

Perspective

Artificial Intelligence and Liability in Medicine: Balancing Safety and Innovation

GEORGE MALIHA,^{*,†} SARA GERKE,[‡]
I. GLENN COHEN,^{‡,§} and RAVI B. PARIKH^{*,†,||,#}

**Perelman School of Medicine, University of Pennsylvania; †Department of Internal Medicine, University of Pennsylvania; ‡Petrie-Flom Center for Health Law Policy, Biotechnology, and Bioethics, Harvard Law School, Harvard University; §Harvard Law School, Harvard University; ||Penn Center for Cancer Care Innovation, University of Pennsylvania; #Corporal Michael J. Crescenz VA Medical Center*

Policy Points:

- With increasing integration of artificial intelligence and machine learning in medicine, there are concerns that algorithm inaccuracy could lead to patient injury and medical liability.
- While prior work has focused on medical malpractice, the artificial intelligence ecosystem consists of multiple stakeholders beyond clinicians. Current liability frameworks are inadequate to encourage both safe clinical implementation and disruptive innovation of artificial intelligence.
- Several policy options could ensure a more balanced liability system, including altering the standard of care, insurance, indemnification, special/no-fault adjudication systems, and regulation. Such liability frameworks could facilitate safe and expedient implementation of artificial intelligence and machine learning in clinical care.

Keywords: Machine learning, artificial intelligence, liability, malpractice.

AFTER MORE THAN A DECADE OF PROMISE AND HYPE, ARTIFICIAL intelligence (AI) and machine learning (ML) are finally making inroads into clinical practice. AI is defined as a “larger

IGC and RBP contributed equally to this article

The Milbank Quarterly, Vol. 99, No. 3, 2021 (pp. 629-647)
© 2021 Milbank Memorial Fund

umbrella of computer intelligence; a program that can sense, reason, act, and adapt,” and ML is a “type of AI that uses algorithms whose performance improves as they are exposed to more data over time.”¹ From 2017 to 2019, the US Food and Drug Administration (FDA) approved or cleared over 40 devices based on AI/ML algorithms for clinical use.^{2–4} Many of these devices improve detection of potential pathology from image-based sources, such as radiographs, electrocardiograms, or biopsies.^{5–8} Increasingly, the FDA has permitted marketing of AI/ML predictive algorithms for clinical use, some of which outperform physician assessment.^{5,9} Despite their promise, AI/ML algorithms have come under scrutiny for inconsistent performance, particularly among minority communities.^{10–13}

Algorithm inaccuracy may lead to suboptimal clinical decision-making and adverse patient outcomes. These errors raise concerns over liability for patient injury.¹⁴ Thus far, such concerns have focused on physician malpractice.¹⁴ However, physicians exist as part of an ecosystem that also includes health systems and AI/ML device manufacturers. Physician liability over use of AI/ML is inextricably linked to the liability of these other actors.^{15,16} Furthermore, the allocation of liability determines not only whether and from whom patients obtain redress but also whether potentially useful algorithms will make their way into practice: Increasing liability for use or development of algorithms may disincentivize developers and health system leaders from introducing them into clinical practice. We examine the larger ecosystem of AI/ML liability and its role in ensuring both safe implementation and innovation in clinical care.

Overview of AI/ML Liability

Physicians, health systems, and algorithm designers are subject to different, overlapping theories of liability for AI/ML systems (see Table 1). Broadly, physicians and health systems may be liable under malpractice and other negligence theories but are unlikely to be subject to products liability. Conversely, algorithm designers may be subject to products liability but are not likely to be liable under negligence theories.

Table 1. Current Landscape of AI/ML Liability

| Type of Liability (Definition) | Implications for Physicians | Implications for Developers or Health Systems |
|---|--|---|
| Medical malpractice (Deviating from the standard of care set by the profession) | Physicians may be liable for failing to critically evaluate AI/ML recommendations. This may change as AI/ML systems integrate into clinical care and become the standard of care. | Health systems or practices that employ or credential physicians and other health care practitioners may be liable for practitioners' errors ("vicarious liability"). |
| Other negligence (Deviating from the norms set by an industry and courts) | Physicians may be liable for (1) their decision to implement an improper AI/ML system in their practice, or (2) their employees' negligent treatment decisions related to AI/ML systems ("vicarious liability"). | Hospital liability for negligent credentialing of physicians could extend to failure to properly assess a new AI/ML system. In general, health systems may be liable for failing to provide training, updates, support, maintenance, or equipment for an AI/ML algorithm. |
| Products liability (Designing a product that caused an injury) | Physicians might be involved in these cases if they work or consult for designers of AI/ML devices. | The law is unsettled in this area. As AI/ML software integrates into care or becomes more complex, algorithm developers may have to contend with liability. |

Abbreviations: AI, artificial intelligence; ML, machine learning.

