

3-14-2021

Who Will Be Liable for Medical Malpractice in the Future? How the Use of Artificial Intelligence in Medicine Will Shape Medical Tort Law

Scott J. Schweikart

Follow this and additional works at: <https://scholarship.law.umn.edu/mjlst>



Part of the [Analytical, Diagnostic and Therapeutic Techniques and Equipment Commons](#), [Biomedical Commons](#), [Patient Safety Commons](#), and the [Torts Commons](#)

Recommended Citation

Scott J. Schweikart, *Who Will Be Liable for Medical Malpractice in the Future? How the Use of Artificial Intelligence in Medicine Will Shape Medical Tort Law*, 22 MINN. J.L. SCI. & TECH. 1 ().

Available at: <https://scholarship.law.umn.edu/mjlst/vol22/iss2/2>

Who Will Be Liable for Medical Malpractice in the Future? How the Use of Artificial Intelligence in Medicine Will Shape Medical Tort Law

Scott J. Schweikart*

ABSTRACT

Artificial intelligence (AI) is a powerful technology that can assist physicians with the practice of medicine. Use of the technology has grown in recent years and has powerful potential. Medical AI typically functions as a type of “machine-learning” that relies on deep neural networks to sift through vast amounts of data to give recommendations or draw conclusions for clinicians. This Article begins by outlining key characteristics of medical AI (e.g., AI’s opacity and “black-box,” and how and with what data the AI was developed) that make assessment of liability under traditional tort paradigms (like negligence) difficult. This Article then highlights several tort paradigms (e.g., medical malpractice, products liability, vicarious liability, and informed consent) and how they might function in the context of medical AI. In conclusion, this Article offers an analysis of how tort law may evolve in the future in response to the challenges created by medical AI today (e.g., legal evolutions that may take the form of new solutions like AI personhood, common enterprise liability, and a new standard of care).

©2021 Schweikart

* Senior Research Associate at the American Medical Association and Legal Editor of the AMA Journal of Ethics. M.B.E., University of Pennsylvania; J.D., Case Western Reserve University; B.A., Washington University in St. Louis. I would like to thank the University of Minnesota Law School for inviting me to participate in their symposium “The Law and Policy of Artificial Intelligence in Health Care.”

I. Introduction	3
II. Nature of AI.....	6
A. Opacity and the “Black-Box”.....	6
B. Development and Control.....	8
III. Challenges in Applying to AI to Tort Law	10
A. Medical Malpractice (Negligence).....	11
B. Products Liability and the Learned Intermediary Doctrine	14
C. Vicarious Liability.....	15
D. Informed Consent Liability	16
IV. How Tort Law May Evolve.....	18
A. AI Personhood	19
B. Common Enterprise Liability	19
C. New Standard of Care	21
V. Conclusion.....	22

I. INTRODUCTION

Artificial intelligence (AI)¹ technology is increasingly used by a large percentage of the healthcare industry.² AI in the medical context is typically a type known as “machine learning,”³ a

1. Artificial Intelligence is also called augmented intelligence (or intelligence augmentation), assisted intelligence, or autonomous intelligence. Sometimes the terms are used interchangeably, but there are distinctions. See Kathleen Walch, *Is There a Difference Between Assisted Intelligence Vs. Augmented Intelligence?*, FORBES (Jan. 12, 2020), <https://www.forbes.com/sites/cognitive-world/2020/01/12/is-there-a-difference-between-assisted-intelligence-vs-augmented-intelligence>. Walch notes the differences in types of artificial intelligence as “a continuum of human-machine intelligence interaction ranging from situations where machines are basically repeating many of the tasks humans are already doing (assisted) to enabling humans to do more than they are currently capable of doing (augmented) to fully accomplishing tasks on their own without human intervention (autonomous).” *Id.* See also Aaron Masih, *Augmented Intelligence, Not Artificial Intelligence, Is the Future*, MEDIUM (Jan. 19, 2019), <https://medium.com/datadriveninvestor/augmented-intelligence-not-artificial-intelligence-is-the-future-f07ada7d4815> (“While the underlying technologies powering AI [artificial intelligence] and IA [intelligence augmentation] are the same, the goals and applications are fundamentally different: AI aims to create systems that run without humans, whereas IA aims to create systems that make humans better. To be clear, this is not a separate category of technology, but simply a different way of thinking about its purpose.”). For purposes of this Article, I use the term artificial intelligence to refer generally to all types of AI, i.e., assisted, augmented, or autonomous. However, the term most often used in the medical context—and hence used in most of the examples discussed in this paper—is augmented intelligence, i.e., a technology assisting or enhancing a physician’s practice of medicine, rather than autonomous AI technology that takes over or supersedes a physician’s practice.

2. TATA CONSULTANCY SERVICES, GETTING SMARTER BY THE SECTOR: HOW 13 GLOBAL INDUSTRIES USE ARTIFICIAL INTELLIGENCE 42 (Nov. 21, 2015), <https://www.tcs.com/global-trend-studies> (“Some 86% [of healthcare companies surveyed] are currently using cognitive technologies; of the 14% that don’t, all plan to do so by 2020.”); Bill Siwicki, *86% of Healthcare Companies Use Some Form of AI*, HEALTHCARE IT NEWS (May 19, 2017), <https://www.healthcareitnews.com/news/86-healthcare-companies-use-some-form-ai> (“About 86 percent of healthcare provider organizations, life science companies and technology vendors are currently using artificial intelligence technology.”).

