of the relevant baseline of appropriate care [29]. Some jurisdictions set a standard of care that reflects the customary practice for physicians with a similar level of training and resources. (Some even follow the historical practice of setting this standard with reference to local professionals.) Some jurisdictions instead impose a standard of “reasonable care,” ideally allowing liability to track the state of the art instead of custom [30]. Standards of care may also be set by statute. Applying the standard of care to technological advances may prove tricky: an established standard is at odds with the use of new technologies, but reasonable care may demand it.

When using AI deviates from the standard of care, the physician risks malpractice liability for resulting injury. By contrast, if the physician adheres to the standard of care—whether it benefits the patient or not—malpractice liability is unlikely. These principles and their implications are played out in Table 7.1 [31]. The punchline is that the current set-up of medical malpractice liability provides an incentive not to use AI when it is contrary to the ordinary course of treatment a physician would undertake. To return to our starting vignette, this would mean a physician is incented to stick with the standard dose of the chemotherapeutic agent instead of the more patient-specific dose recommended by the AI system. When that helps avoid an adverse event, that seems like a good result. But even when the AI produces better outcomes at the population level, and the dose it recommends would be safer and more effective for the patient, the physician is still incented to ignore it. Put otherwise, the chief value of the medical AI is when it tells us to do something other than what the physician would otherwise do, but medical malpractice liability will deter physicians from departing from the standard treatment in precisely such cases.

**Table 7.1 Application of malpractice to medical AI.**

<b>AI recommendation consistency with standard of care</b>	<b>Correctness of AI recommendation</b>	<b>Physician action</b>	<b>Outcome</b>	<b>Legal liability</b>
Consistent with standard of care	Correct	Follows AI	No injury	No liability
		Does not follow AI	Injury	Liability (did not follow standard of care)
	Incorrect	Follows AI	Injury	No liability (followed standard of care)
		Does not follow AI	No injury	No liability
Inconsistent with standard of care	Correct	Follows AI	No injury	No liability
		Does not follow AI	Injury	No liability (followed standard of care)
	Incorrect	Follows AI	Injury	Liability (did not follow standard of care)
		Does not follow AI	No injury	No liability

This unfortunate result is not inevitable. Incorporating the AI system's recommendation into the standard of care can enhance the incentives to use AI. If the AI recommendation is incorporated as a recognized alternative treatment under the "two schools of thought" doctrine, there is limited malpractice liability either for following the AI recommendation or for following the pre-AI standard of care. Litigation risk then provides an incentive to choose the alternative least likely to cause injury, which may be the AI system. A more extreme possibility would be if the AI recommendation becomes a required part of the standard of care. Then it is failing to follow the AI recommendation that creates liability risk by deviating from the standard of care. In that case, malpractice liability provides an incentive to use AI.

How do we get to this better place for standard of care? One possibility is that juror impressions might get us there. One study of mock jurors found that members of the public were likely to view a physician's decision to follow a nonstandard AI recommendation that later turned out to be incorrect as reasonable [32]. On the other hand, the mock jurors were more indecisive about whether a physician's decision to deviate from a correct nonstandard AI recommendation was acceptable. Such views suggest an emerging perception of AI as something like a recognized alternative treatment. But such changes in perception may require time and may be "lumpy" in that they may depend heavily on the way evidence is framed in particular trials as well as orthogonal elements of the case (such as the severity of the injury that resulted). Otherwise put, this mechanism of shifting the standard of care may not be certain enough to quell the concerns of physicians in following AI recommendations that deviate from what they would otherwise do.

Another way the standard of care might evolve is through medical practice guidelines. These guidelines have become increasingly important in malpractice cases as deference to medical customs has eroded. Practice guidelines from professional medical societies like the American Medical Association and specialty boards, health systems, insurers, and federal and state entities like the Agency for Healthcare Research and Quality may all be used as evidence of the standard of care [33]. Guidelines might thus be written to encourage or discourage the use of medical AI in particular contexts depending on its assessed risks and benefits.

The standard of care may vary for different types of AI systems. Some medical software might provide transparent information to support a final physician decision—for example, by reporting patient records or calculating well-understood relationships. This model of physician decision-making resembles the current one, so liability may resemble the current regime. Other AI products transform input data into treatment decisions in a less transparent way. The physician may choose whether to use AI but may lack any basis to assess its recommendations.

