HEALTH LAW, ETHICS, AND HUMAN RIGHTS

Understanding Liability Risk from Using Health Care Artificial Intelligence Tools

Michelle M. Mello, J.D., Ph.D., and Neel Guha, M.S.

Optimism about the explosive potential of artificial intelligence (AI) to transform medicine is tempered by worry about what it may mean for the clinicians being "augmented." One question is especially problematic because it may chill adoption: when AI contributes to patient injury, who will be held responsible?

Some attorneys counsel health care organizations with dire warnings about liability[1] and dauntingly long lists of legal concerns.[2] Unfortunately, liability concern can lead to overly conservative decisions,[3] including reluctance to try new things. Yet, older forms of clinical decision support provided important opportunities to prevent errors and malpractice claims.[4] Given the slow progress in reducing diagnostic errors, not adopting new tools also has consequences and at some point may itself become malpractice.[5] Liability uncertainty also affects AI developers' cost of capital and incentives to develop particular products, thereby influencing which AI innovations become available and at what price.

To help health care organizations and physicians weigh AI-related liability risk against the benefits of adoption, we examine the issues that courts have grappled with in cases involving software error and what makes them so challenging. Because the signals emerging from case law remain somewhat faint, we conducted further analysis of the aspects of AI tools that elevate or mitigate legal risk. Drawing on both analyses, we provide risk-management recommendations, focusing on the uses of AI in direct patient care with a "human in the loop," since the use of fully autonomous systems raises additional issues.[6]

THE AWKWARD ADOLESCENCE OF SOFTWARE-RELATED LIABILITY

Legal precedent regarding AI injuries is rare because AI models are new and few personal-injury claims result in written opinions. As this area of law matures, it will confront several challenges.

Ordinarily, when a physician uses or recommends a product and an injury to the patient results, well-established rules help courts allocate liability among the physician, product maker, and patient. The liabilities of the physician and product maker are derived from different standards of care, but for both kinds of defendants, plaintiffs must show that the defendant owed them a duty, the defendant breached the applicable standard of care, and the breach caused their injury; plaintiffs must also rebut any suggestion that the injury was so unusual as to be outside the scope of liability (Table 1).

Several factors make these determinations difficult with respect to AI and other software, especially for claims against developers (Table 1). First, because software is intangible, courts have been reluctant to apply doctrines of product liability to it — a stance that affects the applicable standard of care. In addition, another doctrine (called "preemption") bars personal-injury claims in state court when they relate to some devices that have been cleared by the Food and Drug Administration (FDA). Although much health care AI never undergoes FDA review, among the AI-enabled devices that do, it is somewhat unclear which devices this doctrine covers.[8]

Another complicating factor is that in most states, plaintiffs alleging that complex products were defectively designed must show that there is a reasonable alternative design that would be safer, but it is difficult to apply that concept to AI. AI models are essentially mathematical equations that encode statistical patterns learned automatically from data. Plaintiffs must show that some such patterns were "defective" and that their injury was foreseeable from the patterns learned. However, because such patterns

Table 1. Challenges in Applying Tort Law Principles to Health Care Artificial Intelligence (AI).

Tort Claim Element	Traditional Approach to Proving	Challenges in Claims Related to AI
Defendant owed plaintiff a legal duty	For malpractice, show evidence that a practitioner (or facility) had established a treatment relationship with the plaintiff. For products, argue that a plaintiff was a foreseeable user or bystander.	Not generally a problem, but if AI is embedded in certain medical devices that had been reviewed by the Food and Drug Administration, product-liability claims may be preempted by federal law.
Defendant's conduct fell below the standard of care	For malpractice, show evidence that care fell below what a reasonable practitioner in the same field (or a facility with similar resources) would have provided in the circumstances. For claims against facilities, argue that equipment or software was negligently selected, maintained, or monitored. For products, show evidence that product had a manufacturing or design defect or that defendant did not supply adequate warnings or instructions.	Model opacity makes it difficult to prove that a physician's decision to accept or depart from output was unreasonable. Wrong model output for a particular patient may not have been foreseeable by a physician. AI may not be considered a product. Difficult to show that a reasonable alternative safer design exists.
Plaintiff had an injury	Show evidence of physical or emotional injury.	Proving algorithmic bias (inferior model performance for some patient subgroups) is insufficient unless actual injury to a plaintiff had resulted.[7]
Defendant's conduct was a factual cause of plaintiff's injury	Usually, show evidence that the injury would not have occurred but for the defendant's conduct (or the defect in the product).	Model opacity makes it difficult to prove that wrong output occurred because of a defect.
Plaintiff was a foreseeable victim injured in a foreseeable way	Rebut the defendant's argument that a very unusual series of events led to the injury.	No distinctive issues at present, but in the future, autonomous AI could make unexpected decisions.[7]