Negligence

Medical Malpractice. Medical malpractice requires an injury caused by a physician's deviation from the standard of care. This standard of care is set by the collective actions of the physician's professional peers, based on local or national standards.^{14,17,18}

A physician who in good faith relies on an AI/ML system to provide recommendations may *still* face liability if the actions the physician takes fall below the standard of care and other elements of medical malpractice are met. Physicians have a duty to independently apply the standard of care for their field, regardless of an AI/ML algorithm output.^{19,20} While case law on physician use of AI/ML is not yet well developed, several lines of cases suggest that physicians bear the burden of errors resulting from AI/ML output. First, courts have allowed malpractice claims to proceed against health professionals in cases where there were mistakes in medical literature given to patients or when a practitioner relies on an intake form that does not ask for a complete history.^{21–24} Second, courts have generally been reluctant to allow physicians to offset their liability when a pharmaceutical company does not adequately warn of a therapy's adverse effect.^{25–27} Third, most courts have been unwilling to use clinical guidelines as definitive proof of the standard of care and require a more individualized determination in a particular case.^{19,28} Fourth, many courts are disinclined to excuse malpractice based on errors by system technicians or manufacturers.^{29–31}

Other Negligence. Outside of the direct physician/practitioner-patient interaction, health systems, hospitals, and practices are also responsible for the well-being of patients.^{17,32} Negligent credentialing theories may hold a health system or physician group liable for not properly vetting a physician who deviates from the standard of care.¹⁸ Some legal experts have suggested that a similar liability may extend to health systems that insufficiently vet an AI/ML system prior to clinical implementation.¹⁸ Indeed, health systems already have a well-defined duty to provide safe equipment and facilities and to train their employees to use health system-provided equipment.^{33–35} Thus, evidence of poor implementation, training, or vetting of an AI/ML system could form a basis of a claim against a health system.

To illustrate how this negligence would apply to a case involving AI/ML, take an example of an algorithm trained on a predominantly young, healthy population that uses vital signs and laboratory readings

to recommend treatment for hypertension. Current guidelines suggest that older adults may benefit from tailored management of hypertension with higher blood pressure thresholds than those used to indicate interventions for younger adults.^{36–40} A health system serving a predominantly older population could be liable for injuries resulting from this off-the-shelf AI/ML system since it will likely provide recommendations that do not conform to the current standard of care for that age group.

Vicarious Liability. Health systems, physician groups, and physician-employers can be liable for the actions of their employees or affiliates.^{17,32} Vicarious liability differs from other negligence in that one is liable for someone else's actions. For example, whereas negligent credentialing means that the hospital itself has been negligent, vicarious liability means the physician is negligent but the hospital is held responsible. One reason for vicarious liability is to distribute costs for injuries among hospitals and groups to compensate a victim.^{41,42}

Vicarious liability is easiest to establish for an employee. For example, a hospital could be sued over the actions of its physician employees for unsafe deployment of an AI/ML algorithm. However, some courts have become more flexible in inferring a relationship that leads to hospital liability in cases that do not involve direct employees.⁴³

To illustrate how vicarious liability might apply in a case involving AI/ML clinical tools, consider a hospital that purchases a sepsis prediction algorithm for the emergency department (ED) or intensive care unit (ICU).⁴⁴ A court could hold the hospital vicariously liable for the negligence of an emergency medicine physician or intensivist who incorrectly interprets an output from the AI/ML system. The rationale for this type of decision would be that patients at risk for sepsis who are seen in the ED or admitted to the ICU usually seek hospital care but do not typically choose their individual ED physician or intensivist; instead, these physicians are usually closely associated with the hospital.^{43,45}

Products Liability

Products liability law related to medical AI/ML products is more unsettled. On the one hand, medical algorithm developers, like the manufacturer of any product, could be liable for injuries that result from poor design, failure to warn about risks, or manufacturing defects.^{17,46} Indeed, if an AI/ML system used by health care practitioners results in

a patient injury, liability for a design defect may implicate issues with the process for inputting data, software code, or output display.

On the other hand, injured patients will face difficulties bringing such claims for medical AI/ML. In many jurisdictions, a patient would need to show that a viable, potentially cost-effective alternate design exists.⁴⁶ Additionally, software—as opposed to tangible items or combinations of software and hardware—does not fit neatly into the traditional liability landscape.^{47–49} A desire to promote innovation has made courts and legislatures reluctant to extend liability to software developers.^{47–49}

As a result, the case law in health care software products liability is inconsistent and sometimes unclear.^{50–52} For instance, one court dismissed a patient's claims against a surgical robot manufacturer because the patient could not show how the robot's error messages and failure caused his particular injuries.⁵⁰ However, another court approved with relatively little discussion a jury award against a developer whose software caused a catheter to persistently ablate heart tissue.^{48,52} This legal uncertainty can lead injured patients to pursue parties other than software developers—including clinicians and health systems—for redress.

Toward a Balanced Liability System

Some clinical software and systems have already failed and caused injury,^{47,53} and future injuries attributable to AI/ML systems are possible and perhaps inevitable. Going forward, courts and policymakers should design liability systems that seek to achieve the following goals:

- balancing liability across the ecosystem;
- avoiding undue burdens on physicians and frontline clinicians; and
- promoting safe AI/ML development and integration.