Alternatively, some products might fall in between in terms of the transparency of the decision. For example, an algorithm that uses a convolutional neural network architecture for image-based diagnosis might provide a heatmap showing where the algorithm is "looking," giving a physician limited information that may nonetheless fail to explain the decision completely [34].

For more opaque types of AI systems, courts could assess liability with reference to the physician's initial decision to use medical AI. Even when an AI system provides a limited explanation, its recommendations may be most valuable when they most supplement existing decision-making—that is, counterintuitively, when its explanation seems most inscrutable [35]. Liability that leads to second-guessing AI recommendations might close off these especially useful applications [36]. Instead of explainability, courts might look to "indicia of reliability," like regulatory approvals or empirical success in randomized trials [37]. On such indicia, a physician's initial decision to use AI may be reasonably calculated to benefit the patient despite a foreseeable

risk that the AI could be wrong. Avoiding negligence liability in such a situation seems to track the question of the standard of care, and avoids imposing strict liability for the AI's errors. Courts seem to have followed such an approach with other products, like pharmaceuticals, where physicians must similarly rely on indicia of reliability like studies and approvals [38].

The standard of care may also come to incorporate other indicia of reliability, like outside validation. For systems where physicians are poorly positioned to second-guess an AI's decisions, validation procedures, like requiring predictions to match third-party computation or requiring different AI systems to agree on a prediction, could help ensure that AI systems are functioning properly [39]. Such procedural standards may be more administrable than requiring the physician to assess an opaque system's recommendations and could be defined through regulation or practice guidelines.

As AI use increases, best practices may emerge that provide evidence of the standard of care. If greater use of AI systems leads patterns to surface in the types of ways that systems can fail, the emergence of predictable forms of harm may establish the foreseeability necessary for courts to discern breach [40].

And apart from breach, the plaintiff's causation and injury burdens may already be challenging to carry in a malpractice case. Algorithmic opacity could make it even more difficult to prove causation, a problem that could create incentives to select less transparent systems [41]. And errors in AI development that moderately impair a model without bankrupting it of predictive force, like dataset bias, could increase the risk of adverse outcomes without causing injury in the traditional doctrinal sense.

### **Regulation and preemption as alternative to the common law of torts**

We have thus far focused on the application to medical AI of the common law of torts that serves as the background principle for liability in the US system. Federal and state regulation can preempt common-law liability in exchange for more predictable regulatory requirements or administrative remedies. FDA regulation is unlikely to preempt malpractice liability for the use of AI in medical practice, but it could preempt some manufacturer liability for product defects. Federal or state legislatures might also preempt malpractice or product liability by establishing specialized adjudication procedures, mandatory insurance schemes, or general compensation funds for AI-related medical injuries. Regulation that diminishes or clarifies liability for medical AI might encourage its adoption.

### **FDA regulation**

The Federal Food, Drug, and Cosmetic Act (FDCA) makes the FDA responsible for regulating medical devices. Broadly speaking, devices are intended to diagnose, cure, mitigate, prevent, or treat diseases [42]. The 21st Century Cures Act excludes from this definition some software functions that provide health-care professionals recommendations to consider in making independent decisions, as opposed to software functions that provide the primary basis for decision-making [43]. The AI products outside this exemption that are intended for use in diagnosis, cure, mitigation, prevention, or medical treatment are likely subject to FDA regulation as medical devices, and AI systems could also be regulated as accessories to regulable devices [44].

FDA device regulation preempts some tort liability. In general, the 1976 Medical Device Amendments (MDA) to the FDCA displace state law concerning medical devices for human use [45], which the Supreme Court has interpreted to preclude some state-law tort claims. In *Medtronic*

*v. Lohr*, the Court considered a device exempt from the rigorous premarket approval (PMA) process because of its substantial equivalence to a pre-1976 device [46]. The Court found that design defect claims were not preempted because the FDA had reviewed the device under the 510(k) process for only equivalence, not product safety, and that federal manufacturing and labeling requirements were too general to preempt related state-law claims [47]. Subsequently, in *Riegel v. Medtronic*, the Court, considering a device that went through PMA [48], applied a two-part test. First, the Court considered “whether the Federal Government ha[d] established requirements applicable to” the device that could preempt state law. Unlike the 510(k) substantial equivalence process in *Lohr*, the PMA process imposed such requirements [49]. Second, the Court considered whether the state-law claims were grounded in safety- and effectiveness-related “requirements with respect to the device that are ‘different from, or in addition to,’ the federal ones,” viewing “common-law duties” as imposing such requirements [50].