are represented with the use of up to billions of variables, identifying the patterns on which a system relies is technically challenging. Plaintiffs can suggest better training data or validation processes but may struggle to prove that these would have changed the patterns enough to eliminate the "defect." The same training data may cause poor performance in one model but not in another that used different learning algorithms.

Models that perform well in general may not perform as well for particular patients or groups. Medical data sets that are used to train and evaluate AI models represent distinctive patient populations and settings, which makes it difficult to estimate how often outputs will be wrong for others.[9] The foreseeability of particular errors will be contested, although plaintiffs may credibly allege that defendants were aware of "distribution shift" (mismatch between the training data and the patients for whom the model is used).[7,10,11] Model opacity and distribution shift can also create causation conundrums. No model is perfect; in a given case, how can it be proved that wrong output occurred because of the alleged defect? Courts also may have conceptual difficulty in deeming physicians and hospitals to be negligent for relying on models that, on average, deliver better results than humans alone.[12]

Furthermore, plaintiffs, when suing clinicians, must show that the decision to accept (or depart from) the model output was unreasonable. Known distribution shift is one avenue for argumentation, but plaintiffs may falter without detailed information about how the model reached its conclusion.[13] In addition, how should liability be allocated when issues relating to the clinical integration of the model had heightened the risk that its errors would reach the patient? This question recalls older conversations about responsibility for medical errors. Individual clinicians and systems both contribute to errors — yet physicians remain the primary locus of liability. This history will not reassure physicians evaluating AI risk.

In summary, AI poses challenges for applying tort principles. Because it is primarily plaintiffs who will struggle, liability worries may be outsized during this period of adolescence for software-related tort doctrine. However, we believe that this situation cannot hold. Tort doctrine will evolve to address needs arising from tech-

nological changes, as it has historically.[14] To investigate whether evolution is already visible, we reviewed relevant judicial decisions.

SOFTWARE-RELATED LIABILITY IN THE COURTS

We collected judicial opinions in tort cases regarding AI and other software in health care and non–health care contexts, supplementing results with searches of jury verdicts, news and scholarly articles, and legal newsletters. We manually reviewed 803 unique cases, extracting information on the dominant issues that courts addressed in the 51 cases that involved software-related errors that caused physical injury (details are provided in the Supplementary Appendix, available with the full text of this article at NEJM.org, and at https://osf.io/zvmku/).

Reported cases involving medical software and AI (Table S2 in the Supplementary Appendix) have clustered around three situations. One situation involves harms to patients caused by defects in software that is used to manage care or resources. Typically, plaintiffs bring product-liability claims against the developer. For example, in *Lowe v. Cerner*, the court held that the plaintiffs had made a viable claim that a defective user interface in drug-management software led physicians to mistakenly believe that they had scheduled medication. Plaintiffs also sue hospitals for the role they play in selecting, maintaining, or updating such systems, as in the 2014 case *Ambrose v. St. Joseph's Hospital of Atlanta*, in which the failure by a hospital to update software on a surgical microscope allegedly harmed patients. The twist in these cases is that plaintiffs may bring ordinary negligence (rather than medical malpractice) claims because the software had been handling administrative functions.

A second situation involves physicians having consulted software in making care decisions (e.g., to screen patients for certain conditions or generate medication regimens). When physicians adhere to erroneous software recommendations, patients bring malpractice claims alleging that the physician should have ignored the recommendation or independently reached the correct decision. For instance, in the 2023 case *Sampson v. HeartWise Health Systems Corporation*, physicians followed the output of a software program for cardiac health screening, which classified a young adult patient with a family history of congenital heart defects as "normal" on the basis of clinical test results. When the man died weeks later of a congenital heart condition, the family sued the physicians, alleging that they should have scrutinized the output of the software more closely and relied on their own interpretation of the tests. The court denied the defendant's motion for summary judgment, which means that the case could proceed to trial.