The precise solution will depend on how much AI/ML innovation one considers optimal, balanced against one's views on compensation for injured patients. The level of liability directly influences the level of development and implementation of clinical algorithms. Increasing liability may discourage physicians, health systems, and designers from developing and implementing these algorithms. Increasing liability may also discourage experimentation because increasing the cost of errors

encourages actors to pursue other ways of improving medical care. In fact, some commentators have observed that the traditional liability system encourages innovation because it accommodates medical advances to foster growth.^{54–57}

Nonetheless, in service of these goals, policymakers must answer several perennial questions of liability design.

Is the Traditional Liability System the Right Way to Process AI/ML Injury Claims?

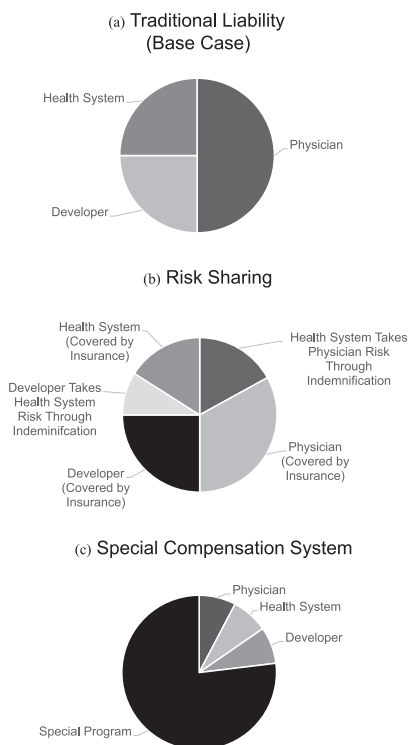
The traditional liability system incentivizes practitioners and stakeholders to invest in care-enhancing activities and safer products (Figure 1).^{49,58,59} The precise distribution of liability among actors varies depending on specialty, practice type, locality, and time. Regardless of the baseline, in the traditional liability system, physicians encounter conflicting signals about how the use of clinical AI/ML systems could affect their liability. On the one hand, existing liability structures could encourage physicians to adopt AI/ML in order to improve diagnosis or prediction and reduce the potential for misdiagnoses in clinical care. On the other hand, physicians may be reluctant to adopt opaque AI/ML systems that expose them to liability if injuries occur.¹⁴

Physicians are not alone in facing conflicting liability signals around AI/ML. Health systems might be held liable for adopting AI/ML systems that they do not have the expertise to fully vet, but they could also be liable for failing to adopt AI/ML systems that improve care.

Developers face other challenges. Because software is traditionally shielded from products liability, developers might introduce AI/ML systems too quickly into clinical care, offloading liability onto physicians and health care systems. However, inconsistent or unclear legal decisions against developers could stifle innovation. Furthermore, liability provides just one set of incentives for developers: regulators, payers, and other stakeholders introduce others.^{60,61}

One solution to the challenges of the traditional liability system that requires no legislative or regulatory action is changing the standard of care. In articulating the standard of care used in the traditional liability system (Figure 1A), courts often look to professional norms.^{55,59} Thus, stakeholders—acting separately or in concert—could alter the standard of care by encouraging the rigorous evaluation of AI/ML algorithms

Figure 1. Models of Liability in Artificial Intelligence and Machine Learning



The precise division of liability in the figure models is arbitrary and not meant to reflect empirical liability, which varies by specialty, practice type, and locality, among other factors. In some cases, the precise distribution of liability is not known for a particular group due to the fragmentary nature of data in this area (e.g., settlements under individual insurance companies with nondisclosure agreements). Model A shows the division of liability under the traditional liability model. In model B, stakeholders enter into contracts to transfer and shift liability among themselves. These agreements can take the form of indemnification (one party assuming some or all of the liability of another) or insurance (spreading risk over policyholders), among others. Further, indemnification agreements can be insured (not shown). In model C, a legislature exempts AI/ML partially or completely from the traditional liability system and the government takes over all or some of the risk. Often, these systems rely on taxes or fees on relevant stakeholders. These modifications do not exist in isolation. For instance, if the legislature enacts a program to shield stakeholders from most of the liability risk (model C), stakeholders could still purchase insurance on their residual risk (model B).

in prospective trials and pressing hospitals and medical practices to appropriately vet AI/ML before implementation.⁶² All stakeholders—physicians, health systems, AI/ML developers, and others—could affect their desired reform by changing their behavior and convincing their peers. Some commentators have already begun to develop frameworks to assess models and algorithms.^{63–65} This process will both smooth clinical integration of AI/ML systems and encourage the development of standardized safety approaches, reducing mismatched incentives and protecting patients.