These decisions could leave developers subject to some liability for defective medical AI. First, outside the tort system, developers are subject to liability under the FDCA for regulatory infractions [51]. Second, express preemption under the MDA applies a priori only to devices for human use. Nondevice software, such as under the 21st Century Cures Act exclusion, is shielded by more limited forms of preemption, if at all [52]. Third, most AI-based medical devices have entered the market via the 510(k) process [53]. Thus, manufacturers are not shielded from state-law tort claims challenging the safety and effectiveness of such devices. Substantial equivalence, as in *Lohr*, does not impose “requirements” at the first step of *Riegel*. And fourth, *Riegel* noted that state courts could still provide causes of action that paralleled the violation of FDA requirements rather than being “different from, or in addition to” those requirements [54].

One open question in this area has been whether parallel state-law claims may be grounded in industry-wide federal standards, like the FDA’s Quality System Regulations or Current Good Manufacturing Practices (CGMPs). The U.S. Court of Appeals for the Eighth Circuit answered this question in the negative, citing the generic nature of such requirements [55]. (The district court had pointed to the fact that the FDA required “manufacturers to develop *their own* quality-system controls” [56].) But rejecting this functionalist analysis, the Seventh Circuit subsequently noted that the language of the preemption provisions did not distinguish between “general” and “device-specific” requirements [57].

Another salient question is the degree to which claims are preempted under FDCA Section 310(a), which limits enforcement of the FDCA to the federal government [58]. Appellate courts have split on whether some state-law claims grounded in FDCA violations should be viewed as “attempt[s] by private parties to enforce” FDCA requirements [59] or as claims for “breach of a recognized state-law duty” [60].

Emerging requirements for software developers seem to tread near these fault lines, indicating the potential for disputes about whether FDCA requirements preempt liability for defective medical AI. For instance, proposed developer precertification standards may delegate requirements for “quality and organizational excellence” to industry [61], similarly to the CGMPs. That may invite litigation as to whether those standards are too generic to support parallel state-law claims. Similarly, reliance on post-market monitoring could lead to state-law failure-to-warn claims grounded in violations of federal reporting requirements, but the implied preemption of such claims under FDCA Section 310(a) may be in dispute [62]. At a general level, if the need to accommodate the quicker innovation pace or adaptive capabilities of AI systems leads to the backgrounding of premarket review and a heavier emphasis on postmarket monitoring, the rationale that ex ante regulation by the FDA substitutes for ex post regulation by



tort may come into question, leading courts to interpret preemption narrowly.

Similarly, the 21st Century Cures Act's exclusion of some software functions from the medical device definition under the FDCA removes the shield of express preemption. That has the effect of shifting regulatory responsibility for such software onto the tort system. Yet the exclusion only applies to some software functions that are transparent clinical decision support tools [63]. The medical device exception might thus enable ex post liability for software functions with more ex ante transparent effects while leaving the systems whose effects are more challenging to ascertain ex ante potentially shielded from liability ex post provided they will be approved via PMA (rather than 510(k)).

### **Remedies outside of tort law**

Governments could also impose administrative programs that preempt liability for the use of medical AI while providing nontort remedies. One such model is the National Vaccine Injury Compensation Program, which has vaccine manufacturers pay into a system that limits common-law tort claims and provides an administrative forum for claims instead [64]. Likewise, Florida and Virginia programs provide no-fault compensation for some birth-related injuries [65]. These programs were intended to increase the provision of covered services by preempting tort liability. But if such programs are not designed carefully and imposed by clear statutes, courts can construe the preemption narrowly [66], sometimes putting defendants in the unhappy position of both funding the program *and* being called to account in court under ordinary tort law.

If designed carefully, administrative adjudication and no-fault compensation programs might shape incentives for AI innovation and adoption. The special difficulties in showing causation of injury where opaque AI systems are involved could make administrative adjudication or compensation funds more responsive fora in some respects. Accordingly, the European Parliament has suggested considering compulsory insurance or a general compensation fund to address “the complexity of allocating responsibility” when autonomous systems inflict injury [67].

### **Conclusion**

In this chapter, we have tried to demonstrate that the question of how traditional tort law will accommodate the use of medical AI is complex and very much evolving. There are multiple forms of liability that apply differently to different players in the ecosystem (physicians, hospitals, and developers). Many decisions about what to develop and what to adopt are driven by expectations about what tort law will do. Through contractual apportionment of liability and indemnification, the players have some freedom to redesign the liability terms, but it may not be enough to harness the full value of medical AI. Preempting some forms of tort liability and adopting alternative compensation schemes through administrative processes may prove more attractive in the long run.