3. Michael E. Matheny et al., *Artificial Intelligence in Health Care: The Hope, the Hype, the Promise, the Peril*, in ARTIFICIAL INTELLIGENCE IN HEALTH CARE: THE HOPE, THE HYPE, THE PROMISE, THE PERIL 15 (Michael Matheny, Sonoo Thadaney-Israni, Mahnoor Ahmed, & Danielle Whicher eds., 2019) (“Machine learning is a family of statistical and mathematical modeling techniques that uses a variety of approaches to automatically learn and improve the prediction of a target state, without explicit programming (e.g., Boolean rules).”). Additionally, I will note that this book from the National Academy of Medicine is an excellent compendium regarding issues of AI in health care.

subset of AI where the technology uses algorithms to find patterns in “vast amounts of data.”⁴ One such example of medical machine learning AI is that of Watson for Oncology (produced by IBM) which “uses cognitive computing to ‘interpret cancer patients’ clinical information and identify individualized, evidence-based treatment options.”⁵ How this works is that Watson is “fed reams of information on oncological matters” and then “scour[s] them to present doctors with treatment options, and recommended drugs and instructions for administration.”⁶ The power of such machine learning is that the AI is able to go through vast amounts of data (more data than any one clinician can hope to scan and evaluate) and give relevant input to the physician in a very short period of time,⁷ all while simultaneously finding patterns in the data that are impossible for humans to find or understand.

The fundamental challenge of AI from a medical liability standpoint is assessing liability when some error and injury inevitably occurs. For example, what happens when an AI algorithm recommends one course of action but a physician believes another course of action, based on their own substantial expertise, is more prudent?⁸ What happens if the AI recommends a course of treatment and the physician adopts the recommendation, but the AI was wrong and the patient is injured? These questions are challenging because it is often impossible to understand how the AI made its decision; its reasoning is sealed off

4. W. Nicholson Price II, *Artificial Intelligence in Health Care: Applications and Legal Implications*, 14 SCITECH LAW. 10, 10 (2017) (“[D]ata come from many sources: electronic health records, medical literature, clinical trials, insurance claims data, pharmacy records, and even the data entered by patients into their smartphones or recorded on fitness trackers.”).

5. Jason Chung, *What Should We Do About Artificial Intelligence in Health Care?*, 22 HEALTH L. J. 37, 37 (2017).

6. *Id.* at 37 (“Watson does not seek answers to problems on its own but can draw upon disparate data sources to synthesize potential solutions.”).

7. Jason Chung & Amanda Zink, *Hey Watson – Can I Sue You for Malpractice?— Examining the Liability of Artificial Intelligence in Medicine*, 11 ASIA PAC. J. HEALTH L. & ETHICS 51, 57 (2018) (“[AI] can parse through unfathomable amounts of data, rank its quality, and integrate new information into its adaptive programming. It can do so at a rate no human can ever hope to match.”).

8. Danny Tobey & Allie Cohen, *Medical Frontiers in AI Liability*, 24 AHLA CONNECTIONS 22, 22 (Feb. 2020). Tobey and Cohen note the key question when a physician is confronted with such a conundrum: “Is the machine malfunctioning or is it seeing further?” *Id.* The answer to this question is often difficult (sometimes impossible) to know but is critical in understanding liability.

in the so-called “black-box.”⁹ If the reasoning of how an AI makes its determination cannot be understood, assigning liability may be near impossible under the traditional tort paradigms. For example, how can a breach of duty be established if the reasoning of the AI is unknown? If the reasoning is unknowable, then it may be impossible to determine if the AI’s reasoning was sound or unsound with regard to a breach of tort duty. Additionally, when should a physician be duty bound to rely or not rely on an AI recommendation? The question is especially fraught when the physician cannot know the reasoning behind the recommendation.

Also, the difficulty in analyzing liability is compounded by the fact that AI in the medical context is typically of the “augmented” variety,¹⁰ meaning that it is more wholly integrated with the physician within the practice of medicine and ultimately functions as a tool to help the physician treat patients.¹¹ The physician functions as an intermediary to the technology and hence still has a key role in the practice of medicine.¹² However, AI also has a role, and the mixed nature of the AI’s role and the physician’s role can make establishing a legal duty or parsing and assigning liability challenging.

This Article aims to outline the aspects of medical negligence (in the American jurisprudential context) most at issue and how they will be impacted by AI. Part II will discuss fundamental aspects about the nature of AI relevant to tort liability.

9. Price, *supra* note 4, at 10. Price notes that AI “algorithms themselves are often too complex for their reasoning to be understood or even stated explicitly. Such algorithms may be best described as a ‘black-box’”, i.e., the AI’s reasoning is completely sealed off from the creators and users of the technology. *Id.*

10. See *supra* text accompanying note 1 for a brief explanation of different types of AI.

11. Matheny et al., *supra* note 3, at 16. As noted in the National Academy of Medicine book about AI, “[c]ombining human intelligence and AI into augmented intelligence focuses on a supportive or assistive role for the algorithms, emphasizing that these technologies are designed to enhance human processing, cognition, and work, rather than replace it.” *Id.*

12. Chung & Zink, *supra* note 7, at 77–78. Authors note the difference between medical AI, like Watson for Oncology, and the AI of self-driving cars which function autonomously. Rather than being truly autonomous, medical AI of the augmented variety “relies on a human intermediary to interact with” and is detach[ed] from administering treatment based on its diagnoses, which means that “difficult moral and practical decisions involving patient care are left to human intermediaries [i.e., physicians and other clinicians].” *Id.* at 78.