Further, stakeholders can shift liability to other stakeholders or share liability with them. Because injured patients seek recovery from as many stakeholders as possible, liability can sometimes be a zero-sum game: placing liability on one stakeholder can release others from liability (Figure 1B). Although courts and traditional liability principles provide some guidance on dividing liability, stakeholders can use contracts to provide more certainty before a lawsuit or to promote other objectives. Some people might argue that tort liability should be focused (at least in part) on imposing safety costs on those most able to bear them. However, we recognize that some actors may wish to promote AI/ML system adoption or innovation by using contracts to shift liability burdens. As such, physicians, health systems, and AI/ML designers might choose to reallocate liability by using contractual indemnification clauses, which allow one party to compensate another in litigation.⁶⁶ Insurance provides a similar function by spreading liability and imposing certain safety norms across policyholders, a practice that is familiar to physicians. The precise division of liability would depend on the negotiating and market power of the stakeholders; of course, some level of liability on all actors is desirable to promote safety. To facilitate this type of risk sharing in AI/ML, stakeholders and their trade/professional societies can provide standard indemnification clauses and can partner with insurance companies to offer insurance plans and risk pools.

Should Specialized Adjudication Systems Be Used?

Legislatures could exempt AI/ML product use from the traditional liability system through specialized adjudication systems. Although such changes are difficult because they require concerted political action,

there are precedents for this strategy. For example, Florida and Virginia have neonatal injury compensation programs,⁶⁷ and Congress enacted a similar program for vaccine manufacturers.⁶⁸ These programs tend to collect revenue from a large group, provide relief to specific groups, streamline adjudication, and provide compensation to more individuals than traditional litigation.^{67,69} For instance, the Florida Birth-Related Neurological Injury Compensation Program collects fees from physicians and taxes each birth to provide a specialized adjudication system that may reduce practice costs for obstetricians while compensating patients for injury.⁶⁷ However, specialized adjudication systems may have unintended consequences.^{70,71} In Florida, jurisdictional technicalities have provoked duplicative litigation or placed liability on hospitals for physician errors.^{72,73}

The need for specialized adjudication systems will likely depend on the type of AI/ML deployed. Two algorithm characteristics can assist in making this determination: “opacity” and “plasticity.”⁷⁴ Opacity refers to the user or designer’s ability to understand how an algorithm weighs inputs in determining an answer, with fully opaque systems not being understandable by any observer. Plasticity refers to the algorithm’s ability to self-learn and change based on inputs, with highly plastic systems dynamically changing its weightings based on every new input.

On the one extreme, a non-updating algorithm based on a fixed, limited number of variables with known weights—i.e., an algorithm that is neither plastic nor opaque—is akin to a clinical scoring tool and may not require special treatment.⁵⁷ This is not to say that such an algorithm is error proof, but it can be assessed and checked by physicians, health systems, and payers independently and transparently. Traditional liability systems are a better fit to address errors arising from closed, non-updating algorithms because these algorithms are more easily assessed by actors.

At the other extreme, a “black-box” algorithm—i.e., one that is both plastic and opaque—poses more difficult questions. A black-box algorithm is self-learning and generally cannot provide the precise reasons for a particular recommendation.^{18,75,76} The traditional liability system could respond by imposing a duty on designers, hospital systems, or other stakeholders to test or retrain such algorithms or perhaps to engage in some forms of adversarial stress testing, a well-established form of assessment in cybersecurity.^{74–76} These stress tests would “stress” the algorithm with diagnostic or therapeutic challenges to assess the

algorithm's responses to difficult situations not necessarily envisioned by its designers.⁷⁴ That said, it is a formidable challenge to craft an administrable standard of negligence for such cases that would give algorithm makers fair warning about how often they should test or retest in such a case, especially given the heterogeneous nature of algorithm design and patient populations. Moreover, at the level of physician liability, black-box algorithms may not be a good fit for traditional liability theories because the actors potentially paying for algorithm-driven mistakes—physicians, health systems, or even algorithm designers—may not be able to avoid making these mistakes when using an opaque and plastic system for a particular patient.^{18,74,75} Some may argue that it is appropriate for liability risks to deter adoption of black-box algorithms, as a way to encourage transparency of algorithms. However, the positive contribution of these black-box algorithms arises from their ability to constantly update and improve prediction without human intervention across ever more diverse clinical scenarios.

Is there a way to design liability regimes that do not deter promising, beneficial clinical algorithms while also ensuring their appropriate design and deployment? For at least a subset of these algorithms, policymakers could consider more fundamental legal reforms: Courts (or, more plausibly, legislatures) could depart from the traditional liability system or change the standard of proof to distribute liability for a patient injury equitably between all stakeholders.^{74,75} Legislatures may choose to go further and create special agencies to consider claims. These special adjudication systems could develop the expertise to adjudicate AI/ML liability among stakeholders. Alternatively, legislatures could create a compensation program that does not consider liability (no-fault systems), instead assessing fees on stakeholders (e.g., a fee per patient affected by an algorithm or a fee on physicians/practitioners) (Figure 1C). In the black-box context, these no-fault systems would have the added advantage of avoiding the difficult question of what precisely caused an error. Indeed, policymakers could fine-tune the program depending on the desire to promote AI/ML, ensure safety through manageable personal liability, allocate responsibility among stakeholders, and facilitate other public policy goals. Whether such alterations to liability are politically feasible in a world where tort reform is politically fraught remains an open question.