### **Article Information**

**Published:** I. Glenn Cohen, Andrew Slottje, Sara Gerke,

Chapter 7 - Medical AI and tort liability,

Editor(s): Joseph J.Y. Sung, Cameron Stewart,

Artificial Intelligence in Medicine,

Academic Press, 2024, Pages 89-104,

ISBN 9780323950688, <https://doi.org/10.1016/B978-0-323-95068-8.00007-8>.

## Acknowledgments

I. Glenn Cohen's work on this chapter was supported by a Novo Nordisk Foundation Grant for a scientifically independent International Collaborative Bioscience Innovation & Law Programme (Inter-CeBIL programme - grant no. NNF23SA0087056). Sara Gerke's work was funded by the European Union (Grant Agreement no. 101057321). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the Health and Digital Executive Agency. Neither the European Union nor the granting authority can be held responsible for them. S.G. also reports grants from the European Union (Grant Agreement no. 101057099), the National Institute of Biomedical Imaging and Bioengineering (NIBIB) and the National Institutes of Health Office of the Director (NIH OD) (Grant Agreement no. 3R01EB027650-03S1 and no. 1R21EB035474-01), and the National Institute on Drug Abuse (NIDA) and the National Institutes of Health (NIH) (Grant Agreement no. 1U54DA058271-01).

## References

- [1] Maier C, Hartung N, Kloft C, Huisinga W, de Wiljes J. Reinforcement learning and Bayesian data assimilation for model-informed precision dosing in oncology. *CPT Pharmacometrics Syst Pharmacol* 2021;10:241.
- [2] Maliha G, Gerke S, Cohen IG, Parikh R. Artificial intelligence and liability in medicine. *Milbank Q* 2021;99:629–34.
- [3] Gerke S, Minssen T, Cohen IG. Ethical and legal challenges of artificial intelligence-driven healthcare. In: Bohr A, Memarzadeh K, editors. *Artificial intelligence in health-care*; 2020. p. 299.
- [4] Benjamins S, Dhunoo P, Mesko B. The state of artificial intelligence-based FDA-approved medical devices and algorithms. *NPJ Digit Med* 2020;3. <https://www.nature.com/articles/s41746-020-00324-0>.
- [5] Babic B, Gerke S, Evgeniou T, Cohen IG. Algorithms on regulatory lockdown in medicine. *Science* 2019;366:1202.
- [6] Babic B, Gerke S, Evgeniou T, Cohen IG. Beware explanations from AI in health care. *Science* 2021;373:284. Babic B, Gerke S. Explaining medical AI is easier said than done. *STAT* 2021;(July 21). <https://www.statnews.com/2021/07/21/explainable-medical-ai-easier-said-than-done/>.
- [7] Gerke S, Babic B, Evgeniou T, Cohen IG. The need for a system view to regulate artificial intelligence/machine learning-based software as medical device. *NPJ Digit Med* 2020;3. <https://www.nature.com/articles/s41746-020-0262-2>. IDx-DR, Digit Diagnostics, <https://www.digitaldiagnostics.com/products/eye-disease/idx-dr/> (last visited Feb. 26, 2022).
- [8] See Gerke et al., *supra* note 7.
- [9] Finlayson SG, Bowers JD, Ito J, Zittrain JL, Beam AL, Kohane IS. Adversarial attacks on medical machine learning. *Science* 2019;363:1287.
- [10] Note that the existence and content of a right to explanation under the General Data Protection Regulation is contested. See, e.g., Wachter S, Mittelstadt B, Floridi L. Why a right to explanation of automated decision-making does not exist in the general data protection regulation. *Int Data Privacy Law* 2017;7:76.
- [11] Cohen IG. Informed consent and medical artificial intelligence. *Georgetown Law J* 2020;108:1425–32.
- [12] See Price WN II, S. Gerke, I.G. Cohen. Liability for use of artificial intelligence in

medicine. In: B. Solaiman, I.G. Cohen, editors. Research handbook in health, AI and the law; forthcoming. Alstott A. Hospital liability for negligence of independent contractor physicians under principles of apparent agency. *J Legal Med* 2004;25:487–90.