Part III will discuss challenges current tort law paradigms have in the context of medical AI. Part IV will analyze how medical tort law may evolve in response to some of these key aspects of AI. Part V will conclude with a discussion of the possible evolution of medical tort law and what considerations may underlie which legal solutions will be chosen by society going forward.

II. NATURE OF AI

The nature of AI and how it is used in medicine is critical in understanding how it will impact tort law. This Part will examine two key characteristics of AI: opaqueness (i.e., the inherent lack of transparency and explainability of AI) and its development and control (i.e., the quality of data used in its creation, the multiple parties and entities involved in the creation of AI, and those parties who also have dominion over its function, existence, and use).

A. OPAQUENESS AND THE “BLACK-BOX”

Most AI used in the medical context functions with a characterization known as the “black-box.” Nicholson Price describes the “black-box” in the healthcare context—what he calls “black-box medicine”—as the use of opaque computational models to make decisions related to healthcare.¹³ He explains that “[a] defining feature of black-box medicine is that those algorithms are non-transparent—that is, the relationships they capture cannot be explicitly understood, and sometimes cannot even be explicitly stated.”¹⁴ The phenomenon of the “black-box” is a product of machine learning known as “deep learning,” so-called from the machine’s vast “deep” neural network.¹⁵ The power of “deep” machine learning is becoming more and more evident. For example,

13. W. Nicholson Price II, *Black-Box Medicine*, 28 HARV. J.L. & TECH. 419, 421 (2015) (emphasizing that “this type of medicine is ‘black-box’ to everyone by nature of its development; it is not ‘black-box’ because its workings are deliberately hidden from view”).

14. *Id.*

15. Will Knight, *The Dark Secret at the Heart of AI*, MIT TECH. REV., (Apr. 11, 2017), <https://www.technologyreview.com/2017/04/11/51113/the-dark-secret-at-the-heart-of-ai/>. Knight provides an excellent synopsis on how AI’s “deep learning” works and how “by its nature” it is a “particularly dark black-box.” *Id.* As Knight explains, “[y]ou can’t just look inside a deep neural network to see how it works. A network’s reasoning is embedded in the behavior of thousands of simulated neurons, arranged into dozens or even hundreds of intricately interconnected layers.” *Id.*

an AI known as Deep Patient was trained using data from about 700,000 individuals, and when testing on new records, it proved incredibly good at predicting disease. Without any expert instruction, Deep Patient had discovered patterns hidden in the hospital data that seemed to indicate when people were on the way to a wide range of ailments, including cancer of the liver.¹⁶

Additionally, opaqueness as a problem with AI goes beyond simple “black-box” opacity and hinges more broadly on the way AI is developed. Matthew Sherer explains how the research and development of AI can manifest itself in four ways (i.e., discreetness, diffuseness, discreteness, and opacity) that impact the technology’s notable characteristic of lacking transparency and explainability:

Discreetness refers to the fact that AI development work can be conducted with limited visible infrastructure. *Diffuseness* means that the individuals working on a single component of an AI system might be located far away from one another. A closely related feature, *discreteness*, refers to the fact that the separate components of an AI system could be designed in different places and at different times without any conscious coordination. Finally, *opacity* denotes the possibility that the inner workings of an AI system may be kept secret and may not be susceptible to reverse engineering.¹⁷

Indeed, there is fear of the “black-box” and its opacity. Some clinicians may be wary about implementing this technology into their medical practice if doing so means relying on the AI’s judgment without a full understanding of its reasoning. However, some argue that there is nothing truly to fear about AI’s opacity. Vijay Pande notes that the “black-box” is also endemic of human intelligence, in that “[h]uman intelligence can reason and make arguments for a given conclusion, but can’t explain the complex, underlying basis for how we arrived at a particular conclusion.”¹⁸

16. *Id.* Knight interviewed a clinician who had experience using Deep Patient. The clinician noted that, while he is impressed that the AI can successfully anticipate diseases and disorders, “we [still] don’t know how they [i.e., AI] work,” leading users to lack complete assurance in the technology’s accuracy.

17. Matthew U. Scherer, *Regulating Artificial Intelligence Systems: Risks, Challenges, Competencies, and Strategies*, 29 HARV. J.L. & TECH. 353, 369 (2016).

18. Vijay Pande, *Artificial Intelligence’s ‘Black Box’ Is Nothing to Fear*, N.Y. TIMES (Jan. 25, 2018), <https://www.nytimes.com/2018/01/25/opinion/artificial-intelligence-black-box.html> (“Perhaps that real source of the critics’ concerns isn’t that we can’t ‘see’ A.I.’s reasoning but that as A.I. gets more powerful, the

Indeed, it may just be the nature of intelligence—human or artificial—“that only a part of it is exposed to rational explanation.”¹⁹

B. DEVELOPMENT AND CONTROL

The development and control of AI is another characteristic of the technology. A key aspect of AI, as it relates to medical tort law, is that who is in control of the technology and in control of how the technology is created (e.g., the quality of data inputted during development), may influence liability across a spectrum of different parties and additionally influence the effectiveness and proper functioning of AI when deployed.²⁰

For example, a fundamental aspect of AI is that it takes many different parties and entities to create it, e.g., many different hardware and software developers. And as explained above, there is an element of *diffuseness* (i.e., developers working on AI might be located far away from each other and not working tightly in conjunction) and *discreteness* (i.e., different parts of the AI are created at different times and locations with potentially no coordination).²¹ The diffuseness and discreteness in the development of AI make questions about assigning liability difficult, as issues of who “controls” the technology become difficult to parse, especially when also factoring in that the end-user of the AI (i.e., a hospital or physician) could also be a party who also potentially has control. Hence, answers related to agency and tort duties may be murky.