How Do Regulators Fit into Liability?

AI/ML systems hold lessons for regulators. The National Institutes of Health and several radiology professional societies have collaborated on an agenda to foster AI/ML innovation that includes identifying knowledge gaps and facilitating processes to validate new algorithms.⁷⁷ These efforts indirectly affect liability by providing potential benchmarks that courts and regulators can use to assess a particular system.

However, in another vein, the FDA has published several guidance and discussion documents that treat AI/ML systems as devices.^{78,79} These regulatory efforts also provide information to the liability system about how to assess a particular AI/ML system.

Additionally, binding regulatory actions by the FDA and other agencies can directly affect the liability system. The FDA's decision to regulate clinical AI/ML tools as devices has important—and unsettled—implications for whether regulation displaces traditional liability systems.^{80–83} Although FDA regulation will likely not directly affect physician or health system liability, developer liability could decrease with more stringent regulatory requirements. This could displace or alter state liability reforms (altering the standard of care or special adjudication systems). Decreased developer liability may also have indirect implications for physician and health system liability under a system of shared liability among stakeholders. Furthermore, FDA regulation might, at least in theory, also work to ensure that AI/ML training populations are representative of patient populations, akin to guidance requiring clinical trials to enroll participants from groups that have been underrepresented in biomedical research studies, including racial minorities and older adults.⁸⁴

So far, the FDA has acted cautiously in this area, instituting incremental reforms or pilot programs to learn more. For instance, the Software Precertification Pilot Program has identified particular organizations or divisions of companies that are committed to the safety and efficacy of their various software applications and is studying how a reformed regulatory process might apply to submissions from these types of organizations. Formally, the FDA's overarching principles for the program are patient safety, product quality, clinical responsibility, cybersecurity responsibility, and proactive culture.⁸⁵ If a precertification program is instituted, the FDA envisions that certain applications from precertified entities might undergo a streamlined approval process or require no

FDA review at all. Although such a program may not formally displace state liability schemes, the practices of FDA precertified entities would likely inform courts, trade organizations, policymakers, and others in developing standards of conduct and safety.

Conclusion

AI/ML systems have the potential to radically transform clinical care. The legal system moves at a slower pace, but it cannot be static with regard to this innovation. The relatively unsettled state of AI/ML and its potential liability provide an opportunity to develop a new liability model that accommodates medical progress and instructs stakeholders on how best to respond to disruptive innovation. To fully realize the benefits of AI/ML, the legal system must balance liability to promote innovation, safety, and accelerated adoption of these powerful algorithms.

References

1. National Institutes of Health. NIH workshop: harnessing artificial intelligence and machine learning to advance biomedical research. July 23, 2018. https://datascience.nih.gov/sites/default/files/AI_workshop_report_summary_01-16-19_508.pdf. Accessed February 3, 2021.
2. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. *Nat Med*. 2019;25(1):44-56. <https://doi.org/10.1038/s41591-018-0300-7>.
3. The Medical Futurist. FDA approvals for smart algorithms in medicine in one giant infographic. June 6, 2019. <https://medicalfuturist.com/fda-approvals-for-algorithms-in-medicine>. Accessed February 3, 2021.
4. Hwang TJ, Kesselheim AS, Vokinger KN. Lifecycle regulation of artificial intelligence- and machine learning-based software devices in medicine. *JAMA*. 2019;322(23):2285. <https://doi.org/10.1001/jama.2019.16842>.
5. Abramoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *NPJ Digital Med*. 2018;1(1):39. <https://doi.org/10.1038/s41746-018-0040-6>.
6. Perez MV, Mahaffey KW, Hedlin H, et al. Large-scale assessment of a smartwatch to identify atrial fibrillation. *N*