[13] Furrow B. Medical malpractice liability: of modest expansions and tightening standards. In: Cohen IG, Hoffman AK, Sage WM, editors. *The Oxford Handbook of U.S. Health Law*; 2017. p. 421–34.

[14] Gerke et al., *supra* note 7.

[15] Price WN II. Medical malpractice and black-box medicine. In: Cohen IG, Lynch HF, Vayena E, Gasser U, editors. *Big data, health law, and bioethics*; 2018. p. 295–303.

[16] Manufacturers may also be held liable for product defects on other theories, such as negligence or breach of warranty.

[17] *Hollander v. Sandoz Pharm. Corp.*, 289 F.3d 1193, 1217 n.22 (10th Cir. 2002).

[18] Schweikart SJ. Who will be liable for medical malpractice in the future? How the use of artificial intelligence in medicine will shape medical tort law. *Minn J Law Sci Technol* 2021;22:1. Selbst AD. Negligence and AI's human users. *Boston Univ Law Rev* 2020;100:1315.

[19] Evans BJ, Pasquale F. Product liability suits for FDA-regulated AI/ML software. In: Cohen IG, Minssen T, Price WN II, Robertson CT, Shachar C, editors. *The future of medical device regulation: Innovation and protection*. Cambridge University Press; 2022.

[20] See, e.g., *Temple v. Synthes Corp.*, 498 U.S. 5, 7 (1990) (*per curiam*).

[21] Restatement (third) of torts: Apportionment of liability. § E18 cmt. b (2000).

[22] Koppel R, Kreda D. Commentary, health care information technology vendors' "hold harmless" clause. *JAMA* 2009;301:1276.

[23] See Ben-Shahar O, Logue KD. Outsourcing regulation: how insurance reduces moral hazard. *Mich Law Rev* 2012;111:197–203.

[24] Panel: Understanding algorithms, artificial intelligence, and predictive analytics through real world applications. Federal Trade Commission; November 13, 2018 [comments of Michael D. Abramoff] [https://www.ftc.gov/system/files/documents/public\\_comments/2018/11/ftc-2018-0101-d-0004-162932.pdf](https://www.ftc.gov/system/files/documents/public_comments/2018/11/ftc-2018-0101-d-0004-162932.pdf).

[25] Furrow, *supra* note 13, at 431.

[26] *Id.* at 432 & n.79.

[27] Khullar D, Casalino LP, Qian Y, Yuan L, Chang E, Aneja S. Public vs physician views of liability for artificial intelligence in health care. *J Am Med Inform Ass'n* 2021;28:1574.

[28] E.g., *Bush v. Thoratec Corp.*, 13 F. Supp. 3d 554, 575 (E.D. La. 2014); see Maliha et al., *supra* note 2, at 632.

[29] Greenberg M. Medical malpractice and new devices: defining an elusive standard of care. *Health Matrix* 2009;19:423. Price WN II, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *JAMA* 2019;322:1765.

[30] E.g., *Nowatske v. Osterloh*, 543 N.W.2d 265, 272 (Wis. 1996).

[31] Adapted from Price et al., *supra* note 29.

[32] Price WN II, Gerke S, Cohen IG. How much can potential jurors tell us about liability for medical artificial intelligence? *J Nucl Med* 2021;62:15. Tobia K, Nielsen A, Stremitzer A. When does physician use of AI increase liability? *J Nucl Med* 2021;62:17.

[33] Mello MM. Of swords and shields: the role of clinical practice guidelines in medical malpractice litigation. *Univ Pennsylvania Law Rev* 2001;149:645–65. Government agency guidelines may assume particular influence in light of the Affordable Care Act. See Furrow, *supra* note 13, at 426–27.