In such cases, physicians' decisions are evaluated against what other specialists would have done. In *Skounakis v. Sotillo* in 2018, for instance, the plaintiff was prescribed a harmful combination of drugs by a physician who followed the recommendation from a software program. The court reversed an earlier summary judgment for the defendants, finding that the plaintiff had produced sufficient expert testimony on the operative legal question of whether the physician's interactions with the patient, understanding of the patient's history, and drug knowledge met the customary standard of care.

Approaches taken by courts in claims against software developers are varied. In at least one (nonmedical) case (*Rodgers v. Christie* in 2020), a court dismissed claims of design defect on the basis that algorithms are not products and that applying tort doctrine to algorithmic recommendations might implicate speech rights. In contrast, the court in *Sampson* suggested that plaintiffs could have brought design-defect claims over the developer's choice of clinical tests to include in the algorithm. Courts have also disagreed on whether plaintiffs may bring negligence claims for the content of software recommendations. In *Skounakis*, the court allowed both ordinary and medical-negligence claims against the software developer to proceed to trial on the theory that the medication recommendation made by the software violated the standard of care. In *Sampson*, however, in response to the plaintiff's claim of ordinary negligence, the court granted summary judgment for the defendants on two grounds: first, the developer had no duty to give the patient an accurate diagnosis because its licensing agreement with the clinic gave physicians responsibility for final decision making, and second, a state statute barred the claim because the parties had agreed that the developer

was not a "health care provider" under Alabama law. Although it is unclear what drove the disparate outcomes in these cases, they collectively suggest a future in which developer liability may vary depending on private contracting and jurisdictional variation.

A third situation in which cases arise involves apparent malfunctions of software embedded within devices, such as implantables, surgical robots, or monitoring tools. Plaintiffs may assert malpractice claims against physicians and hospitals, alleging negligent use, installation, or maintenance of these devices, as in the 2006 case of *Seargeant v. Orthopedic Associates Medical Clinic*, in which physicians, a technician, and a clinic were sued after human error during routine reprogramming of an infusion pump led to lethal morphine administration. Plaintiffs may also sue developers, alleging defects in manufacturing, design, and warnings. In these claims, plaintiffs must navigate preemption defenses as well as device complexity. Courts are skeptical when plaintiffs frame the failure of the device as its defect and instead demand that they identify specific design flaws. Our review suggests that this is no easy task: plaintiffs often fail to identify defects specific enough to survive defendants' motions for summary judgment.

Software-related cases to date signal three emerging trends. First, cases involving software defects in implantable devices suggest that plaintiffs struggle to sustain claims when diminished visibility into the workings of the device makes identifying a specific design defect difficult. The complexity and opacity of AI lead to similar issues. Second, the cases involving software recommendations suggest that the varying performance of AI for different patient groups will force courts to grapple with determining when a physician reasonably should have known that the output was not reliable for particular patients. Third, across all case clusters, the reluctance by courts to distinguish "AI" from "traditional" software suggests that rules or approaches that courts create in AI-related cases may have spillover effects on non-AI software (and vice versa), although technical differences may make them ill-suited to another type of model. For example, courts might relax requirements for proving design defects, although not all software models present opacity problems.

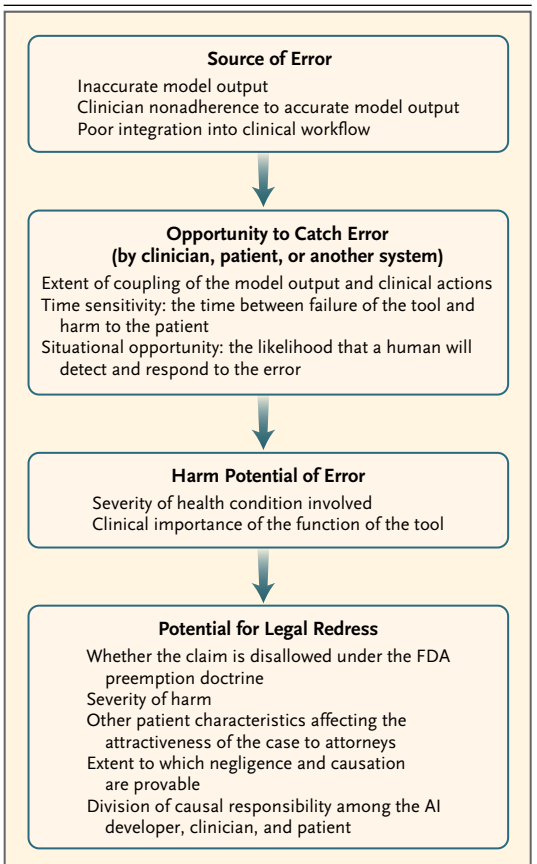

Figure 1. Typology of Factors Influencing Liability Risk of Health Care Artificial Intelligence Tools.