- Engl J Med.* 2019;381(20):1909-1917. <https://doi.org/10.1056/NEJMoa1901183>.
7. Mills TT, Food and Drug Administration 510(k) clearance of Zebra Medical Vision HealthPNX device (letter). May 6, 2019. https://www.accessdata.fda.gov/cdrh_docs/pdf19/K190362.pdf. Accessed February 3, 2021.
 8. Ström P, Kartasalo K, Olsson H, et al. Artificial intelligence for diagnosis and grading of prostate cancer in biopsies: a population-based, diagnostic study. *Lancet Oncol.* 2020;21(2):222-232. [https://doi.org/10.1016/S1470-2045\(19\)30738-7](https://doi.org/10.1016/S1470-2045(19)30738-7).
 9. Pennic F. Biofourmis analytics engine receives FDA clearance for ambulatory physiologic monitoring. HIT Consultant blog. October 4, 2019. <https://hitconsultant.net/2019/10/04/biofourmis-biovitals-analytics-engine-receives-fda-clearance/#.Xi3xOesnq0>. Accessed February 3, 2021.
 10. Parikh RB, Teeple S, Navathe AS. Addressing bias in artificial intelligence in health care. *JAMA.* 2019;322(24):2377. <https://doi.org/10.1001/jama.2019.18058>.
 11. Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science.* 2019;366(6464):447-453. <https://doi.org/10.1126/science.aax2342>.
 12. Challen R, Denny J, Pitt M, Gompels L, Edwards T, Tsaneva-Atanasova K. Artificial intelligence, bias and clinical safety. *BMJ Qual Saf.* 2019;28(3):231-237. <https://doi.org/10.1136/bmjqs-2018-008370>.
 13. Gianfrancesco MA, Tamang S, Yazdany J, Schmajuk G. Potential biases in machine learning algorithms using electronic health record data. *JAMA Intern Med.* 2018;178(11):1544-1547. <https://doi.org/10.1001/jamainternmed.2018.3763>.
 14. Price WN, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *JAMA.* 2019;322(18):1765. <https://doi.org/10.1001/jama.2019.15064>.
 15. Parasidis E. Clinical decision support: elements of a sensible legal framework. *J Health Care Law Pol.* 2018;20(2):183.
 16. Kesselheim AS, Cresswell K, Phansalkar S, Bates DW, Sheikh A. Clinical decision support systems could be modified to reduce 'alert fatigue' while still minimizing the risk of litigation. *Health Aff (Millwood).* 2011;30(12):2310-2317. <https://doi.org/10.1377/hlthaff.2010.1111>.
 17. Ridgely MS, Greenber MD. Too many alerts, too much liability: sorting through the malpractice implications of drug-drug interaction clinical support. *St Louis U J Health Law Pol.* 2012;5(2):257-295.

18. Price WN. Medical malpractice and black-box medicine. In: Cohen IG, Lynch HF, Vayena E, Gasser U, eds. *Big Data, Health Law, and Bioethics*. New York: Cambridge University Press; 2018:295-306.
19. Mehlman MJ. Medical practice guidelines as malpractice safe harbors: illusion or deceit? *J Law Med Ethics*. 2012;40(2):286-300. <https://doi.org/10.1111/j.1748-720X.2012.00664.x>.
20. *Tesauro v Perrige*, 650 A2d, 1079 (Pa Super Ct 1994).
21. *Roman v City of New York*, 442 NYS, 945 (NY Sup Ct 1981).
22. *Appleby v Miller*, 554 NE.2d, 773 (Ill App Ct 1990).
23. *Bailey v Huggins Diagnostic & Rehabilitation Ctr*, 952 P2d, 768 (Colo App Ct 1997).
24. *Smith v Linn*, 563 A2d, 123 (Pa Super Ct 1989).
25. *Ross v Jacobs*, 684 P2d, 1211 (Okla App Ct 1984).
26. *Bristol-Myers v Gonzales*, 561 SW2d, 801 (Tex 1978).
27. *Sacher v Long Island Jewish-Hillside Med Ctr*, 530 NYS2d, 232 (App Div Ct 1988).
28. *Conn v United States*, 880 F Supp 2d, 741 (SD Miss 2012).
29. *Toth v Cmty Hosp at Glen Cove*, 239 NE2d, 368 (NY 1968).
30. *Sander v Geib, Elston, Frost Pro Assn*, 506 NW2d, 107 (SD 1993).
31. *Bush v Thoratec Corp*, 13 F Supp 3d, 554 (ED La 2014).
32. Abraham KS, Weiler PC. Enterprise medical liability and the evolution of the american health care system. *Harvard Law Rev*. 1994;108(2):381-436. <https://doi.org/10.2307/1341896>.
33. *S Highlands Infirmary v Camp*, 180 So2d, 904 (Ala1965).
34. Shipley WE. Hospital's liability to patient for injury sustained from defective equipment furnished by hospital for use in diagnosis or treatment of patient. 14 ALR3d, 1254 (1967).
35. *Berg v United States*, 806 F2d, 978 (10th Cir 1986).
36. Whelton PK, Carey RM, Aronow WS, et al. 2017 ACC/AHA/AAPA/ABC/ACPM/AGS/APhA/ASH/ASPC/NMA/PCNA guideline for the prevention, detection, evaluation, and management of high blood pressure in adults: executive summary: a report of the American College of Cardiology/American Heart Association Task Force on Clinical Practice Guidelines. *J Am Coll Cardiol*. 2018;71(19):2199-2269. <https://doi.org/10.1016/j.jacc.2017.11.005>.
37. Benetos A, Petrovic M, Strandberg T. Hypertension management in older and frail older patients. *Circ Res*. 2019;124(7):1045-1060. <https://doi.org/10.1161/CIRCRESAHA.118.313236>.
38. Naranjo M, Mezue K, Rangaswami J. Outcomes of intensive blood pressure lowering in older hypertensive patients. *J Am*