Additionally, the quality of data used in the development of medical AI is also of profound concern as “bad data quality adversely affects patient care and outcomes.”²² Experts note that

human mind becomes the limiting factor . . . Doctors will no longer ‘drive’ the primary diagnosis; instead, they’ll ensure that the diagnosis is relevant and meaningful for a patient and oversee when and how to offer more clarification and more narrative explanation.”)

19. Knight, *supra* note 15 (“Some of it [i.e., intelligence] is just instinctual, or subconscious, or inscrutable.”).

20. See Sara Gerke, Timo Minssen & Glenn Cohen, *Ethical and Legal Challenges of Artificial Intelligence-Driven Health Care*, in *ARTIFICIAL INTELLIGENCE IN HEALTHCARE* 295, 313–14 (Bohr & Memarzadeh eds., 2020) (discussing the challenges AI-based technologies raise for medical tort liability regimes).

21. Scherer, *supra* note 17, at 369.

22. Hongfang Liu et al., *AI Model Development and Validation*, in *ARTIFICIAL INTELLIGENCE IN HEALTH CARE: THE HOPE, THE HYPE, THE PROMISE, THE*

while it is “widely accepted that the successful development of an AI system requires high-quality data,” a significant problem exists because the means to assess “the quality of data that are available and the methodology to create a high-quality dataset are not standardized or often are nonexistent.”²³ Another notable problem regarding the data sets used by AI is that they may not be properly representative of the population the AI is ultimately used on, leading to poor outcomes.²⁴ In the United States, for example, while there exist some datasets that are adequately representative of the population, “in many instances AI is being developed with data that is not population-representative.”²⁵ Matheny et al. note that “AI should be trained and validated on population-representative data to ensure accuracy for all populations and to achieve performance levels necessary for scalable success.”²⁶ Experts conclude that in order to “effectively use AI, it is essential to follow good data practices in both the creation and curation of retrospective datasets for model training and in the prospective collection of the data.”²⁷

The problem regarding data is different from the problem regarding opacity, as it not necessarily inherent in the nature of machine learning AI (like opacity). It is instead driven by devel-

PERIL 136 (Michael Matheny, Sonoo Thadaney-Israni, Mahnoor Ahmed, & Danielle Whicher eds., 2019); *see also id.* at 133 (“Developing AI based on bad data further amplifies the potential negative impacts of poor-quality data. Consider, for example, that race and ethnicity information is simply not recorded, is missing, or is wrong in more than 30 to 40 percent of the records at most medical centers.”).

23. *Id.* at 133.

24. Michael E. Matheny et al., *Artificial Intelligence in Health Care: Hope Not Hype, Promise Not Peril*, in *ARTIFICIAL INTELLIGENCE IN HEALTH CARE: THE HOPE, THE HYPE, THE PROMISE, THE PERIL* 214, 218 (Michael Matheny, Sonoo Thadaney Israni, Mahnoor Ahmed & Danielle Whicher eds., 2019).

25. *Id.*

26. Michael E. Matheny et al., *Artificial Intelligence in Health Care: A Report From the National Academy of Medicine*, 323 *J. AM. MED. ASS'N* 509, 509 (2020) (“[T]here continue to be issues of data quality, appropriate consent, interoperability, and scale of data transfers. The current challenges [regarding population-representative data] are grounded in patient and health care system preferences, and political will rather than technical capacity or specifications. It is prudent to engage AI developers, users, and patients and their families in discussion about appropriate policy, regulatory, and legislative solutions.”).

27. Liu et al., *supra* note 22, at 133. The authors additionally note that “[t]he quality of these data practices affects the development of models and the successful implementation at the point of care.” *Id.*

opers and users who either do not develop the technology adequately (e.g., feeding the AI with poor quality data) or use AI improperly (e.g., using AI developed with poor data or using AI on an unrepresentative population). If the AI is not using quality or relevant data, it will not operate as ideally intended and, if its results are relied upon by clinicians, poor or disastrous outcomes will likely follow. The solution to the data problem requires coordination of several stakeholders, including developers, government regulators, and end-users (i.e., physicians and other clinicians) of the technology.

III. CHALLENGES IN APPLYING AI TO TORT LAW

Difficulties will undoubtedly arise when cases unfold where the AI is alleged to have caused harm. Considering the nature of AI, complicated questions emerge when trying to ask courts “to unravel novel technology and apply ill-fitting case law to make determinations of liability” when traditional tort paradigms applicable to humans may not easily apply to AI.²⁸ Note that the legal challenges (and possible legal solutions) go beyond mere tort law. The legal impact of AI in the medical context is wide in scope and relevant to numerous legal spheres, such as intellectual property and privacy.²⁹ Additionally, it should be noted that while tort law is primarily governed by state common law, regulatory law can have strong influence over torts.³⁰ However, this

28. *Artificial Intelligence Litigation: Can the Law Keep Pace With the Rise of the Machines?*, QUINN EMANUEL: BUSINESS LITIGATION REPORTS <https://www.quinnemanuel.com/the-firm/publications/article-december-2016-artificial-intelligence-litigation-can-the-law-keep-pace-with-the-rise-of-the-machines/> (last visited Nov. 25, 2020).