FDA denotes Food and Drug Administration.

ASSESSING LEGAL RISK IN AI DEPLOYMENTS

Although courts have lumped together different types of software, health care organizations should not follow suit when evaluating the risks and benefits of AI adoption. AI is not one technology but a heterogeneous group with varying liability risks. Identifying AI tools with the greatest risk can help target risk-management interventions and clinical oversight.

A framework to support health care organizations and clinicians in assessing AI liability risk is provided in Figure 1. The framework incorporates our findings regarding how courts evaluate claims related to software errors and broadens the lens to include assessment of the likelihood that claims will be brought. Drawing on previous

conceptual work in safety science[15,16] and malpractice claiming dynamics,[17] we conceptualized risk as a function of the following four factors: the likelihood and nature of model errors, the likelihood that humans or another system will detect the errors and prevent harm, the potential harm if errors are not caught, and the likelihood that injuries would garner compensation in the tort system.

In the first step of liability risk assessment, reasonable expectations about the likelihood and nature of errors should be identified on the basis of the model, its training data, and its task design. For instance, what is known about disparities between the training data and the patients for whom the model will be deployed? The plan for clinical integration is also important — for example, will some clinicians be inclined to dismiss model recommendations because of engrained habits or distaste for AI? Research suggests that users' trust in machine learning models predicts their willingness to adjust their judgments on the basis of model output.[18,19] Because model performance and human reactions to model output may change over time, they must be reassessed periodically.

The second step involves assessing the scope of the opportunity for catching errors introduced by the model. How much buffer or time exists between failure of the tool and harm to the patient? How rich is the situational opportunity for intervention? That is, what is the likelihood that, even with time permitting, humans will detect and respond to the error? Situational opportunity is greatest when AI tools are highly visible (i.e., when they and their warnings and guidelines are not easily forgotten) and when clinicians have considerable information about the model to help focus their vigilance (i.e., the model is not so opaque as to give scant basis for evaluating the correctness of its output). Information about false negative and false positive rates and patient-specific probability scores can help clinicians scrutinize predictions.[20] Information about training data facilitates assessment of whether a given patient is well represented and of the task design. For example, for reasons of data availability, an algorithm may be tasked with sorting patients according to clinical need, but the algorithm may be trained on proxy measures such as health care costs; understanding the shortcomings of the proxy measures can help anticipate inequitable results[21] and errors.

For situational opportunity to be effective, it must not rest on unrealistic assumptions about how humans behave.[22,23] Research in human–computer interactions shows that decision makers who are assisted by computer models frequently overrely on model output, fail to recognize when it is incorrect, and do not intervene when they should, a phenomenon known as automation bias.[24,25] Can busy physicians, for example, be counted on to thoughtfully edit large language model–generated draft replies to patients' emails, investigate whether model-recommended drugs are indeed appropriate for a given patient, or catch errors in visit notes produced by speech-to-text models? Humans are also prone to automation-induced complacency or failure to appropriately monitor computer-based decision-support systems.[25] AI deployments that assume a high level of oversight by clinicians battling time pressure and overwork may not provide meaningful catch opportunities.

The third step involves asking how serious the potential harm would be to patients if model-related errors are not caught. Tools that perform critical clinical functions and those used in the care of patients with serious health conditions are of particular concern.[16]

The final step is to assess how likely it would be for injured patients to find legal redress. Although all preventable injuries are regrettable, most never become claims. Patients who obtain legal representation are overwhelmingly those with serious injuries.[17] AI tools with straightforward causal paths to injury are attractive subjects for claims. A further consideration for AI–device combinations is whether preemption applies.[8] In addition, how steep are the barriers to proving negligence? Models with the highest risk of error may not pose the greatest liability risk because of the problems that plaintiffs may encounter in proving design defect. Higher-performing models in which errors are more easily identifiable to plaintiffs may involve greater liability risk than poorer-performing models in which the operation is more opaque.[26] In addition, how will responsibility for the harm probably be allocated? Health care organizations may face greater liability for situations in which errors are more likely to have resulted from human

conduct or clinical integration than from erroneous output, for injuries in which patients' own decisions had no bearing, and for homegrown AI tools, as opposed to externally developed ones.