- Coll Cardiol.* 2017;70(1):119-120. <https://doi.org/10.1016/j.jacc.2017.02.078>.
39. Benetos A, Labat C, Rossignol P, et al. Treatment with multiple blood pressure medications, achieved blood pressure, and mortality in older nursing home residents: the PARTAGE study. *JAMA Intern Med.* 2015;175(6):989-995. <https://doi.org/10.1001/jamainternmed.2014.8012>.
 40. James PA, Oparil S, Carter BL, et al. 2014 evidence-based guideline for the management of high blood pressure in adults: report from the panel members appointed to the eighth joint national committee (JNC 8). *JAMA.* 2014;311(5):507-520. <https://doi.org/10.1001/jama.2013.284427>.
 41. *Rose v Ryder Truck Renter*, 662 F Supp, 177 (D Me 1987).
 42. Overton PR. Rule of “Respondeat Superior.” *JAMA.* 1957;163(10):847-852.
 43. Alstott A. Hospital liability for negligence of independent contractor physicians under principles of apparent agency. *J Leg Med.* 2004;25(4):485-504. <https://doi.org/10.1080/01947640490887607>.
 44. Mao Q, Jay M, Hoffman JL, et al. Multicentre validation of a sepsis prediction algorithm using only vital sign data in the emergency department, general ward and ICU. *BMJ Open.* 2018;8(1):e017833. <https://doi.org/10.1136/bmjopen-2017-017833>.
 45. *Mcrosky v Carson Tahoe Regional Med Ctr*, 408 P3d, 149 (Nev 2017).
 46. Chagal-Feferkorn KA. Am I an algorithm or a product? When products liability should apply to algorithmic decision-makers. *Stanford Law Pol Rev.* 2019;30(1):61-114.
 47. Brown SH, Miller RA. Legal and regulatory issues related to the use of clinical software in health care delivery. In: Greenes, R, ed. *Clinical Decision Support*, 2nd ed. New York: Elsevier; 2014:711-740.
 48. Choi BH. Crashworthy code. *Washington Law Rev.* 2019;94(1):39-117.
 49. Terry NP. When the “machine that goes ‘ping’” causes harm: default torts rules and technologically mediated health care injuries. *St Louis U Law J.* 2002;46(1):37.
 50. *Mracek v Bryn Mawr Hosp*, 610 F Supp 2d, 401 (ED Pa 2009).
 51. *Corley v Stryker Orthopaedics*, No. 13-2571, 2014 WL 3125990 (WD La 2014).
 52. *Singh v Edwards Lifesciences Corp*, 210 P3d, 337 (Wash App Ct 2009).

53. Graber M, Siegal D, Riah H, Johnston D, Kenyon K. Electronic health record–related events in medical malpractice claims. *J Patient Saf.* 2019;15(2):77-85. <https://doi.org/10.1097/PTS.0000000000000240>.
54. Henderson JA. Tort vs. technology: accommodating disruptive innovation. *Ariz St Law J.* 2015;47(4):1145-1180.
55. Laakmann AB. When should physicians be liable for innovation? *Cardozo Law Rev.* 2015;36(3):913-968.
56. Mastroianni AC. Liability, regulation and policy in surgical innovation: the cutting edge of research and therapy. *Health Matrix Clevel.* 2006;16(2):351-442.
57. De Ville K. Medical malpractice in twentieth century United States. The interaction of technology, law and culture. *Int J Tech Assess Health Care.* 1998;14(2):197-211. <https://doi.org/10.1017/s0266462300012198>.
58. Stein A. The domain of torts. *Columbia Law Rev.* 2017;117(3):535-611.
59. Froomkin AM, Kerr I, Pineau J. When AIs Outperform doctors: confronting the challenges of a tort-induced over-reliance on machine learning. *Ariz Law R.* 2019;61(1):33-100.
60. Panch T, Mattie H, Celi LA. The “inconvenient truth” about AI in healthcare. *NPJ Digit Med.* 2019;2(1):77. <https://doi.org/10.1038/s41746-019-0155-4>.
61. Syverud KD. The duty to settle. *Va Law Rev.* 1990;76(6):1113-1209. <https://doi.org/10.2307/1073190>.
62. Parikh RB, Obermeyer Z, Navathe AS. Regulation of predictive analytics in medicine: algorithms must meet regulatory standards of clinical benefit. *Science.* 2019;363(6429):810-812. <https://doi.org/10.1126/science.aaw0029>.
63. Khalifa M, Magrabi F, Gallego B. Developing a framework for evidence-based grading and assessment of predictive tools for clinical decision support. *BMC Med Inform Decis Mak.* 2019;19(1):207. <https://doi.org/10.1186/s12911-019-0940-7>.
64. Localio AR, Stack CB. TRIPOD: a new reporting baseline for developing and interpreting prediction models. *Ann Intern Med.* 2015;162(1):73-74. <https://doi.org/10.7326/M14-2423>.
65. Moons KGM, Wolff RF, Riley RD, et al. PROBAST: a tool to assess risk of bias and applicability of prediction model studies: explanation and elaboration. *Ann Intern Med.* 2019;170(1):51-58. <https://doi.org/10.7326/M18-1376>.
66. Parker PL, Slavich J. Contractual efforts to allocate the risk of environmental liability: is there a way to make indemnities