29. See generally Price, *supra* note 4 (noting that several areas of law besides tort are relevant to AI in healthcare, including regulation (e.g., how the FDA will regulate AI algorithms), intellectual property (e.g., protection of algorithms via patent law and/or trade secrecy), and privacy (e.g., adherence to privacy laws like HIPAA)). See also, I. Glenn Cohen & Michelle M. Mello, *Big Data, Big Tech, and Protecting Patient Privacy*, 322 J. AM. MED. ASS'N 1141 (2019) (discussing privacy concerns related to AI and arguing that HIPAA is outdated and needs to be modernized to address concerns of data sharing and privacy in the 21st century).

30. Douglas McNair & W. Nicholson Price II, *Health Care AI: Law, Regulation, and Policy*, in ARTIFICIAL INTELLIGENCE IN HEALTH CARE: THE HOPE, THE HYPE, THE PROMISE, THE PERIL 181, 200 (Michael Matheny, Sonoo Thadaney Israni, Mahnoor Ahmed & Danielle Whicher eds., 2019) (“Different regulatory pathways influence the availability of state tort lawsuits against AI developers and, indirectly, the ability of state tort law (and liability insurers reacting to that law) to create independent incentives for the safe and effective

Article will primarily focus on the impact of medically-related American³¹ tort law in the realm of traditional negligence, without examining in detail the regulatory complications in a tort analysis.

A. MEDICAL MALPRACTICE (NEGLIGENCE)

Applying tort law (i.e., claims of negligence or medical malpractice³²) to AI is challenging in that “[c]ourts have traditionally deemed it impossible for machines to have legal liability as they are not legal persons.”³³ Hence, absent a model AI personhood (discussed later in this Article), the possible defendants in an AI malpractice case are the AI developers and manufacturers or the users of the AI (i.e., hospital, physicians, or other clinicians). However, parsing who is liable—either by determining causation, existence of duty, or breach of the standard of care³⁴—may be difficult in the context of medical AI.

To satisfy a claim of malpractice, the plaintiff must demonstrate that the physician’s breach is the “proximate cause” of the plaintiff’s injury.³⁵ Yavar Bathaee explains that:

development of clinical AI systems. In general, states may not establish statutory requirements that are ‘different from, or in addition to,’ FDA requirements regulating devices (21 U.S.C. § 360k). The U.S. Supreme Court has also held that this preempts certain state tort lawsuits alleging negligent design or manufacturing.”)

31. While the focus of this Article is on the tort law in the American context, if one has interest in delving deeper into other jurisdictions globally, the European Commission published an excellent report examining this issue from the context of European tort law. See LIABILITY FOR ARTIFICIAL INTELLIGENCE AND OTHER EMERGING DIGITAL TECHNOLOGIES, EXPERT GROUP ON LIABILITY AND NEW TECHNOLOGIES – NEW TECHNOLOGIES FORMATION, DS-03-19-853-EN-N (2019).

32. Hannah R. Sullivan & Scott J. Schweikart, *Are Current Tort Liability Doctrines Adequate for Addressing Injury Caused by AI?*, 21 AMA J. ETHICS, 160, 161 (2019) (“Typical tort claims in the realm of medicine and health include medical malpractice (negligence), *respondeat superior* (vicarious liability) and products liability.”).

33. Chung & Zink, *supra* note 7, at 51.

34. B. Sonny Bal, *An Introduction to Medical Malpractice in the United States*, 469 CLINICAL ORTHOPAEDICS AND RELATED RES. 339, 342 (2009) (stating that the elements of medical malpractice are “(1) the existence of a legal duty on the part of the doctor to provide care or treatment to the patient; (2) a breach of this duty by a failure of the treating doctor to adhere to the standards of the profession; (3) a causal relationship between such breach of duty and injury to the patient; and (4) the existence of damages that flow from the injury such that the legal system can provide redress”).

35. *Id.*

Fundamentally, proximate cause asks whether the result of the conduct was one that could have been foreseen by a reasonable person. At its core is the assumption that a person should not be liable for results having nothing to do with what he could have done to limit the risk of harm, nor should there be liability for the flukes of chance.³⁶

However, in the context of “black-box” AI, “the result of the AI’s decision or conduct may not have been in any way foreseeable by the AI’s creator or user.”³⁷ With an inability to establish causation, it can become exceedingly difficult to successfully assign liability in context of medical AI. Danny Tobey further notes that “[c]ausation and fault are increasingly opaque too, because of the number of human and machine interactions along the spectrum of AI usage.”³⁸

Sometimes, when causation may be difficult to prove, negligence claims may be resolved by the doctrine of *res ipsa loquitur*, a method where the common law has allowed negligence to be inferred to a particular defendant “when the accident causing the plaintiff’s physical harm is a type of accident that ordinarily happens because of the negligence of the class of actors of which the defendant is the relevant member.”³⁹ In the context of AI’s issues with opacity and explainability, *res ipsa loquitur* may be a possible solution in establishing liability, though issues remain. As Brandon Jackson notes, “[i]f the harm in question is unexplainable, untraceable, and rare, then the elements of *res*

36. Yavar Bathaee, *The Artificial Intelligence Black Box and the Failure of Intent and Causation*, 31 HARV. J.L. & TECH. 889, 923 (2018).