RISK-MANAGEMENT RECOMMENDATIONS

While awaiting clarification of how tort doctrine will evolve to address AI, health care organizations and clinicians can take several steps to manage liability uncertainty. One such step is to resist the temptation to lump all applications of AI together. Adoption decisions and postdeployment monitoring should reflect the fact that some tools are riskier than others. When tools have the hallmarks of high liability risk that we have identified (e.g., low opportunity to catch the error, high potential for patient harm, and unrealistic assumptions about clinician behavior), organizations should expect to allocate substantial time and resources to safety monitoring and gather considerable information from model developers and implementation teams. In contrast, for lower-risk tools, organizations may be able to apply more generalized, lower-touch monitoring.

Another step is to recognize that health care organizations are currently in a buyer's market. With so many AI developers jockeying to gain footholds in health systems and access their patient data, opportunities exist to bargain for terms that minimize purchasers' liability risk. Licensing agreements should, for instance, require developers to provide information necessary for effective risk assessment and monitoring. Such information includes developers' assumptions regarding the data that models will ingest, processes for validating models, and recommendations for auditing model performance (perhaps with statistical indications that constitute early warning signs for systematic errors). However, it should be noted that although disclosures can improve safety monitoring, they could increase liability risk if they shunt auditing responsibilities onto purchasers who do not follow through.

Purchasers should also insist on favorable terms governing liability, insurance, and risk management in AI licensing contracts. Although courts and legislatures set the rules about when injured persons are eligible for compensation, contracting parties have wide latitude to use indemnification clauses to establish which of them pays in the event of a qualifying injury.[27,28] Indemnification provisions can require that developers pay for errors in model output, whereas hospitals or clinicians pay for those arising from poor deployment or misuse. Certainly, purchasers should refuse clauses that explicitly or implicitly (as in *Sampson*) immunize developers from liability or cap their financial responsibility. Contracts can also specify minimum insurance requirements and postdeployment monitoring responsibilities.

When models are developed in house, there is no external developer to assume legal obligations; having adequate insurance is therefore critical.[29] Professional liability insurers may impose coverage exclusions for AI-related injuries, and cyber policies may cover only economic losses, not physical injuries. Organizations should ensure that their coverage is not limited in these ways and is deep enough to cover worst-case scenarios in which a systematic error affects many patients.

We expect that insurance designed specifically for AI will become increasingly available. Such insurance could have benefits beyond risk spreading: these insurers will have data, research capacity, and expertise to give the persons and entities they insure sophisticated, tailored loss-prevention advice. Insurers also apply their expertise when setting premiums, which should incentivize AI developers and adopters to provide robust evidence that AI models are safe.[29-31]

Given that courts seem disinclined to create special new rules for AI, another useful step is to apply lessons learned from older forms of decision support. In cases involving clinical practice guidelines and alerts in electronic health records, for instance, courts examine whether the recommendation was evidence-based and whether the physician should have heeded it for the patient in question.[32,33] Some problems that seem distinctive to AI actually echo these older determinations. For example, the problem of distribution shift in AI resembles arguments that a practice guideline is based on outdated or unrepresentative studies. Another example is that physicians' decisions not to follow AI output may reflect reasoning similar to decisions not to follow guidelines. Our case review suggests that

courts can be expected to adopt similar modes of analysis.

Health care organizations should also anticipate the evidentiary problems that may arise in AI litigation. AI models may be frequently updated in order to account for distribution shift, yet litigation will require that parties be able to reproduce past predictions. Our reviewed cases included instances in which failure to appropriately track software versions or types prolonged litigation. Model inputs, outputs, and versions should be documented at the time of care, along with the reasons that clinicians followed or departed from model recommendations.