- [34] Harned Z, Lungren MP, Rajpurkar P. Machine vision, medical AI, and malpractice. *Harv J Law Technol Digest* 2019;5. <https://jolt.law.harvard.edu/assets/digestImages/PDFs/Harned19-03.pdf>.
- [35] Tobey D. Explainability: where AI and liability meet. *DLA Piper* 2019;(Feb. 25). <https://www.dlapiper.com/en/us/insights/publications/2019/02/explainability-where-ai-and-liability-meet/>.
- [36] On the other hand, if a system's explanation of its decision supports a physician's choice, the physician might seek to adduce the explanation as evidence of due care.
- [37] See Cohen, *supra* note 11, at 1443.
- [38] See *id.*; *Richardson v. Miller*, 44 S.W.3d 1, 16–17 & nn.19–20 (Tenn. Ct. App. 2000).
- [39] Price WN II. Black-box medicine. *Harv J Law Technol* 2015;28:419–41.
- [40] Karnow CEA. The application of traditional tort theory to embodied machine intelligence. In: Calo R, Froomkin AM, Kerr I, editors. *Robot law*; 2016. p. 51–76.
- [41] McNair D, Price WN II. Health care AI: law, regulation, and policy. In: Matheny M, Israni ST, Ahmed M, Whicher D, editors. *Artificial intelligence in health care*. National Academy of Sciences; 2019. p. 197–233.
- [42] FDCA Section 201(h).
- [43] FDCA Section 520(o)(1)(E).
- [44] Price WN II. Regulating black-box medicine. *Mich Law Rev* 2017;116:439–40.
- [45] FDCA Section 521(a).
- [46] *Medtronic v. Lohr*, 518 U.S. 470, 493–94 (1996).
- [47] *Id.* at 493–94, 501.
- [48] *Riegel v. Medtronic, Inc.*, 552 U.S. 312, 320 (2008).
- [49] *Id.* at 321, 323.
- [50] *Id.* at 321–24; see also FDCA Section 521(a).
- [51] FDCA Section 303(f).
- [52] Developers of such software could argue based on the 21st Century Cures Act that tort liability would “frustrate the achievement of congressional objectives.” *Wyeth v. Levine*, 555 U.S. 555, 581 (2009). Where it does have authority to regulate, FDA might also preempt state law tort claims by adopting regulations in direct conflict with such claims. *Id.* at 582 (Breyer, J., concurring).
- [53] Artificial intelligence and machine learning (AI/ML)-enabled medical devices. FDA; December 6, 2023. <https://www.fda.gov/medical-devices/software-medical-device-samd/artificial-intelligence-and-machine-learning-aiml-enabled-medical-devices>.
- [54] *Riegel*, 552 U.S. at 330. See generally Tarloff ES. Note, medical devices and preemption. *N.Y.U. Law Rev* 2011;86:1196.
- [55] *In re Medtronic, Inc., Sprint Fidelis Leads Prod. Liab. Litig.*, 623 F.3d 1200, 1206 (8th Cir. 2010).
- [56] *In re Medtronic, Inc. Sprint Fidelis Leads Prod. Liab. Litig.*, 592 F. Supp. 2d 1147, 1157 (D. Minn. 2009) (emphasis in original).
- [57] *Bausch v. Stryker Corp.*, 630 F.3d 546, 555 (7th Cir. 2010).
- [58] *Buckman Co. v. Plaintiffs' Legal Comm.*, 531 U.S. 341, 353 (2001).
- [59] *In re Medtronic*, 623 F.3d at 1205.
- [60] *Bausch*, 630 F.3d at 557–58; accord *Hughes v. Bos. Sci. Corp.*, 631 F.3d 762, 775 (5th Cir. 2011); *Howard v. Zimmer, Inc.*, 718 F.3d 1209, 1210 (10th Cir. 2013).
- [61] Proposed regulatory framework for modifications to artificial intelligence/machine

learning (AI/ML)-based software as a medical device (SaMD). FDA; April 2, 2019. <https://www.fda.gov/media/122535/download>.

[62] Compare *In re Medtronic*, 623 F.3d at 1205, with *Hughes*, 631 F.3d at 775.

[63] FDCA Section 520(o)(1)(E).

[64] 42 U.S.C. § 300aa–10 et seq.; see *Schafer v. Am. Cyanamid Co.*, 20 F.3d 1, 2 (1st Cir. 1994) (Breyer, C.J.).

[65] Fla. Stat. Ann. § 766.303 (West 2021); Va. Code Ann. § 38.2-5002 (West 2021).

[66] See *Moss v. Merck & Co.*, 381 F.3d 501, 503–04 (5th Cir. 2004); *Schafer*, 20 F.3d at 7; Engstrom NF. Exit, adversarialism, and the stubborn persistence of tort. *J Tort Law* 2013;6:75.

[67] European Parliament Resolution of 16 February 2017 with Recommendations to the Commission on Civil Law Rules on Robotics (2015/2103(INL)), Eur. Parl. Doc. P8\_TA(2017)0051 (2017), ¶¶ 57, 59, [https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051\\_EN.html](https://www.europarl.europa.eu/doceo/document/TA-8-2017-0051_EN.html).