37. *Id.* at 924 (“Put simply, if even the creator of the AI cannot necessarily foresee how the AI will make decisions, what conduct it will engage in, or the nature of the patterns it will find in data, what can be said about the reasonable person in such a situation?”).

38. Danny Tobey, *Explainability: Where AI and Liability Meet*, DLA PIPER (Feb. 25, 2019), <https://www.dlapiper.com/en/us/insights/publications/2019/02/explainability-where-ai-and-liability-meet/>. Tobey gives a “real-life” example of the opaqueness problem the context of AI used in the legal profession: “a judge’s reliance on a bail-setting algorithm.” *Id.* “The algorithm suggested the risk of releasing the defendant back into society was low. The judgment followed the machine’s recommendation, only for the defendant to commit murder upon release.” *Id.* After such a tragedy and assuming negligence occurred, without an understanding of *how* the AI made its recommendation, it is not clear how to establish causation or attribute fault.

39. RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 17 (Am. Law. Inst. 2010).

ipsa loquitur likely cannot be satisfied.”⁴⁰ Additionally, courts using *res ipsa loquitur*, when making a negligence analysis, infer that the party with exclusive control over the instrumentality that inflicts harm, is in all likelihood, negligent.⁴¹ This “exclusive-control criterion is often effective in identifying the negligent party.”⁴² Hence, in the AI context, *res ipsa loquitur* may not be able to address all negligence uncertainties, because as previously discussed, AI systems often have a slew of developers and no one entity may have clear “control” over the “instrumentality” (i.e., the AI) that causes that patient’s injury.

Another key factor when considering medical malpractice in the face of AI, is how the standard of care changes in the liability analysis.⁴³ The standard of care is a factor in the “breach” element of negligence and medical malpractice. A physician’s deviation from the standard of care will amount to a breach in their duty to the patient.⁴⁴ Typically, a court evaluates this standard “by reference of what a reasonable *physician* would have done—i.e., a person with the same kind of technical background, training, and expertise as the defendant.”⁴⁵ But when a new technology, such as a medical device or AI, is introduced into practice, “the standard of care can be ambiguous as applied to the adoption by a physician of a new medical device.”⁴⁶ Hence, in this era

40. Brandon W. Jackson, *Artificial Intelligence and the Fog of Innovation: A Deep-Dive on Governance and the Liability of Autonomous Systems*, 35 SANTA CLARA HIGH TECH. L.J. 35, 59 (2019).

41. RESTATEMENT (THIRD) OF TORTS: LIABILITY FOR PHYSICAL AND EMOTIONAL HARM § 17 cmt. b (Am. Law. Inst. 2010).

42. *Id.*

43. Tobey & Cohen, *supra* note 8, at 22. (“[S]oon, with AI raising the achievable standards of care for patients, failure to adopt [an AI’s recommendation] may provide new sources of liability.”). Hence, those physicians who fail to adopt an AI’s determination may in fact become liable themselves, as the standard of care has shifted to include the AI.

44. Michael D. Greenberg, *Medical Malpractice and New Devices: Defining an Elusive Standard of Care*, 19 HEALTH MATRIX J. 423, 427 (2009) (“Establishing a whether a breach has taken place requires a comparison between the physician’s actions and a legal ‘standard of care’, which represents what physicians are obligated by law to do in providing medical services to their patients.”).

45. *Id.* at 428 (“Traditionally, the law has been very deferential to physician custom in determining what qualifies as malpractice. In other words, whatever constituted usual or typical medical care in a region was often formally defined as reasonable conduct, and where a physician’s treatment comported with professional custom, this was sufficient to avoid any breach in the duty of care.”).

46. *Id.* at 434.

of newly deployed AI, physicians are left with some degree of uncertainty (at least early on) about how and if they should be using AI, so as to properly align with the standard of care.⁴⁷

B. PRODUCTS LIABILITY AND THE LEARNED INTERMEDIARY DOCTRINE

Products liability is a common law doctrine that allows patients to be “entitled to recovery when they are injured by products that are ‘not reasonably safe’ due to defective design, manufacture, or warning.”⁴⁸ With regard to medical devices (and there are arguments for and against AI being classified as a medical device for products liability purposes⁴⁹), the common law states that “[a] manufacturer of a prescription drug or medical device who sells or otherwise distributes a defective . . . device is subject to liability for harm to persons caused by the defect.”⁵⁰ While products liability is straightforward, it may be complicated when applied to AI since, as discussed, AI may have multiple developers, and who may be deemed the manufacturer in some instances could be unclear. Likewise, the nature of opacity and the “black-box” may make it impossible in some instances to determine that an AI was actually a “defective” product.