It is also useful for health care organizations to recognize that the defense of AI cases may require different expertise than what malpractice defense counsel are accustomed to needing. Our case review suggests the question of who qualifies as a health care AI expert is far from settled. In addition to cultivating relationships with expert witnesses in computer science,[34] counsel will need to develop sufficient familiarity with AI methods to be able to quarterback a legal defense.

It also may be prudent to inform patients when AI models are used in diagnostic or treatment decisions. In evaluating claims alleging breach of informed consent, many jurisdictions apply a patient-centered standard to decide what constitutes material information that should have been disclosed, and unlike with other software, surveys indicate a majority of U.S. residents feel uncomfortable about AI being used in their care.[35,36] If use of AI is documented in the medical record, it will come to light during litigation; disclosure to patients reduces the risk that plaintiffs will add informed-consent claims in response. A reasonable disclosure might include what function the model serves, what shortcomings are known, how the team uses output in light of shortcomings, and why they believe that its use improves care.

Finally, as with all medical errors, the best risk-management strategy is to prevent injuries. Following emerging guidelines for evaluating AI model safety[37-39] can help minimize the human and financial cost that the leap into AI-informed medicine involves.

Disclosure forms provided by the authors are available with the full text of this article at NEJM.org.

We thank Nicholson Price, Jonathan Chen, Curt Langlotz, and Russ Altman for feedback on an earlier draft of the manuscript.

From Stanford Law School (M.M.M., N.G.), the Department of Health Policy, School of Medicine (M.M.M.), the Freeman Spogli Institute for International Studies (M.M.M.), and the Department of Computer Science (N.G.), Stanford University, Stanford, CA.

1. Keris MP. Artificial intelligence in medicine creates real risk management and litigation issues. J Healthc Risk Manag 2020; 40:21-6.
2. Broccolo BM. Coping with the mystery and reality of artificial intelligence in health care. Washington, DC: American Health Law Association, 2019 (https://www.americanhealthlaw.org/getmedia/ebe19121-1af0-4d26-b208-3aa40b5f4486/I_Overview_Coping-with-the-Mystery-and-the-Reality_June2019.pdf).
3. Carrier ER, Reschovsky JD, Mello MM, Mayrell RC, Katz D. Physicians' fears of malpractice lawsuits are not assuaged by tort reforms. Health Aff (Millwood) 2010;29:1585-92.
4. Zuccotti G, Maloney FL, Feblowitz J, Samal L, Sato L, Wright A. Reducing risk with clinical decision support: a study of closed malpractice claims. Appl Clin Inform 2014;5:746-56.
5. 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.
6. Duffourc MN. Malpractice by the autonomous AI physician. Univ Ill J Law Tech & Pol'y 2023;2023:1-49.
7. Selbst AD. Negligence and AI's human users. Boston Univ Law Rev 2020;100:1315-76.
8. Tschider CA. Medical device artificial intelligence: the new tort frontier. BYU Law Rev 2021;46:1551-617.
9. Kaushal A, Altman R, Langlotz C. Geographic distribution of US cohorts used to train deep learning algorithms. JAMA 2020;324:1212-3.
10. Price WN, Cohen IG. Locating liability for medical AI. SSRN. July 31, 2023. abstract (https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4517740).
11. Finlayson SG, Subbaswamy A, Singh K, et al. The clinician and dataset shift in artificial intelligence. N Engl J Med 2021; 385:283-6.
12. Evans B, Pasquale FA. Product liability suits for FDA-regulated AI/ML software. In: Cohen IG, Minssen T, Price WN, Robertson C, Shachar C, eds. The future of medical device regulation: innovation and protection. Cambridge, United Kingdom: Cambridge University Press, 2022.
13. Price WN. Medical malpractice and black-box medicine. In: Cohen IG, Lynch HF, Vayena E, Gasser U, eds. Big data, health law, and bioethics. Cambridge, United Kingdom: Cambridge University Press, 2018.
14. Gifford DG. Technological triggers to tort revolutions: steam locomotives, autonomous vehicles, and accident compensation. J Tort Law 2018;11:71-143.
15. Reason J. Human error: models and management. BMJ 2000; 320:768-70.
16. International Medical Device Regulators Forum. Software as a medical device: possible framework for risk categorization and corresponding considerations. September 18, 2014 (https://www.imdrf.org/documents/software-medical-device-possible-framework-risk-categorization-and-corresponding-considerations).
17. Studdert DM, Mello MM, Gawande AA, et al. Claims, errors, and compensation payments in medical malpractice litigation. N Engl J Med 2006;354:2024-33.
18. McGrath S, Mehta P, Zytek A, Lage I, Lakkaraju H. When does uncertainty matter? Understanding the impact of predictive uncertainty in ML assisted decision making. Boston: Harvard Business School, June 2023 (https://www.hbs.edu/faculty/Pages/item.aspx?num=64210).