- worth more than the paper they are written on? *SMU Law Rev.* 1991;44(4):1349-1381.
67. Horwitz J, Brennan TA. No-fault compensation for medical injury: a case study. *Health Aff (Millwood)*. 1995;14(4):164-179. <https://doi.org/10.1377/hlthaff.14.4.164>.
 68. *Schafer v Am Cyanamid Co*, 20 F3d, 1 (1st Cir Ct 1994).
 69. Whetten-Goldstein K, Kulas E, Sloan F, Hickson G, Entman S. Compensation for birth-related injury: no-fault programs compared with tort system. *Arch Pediatr Adolesc Med.* 1999;153(1):41-48. <https://doi.org/10.1001/archpedi.153.1.41>.
 70. *Moss v Merck & Co*, 381 F3d, 501 (5th Cir Ct 2004).
 71. *Doe 2 v Ortho-Clinical Diagnostics*, 335 F Supp 2d, 614 (MDNC 2004).
 72. Studdert DM, Fritz LA, Brennan TA. The jury is still in: Florida's birth-related neurological injury compensation plan after a decade. *Jo Health Polit Policy Law.* 2000;25(3):499-526. <https://doi.org/10.1215/03616878-25-3-499>.
 73. *Univ of Miami v Ruiz*, 164 So3d, 758 (Fla App Ct 2015).
 74. Babic B, Gerke S, Evgeniou T, Cohen IG. Algorithms on regulatory lockdown in medicine. *Science.* 2019;366(6470):1202-1204. <https://doi.org/10.1126/science.aay9547>.
 75. Price N. Artificial intelligence in health care: applications and legal issues. *SciTech Lawyer.* 2017;14(1):10-13.
 76. Vladeck DC. Machines without principals: liability rules and artificial intelligence. *Wash Law Rev.* 2014;89(1):117-150.
 77. Langlotz CP, Allen B, Erickson BJ, et al. A roadmap for foundational research on artificial intelligence in medical imaging: from the 2018 NIH/RSNA/ACR/The Academy workshop. *Radiology.* 2019;291(3):781-791. <https://doi.org/10.1148/radiol.2019190613>.
 78. US Food and Drug Administration. Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD). April 2, 2019. <https://www.fda.gov/media/122535/download>. Accessed February 3, 2021.
 79. US Food and Drug Administration. Software as a medical device (SaMD): clinical evaluation: guidance for industry and Food and Drug Administration staff. December 8, 2017. <https://www.fda.gov/media/100714/download>. Accessed February 3, 2021.
 80. *Riegel v Medtronic*, 552 US, 312 (2008).
 81. Boumil MM. FDA approval of drugs and devices: preemption of state laws for "parallel" tort claims. *J Health Care Law Policy.* 2015;18(1):1-42.

82. *Medtronic v Lohr*, 518 US, 470 (1996).
83. Geistfeld MA. Tort law in the age of statutes. *Iowa Law Rev.* 2014;99(3):957-1020.
84. US Food and Drug Administration. Enhancing the diversity of clinical trial populations—eligibility criteria, enrollment practices, and trial designs: guidance for industry. November 2020. <https://www.fda.gov/media/127712/download>. Accessed February 3, 2021.
85. US Food and Drug Administration. Precertification (Pre-Cert) Pilot Program: Frequently Asked Questions. September 14, 2020. <https://www.fda.gov/medical-devices/digital-health-software-precertification-pre-cert-program/precertification-pre-cert-pilot-program-frequently-asked-questions>. Accessed February 3, 2021.

Funding/Support: IGC and SG were supported by a grant from the Collaborative Research Program for Biomedical Innovation Law, a scientifically independent collaborative research program supported by a Novo Nordisk Foundation grant (NNF17SA0027784).

Disclosure: This article is written for academic purposes only and should not be taken as legal advice.

Conflict of Interest Disclosure: All authors completed the ICMJE Form for Disclosure of Potential Conflicts of Interest. IGC reports that he serves as a bioethics consultant for Otsuka for its Abilify MyCite product, outside the submitted work. RP reports personal fees from GNS Healthcare, Nanology, and the Cancer Study Group; an honorarium from Medscape; and grants from the Veterans Health Administration, the National Palliative Care Research Center, the Prostate Cancer Foundation, the National Cancer Institute, the Conquer Cancer Foundation, and the Medical University of South Carolina, outside the submitted work.

Address correspondence to: Ravi B. Parikh, MD, MPP, 423 Guardian Dr, Blockley 1102, Philadelphia, PA 19104 (email: ravi.parikh@penmedicine.upenn.edu).