An additional wrinkle in products liability is that of the learned intermediary doctrine. The doctrine has “potential to protect an AI’s manufacturer from failure-to-warn claims for

47. W. Nicholson Price II, *Medical Malpractice and Black-Box Medicine*, in *BIG DATA, HEALTH LAW, & BIOETHICS* 295, 303 (I. Glenn Cohen, Holly Fernandez Lynch, Effy Vayena, Urs Gasser, eds., Cambridge Univ. Press 2018) (“Providers will typically be held to the level of care of other comparable providers, which will develop over time as black-box medicine enters care.”).

48. Sullivan & Schweikart, *supra* note 32, at 162.

49. Jessica S. Allain, *From Jeopardy to Jaundice: The Medical Liability Implications of Dr. Watson and Other Artificial Intelligence Systems*, 73 *LA. L. REV.* 1049, 1072 (comparing AI systems to “a textbook or the Internet, which are clearly outside the scope of medical device regulation,” but explaining that such systems “make a clear impact on patients through [their] diagnosis and treatment recommendations,” which weighs in favor of including AI as a medical device).

50. *RESTATEMENT (THIRD) OF TORTS: PRODS. LIAB.* § 6(a) (1998); *see also id.* § 6(c) (“A prescription drug or medical device is not reasonably safe due to defective design if the foreseeable risks of harm posed by the drug or medical device are sufficiently great in relation to its foreseeable therapeutic benefits that reasonable health-care providers, knowing of such foreseeable risks and therapeutic benefits, would not prescribe the drug or medical device for any class of patients.”).

risks they have disclosed to intermediaries.”⁵¹ The doctrine functions this way:

The learned intermediary doctrine prevents plaintiffs from suing medical device manufacturers directly. Manufacturers have the duty to warn consumers of potential dangers inherent in a product’s natural use. In the case of medical devices [and hence, some AIs if applicable] manufacturers have a duty to warn the treating physician of the product’s potential dangers. The physician becomes a learned intermediary between the manufacturer and the patient, eliminating any duty the manufacturer may have had directly to the patient.⁵²

Also concerning from the standpoint of medical AI is that in the growing age of telemedicine,⁵³ the learned intermediary doctrine loses some of its initial purpose in allowing the physician to function as the protective “barrier” to the patient, as physicians are more removed and their warning may have less impact.⁵⁴

C. VICARIOUS LIABILITY

Vicarious liability is another aspect of tort law that must be considered when confronting medical AI. It is a tort “regime whereby one individual can be held legally responsible for the acts of another.”⁵⁵ In the employment context under a theory of agency, “the doctrine of *respondet superior* places vicarious liability on employers for the negligent acts of employees acting

51. Tobey & Cohen, *supra* note 8, at 22.

52. Allain, *supra* note 49, at 1068–69.

53. Telemedicine use has been dramatically growing in recent years. See *Telehealth Up 53%, Growing Faster Than Any Other Place of Care*, AMA (May 29, 2019), <https://www.ama-assn.org/practice-management/digital/telehealth-53-growing-faster-any-other-place-care>. Additionally, in light of the COVID-19 pandemic, the frequency of telemedicine is further increasing in addition to the already established growing trend. See Jedrek Wosik et al., *Telehealth Transformation: COVID-19 and the Rise of Virtual Care*, 27 J. AM. MED. INFORMATICS ASS’N 957, 959 (2020).

54. Tobey & Cohen, *supra* note 8, at 24 (“[W]hen physicians don’t interact with patients until *after* the prescription has been written, and then only through online messaging, physician warnings start to resemble labels, and the barrier of the ‘intermediary’ thins.”).

55. Allain, *supra* note 49, at 1064–65. Allain notes that this “is most commonly found in the employer-employee relationship in which an employer can be held liable for the employee’s tortious conduct.” *Id.*; see also RESTATEMENT (THIRD) OF TORTS: APPORTIONMENT OF LIAB., § 13 (2000).

within the scope of their employment.”⁵⁶ In the medical context, vicarious liability or *respondeat superior* is possible as “hospitals can be held vicariously liable for the acts of their employees, including physicians, who commit malpractice.”⁵⁷ In the context of AI, a question emerges whether or not an AI could be deemed an “agent” or “employee” of a hospital into order to impute liability to the “principal,” in this case a hospital.⁵⁸ Answering this question is challenging, as *respondeat superior* is founded on agency theory, premised by the idea that the principal (e.g., a hospital) has some control or power over the agent (e.g., an employee—normally a physician, but potentially also an AI system).⁵⁹ Control is the key to analysis, as an agency relationship is required for vicarious liability to attach and the concept of “control” between the parties, i.e., one party having supervisory power over another, is key to establishing the agency relationship. The difficulty with AI is its ever-growing autonomous nature. As Mark Chinen notes, “[t]he less sophisticated a machine is [as opposed to one that is fully autonomous], the more appropriate it is to focus on the individual human or group of humans who used it, and any harm caused by such a tool is readily attributable to its users.”⁶⁰ If a court deems an AI to be fully autonomous (or, if not autonomous, maybe held to be under the dominion of its designers rather than the hospital who purchased and uses it), then holding a hospital vicariously liable for any injury caused by AI will be impossible, as such an autonomous AI will functionally be outside of the principal’s control.

D. INFORMED CONSENT LIABILITY

While much of the discussion is focused on AI impacting traditional medical negligence, it may also impact informed consent

56. Sullivan & Schweikart, *supra* note 32, at 161–62.

57. Allain, *supra* note 49, at 1065.

58. *Id.* at 1066 (“Just as hospitals may be vicariously liable for a physician’s negligence, courts could likewise hold a hospital vicariously liable for injuries caused by its artificial intelligence systems.”).