19. Lai V, Chen C, Liao QV, Smith-Renner A, Tan C. Towards a science of human-AI decision making: a survey of empirical studies. In: Proceedings and Abstracts of the 2023 ACM Conference on Fairness, Accountability, and Transparency, June 12–15, 2023. Chicago: Association for Computing Machinery, 2023.
20. Kompa B, Snoek J, Beam AL. Second opinion needed: communicating uncertainty in medical machine learning. NPJ Digit Med 2021;4:4.
21. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. Science 2019;366:447-53.
22. Crootof R, Kaminski ME, Price WN. Humans in the loop. Vanderbilt Law Rev 2023;76:429-510.
23. Kostick-Quenet KM, Gerke S. AI in the hands of imperfect users. NPJ Digit Med 2022;5:197.
24. Parasuraman R, Manzey DH. Complacency and bias in human use of automation: an attentional integration. Hum Factors 2010;52:381-410.
25. Chuganova M, Sele D. We and it: an interdisciplinary review of the experimental evidence on how humans interact with machines. J Behav Exp Econ 2022;99:101897.
26. Burrell J. How the machine 'thinks': Understanding opacity in machine learning algorithms. Big Data Soc 2016;3:1-12.
27. Banja JD, Hollstein RD, Bruno MA. When artificial intelligence models surpass physician performance: medical malpractice liability in an era of advanced artificial intelligence. J Am Coll Radiol 2022;19:816-20.
28. Maliha G, Gerke S, Cohen IG, Parikh RB. Artificial intelligence and liability in medicine: balancing safety and innovation. Milbank Q 2021;99:629-47.
29. Stern AD, Goldfarb A, Minssen T, Price WN. AI insurance: how liability insurance can drive the responsible adoption of artificial intelligence in health care. Waltham, MA: NEJM Catalyst March 16, 2022 (https://catalyst.nejm.org/doi/full/10.1056/CAT.21.0242).
30. Baker T, Swedloff R. Regulation by liability insurance: from auto to lawyers professional liability. UCLA Law Rev 2013;60:1412-50.
31. Lior A. Insuring AI: the role of insurance in artificial intelligence regulation. Harv J Law Technol 2022;35:467-530.
32. Mello MM, Guha N. ChatGPT and physicians' malpractice risk. JAMA Health Forum 2023;4(5):e231938.
33. Mangalmurti SS, Murtagh L, Mello MM. Medical malpractice liability in the age of electronic health records. N Engl J Med 2010;363:2060-7.
34. Laddon T, O'Reilly T, Steeb E, Lawson J. Who's to blame? Get smart about smart medical devices. In-House Defense Q 2021;16:15-9.
35. Tyson A, Pasquini G, Spencer A, Funk C. 60% of Americans would be uncomfortable with provider relying on AI in their own health care. Washington, DC: Pew Research Center, February 22, 2023 (https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/).
36. Robertson C, Woods A, Bergstrand K, Findley J, Balser C, Slepian MJ. Diverse patients' attitudes towards Artificial Intelligence (AI) in diagnosis. PLOS Digit Health 2023;2(5):e0000237.
37. Dixit A, Quaglietta J, Gaulton C. Preparing for the future: how organizations can prepare boards, leaders, and risk managers for artificial intelligence. Healthc Manage Forum 2021;34:346-52.
38. National Institute of Standards and Technology. Artificial intelligence risk management framework (AI RMF 1.0). January 2023 (https://doi.org/10.6028/NIST.AI.100-1).
39. White House. Executive Order on the safe, secure, and trustworthy development and use of artificial intelligence. October 30, 2023 (https://www.whitehouse.gov/briefing-room/presidential-actions/2023/10/30/executive-order-on-the-safe-secure-and-trustworthy-development-and-use-of-artificial-intelligence/).

DOI: 10.1056/NEJMhle2308901