59. RESTATEMENT (THIRD) OF TORTS: APPOINTMENT OF LIAB., § 13 cmt. a (2000).

60. Mark A. Chinen, *The Co-Evolution of Autonomous Machines and Legal Responsibility*, 20 VA. J.L. & TECH. 338, 360–61 (2016) (noting that “it can be argued that the sophistication of a machine does impact legal liability and stretches current conceptions of that liability. Put in terms of agency law, a completely autonomous machine would be capable of engaging in the frolic and detour” outside of the scope of the agency relationship.).

claims. Informed consent causes of action are a species of medical negligence that grew out of an emphasis on patient autonomy.⁶¹ Nadia Sawicki explains that “[i]n defining the scope of the informed consent duty, courts uniformly concluded that physicians have a legal obligation to inform patients of material information [i.e., risks, benefits, and alternatives] about a proposed course of treatment.”⁶² Tobey and Cohen note that in the context of medical AI, “treatment plans may be presented in absolutes without a discussion of options or pros and cons,”⁶³ thus reinforcing the physician’s role to ensure that adequate information is imparted to the patient in order to make an informed decision. Price also recognizes the problem of informed consent in the context of AI, noting that “[b]lack-box medicine raises nuanced informed-consent issues. At an intuitive level, it is hard to imagine precisely what ‘informed’ means in the context of a recommendation where no-one knows exactly how it works.”⁶⁴ Looking broadly at the landscape of informed consent cases, Glenn Cohen outlines three types of cases that may be relevant to medical AI: provider experience (e.g., a scenario where a physician fails to disclose use of AI during the informed consent process in an effort to hide inexperience),⁶⁵ substitute physicians (e.g., a scenario where the patient may be “unaware of the role of AI/ML [artificial intelligence/machine-learning] in only a *particular*

61. Nadia N. Sawicki, *Modernizing Informed Consent: Expanding the Boundaries of Materiality*, 2016 U. ILL. L. REV. 821, 822–23 (2016) (noting that “[i]t was not until the early 1960s that most medical and legal professionals began to recognize that malpractice liability could attach to a physician’s failure to properly inform her patient of the risks and benefits of proposed clinical treatment”).

62. *Id.* at 827.

63. Tobey & Cohen, *supra* note 8, at 24.

64. Price II, *supra* note 47, at 299 n.15; *see also* Gerke et al., *supra* note 20, at 301. The notion is also reflected by Gerke, Minssen and Cohen, when they ask: “[t]o what extent, for example, does a clinician need to disclose that they cannot fully interpret the diagnosis/treatment recommendations by the AI? How much transparency is needed?” *Id.*

65. I. Glenn Cohen, *Informed Consent and Medical Artificial Intelligence: What to Tell the Patient?*, 108 GEO. L.J. 1425, 1435–36 (2020). Cohen hypothesizes about a possible case where “failure to disclose AI/ML [artificial intelligence/machine-learning] involvement in decisionmaking or actual procedures, especially when AI/ML is meant to enable a regular doctor to perform at the level of a specialist, goes to the lack of experience/qualification of a doctor that ought to be disclosed.” *Id.*

part of their treatment”),⁶⁶ and financial conflicts of interest (e.g., a scenario where a physician fails to disclose a financial interest in the AI used in the patient’s treatment).⁶⁷

IV. HOW TORT LAW MAY EVOLVE

While it is early and the law is unfolding⁶⁸, undoubtedly tort law will evolve over the coming decades to adapt to the problems created by the nature and characteristics of AI outlined above. These legal “evolutions” can be thought of as legal solutions to the problems created by AI used in the medical context. When considering legal solutions, it is important to understand what the biggest problems raised by AI are. Gerke et al. explain that how AI-driven problems are prioritized will ultimately shape the legal policy solutions:

Setting the optimal liability regime depends heavily on what one thinks the “problem” is. If one is concerned that the deployment of AI-based technology in the clinical space is associated with a high risk for patients to get hurt, one might want to keep the current medical malpractice regime that attempts to meet both tort law’s two functions: (1) deterrence and (2) compensation of the victims. By contrast, if one believes that over the run of cases, reliance on AI promotes patient health, then it may be a problem if physicians prove reluctant to rely on these algorithms, especially the more opaque ones, when they remain on the hook for resulting liability . . . This might drive the policy-maker to a different model.⁶⁹

In light of this, three possible ways the law could evolve and adapt in an attempt to find solutions include: AI personhood, common enterprise liability, and a modification of the standard of care.⁷⁰

66. *Id.* at 1436–39 (“For example, the physician performs the surgery herself, but it is based on an AI/ML [artificial intelligence/machine-learning] recommendation of which surgical technique to use in this specific case.”).

67. *Id.* at 1439.

68. W. Nicholson Price II, *Medical AI and Contextual Bias*, 33 HARV. J.L. & TECH. 65, 86 (2019) (“The tort landscape for medical AI is largely theoretical, as the technology is just entering practice.”).

69. Gerke et al., *supra* note 20, at 313–14.

70. In 2019, Hannah Sullivan and I outlined these possible legal solutions. I will expound upon these further in this Part. *See*, Sullivan & Schweikart, *supra* note 32, at 163–64.

A. AI PERSONHOOD

Traditionally, courts have not viewed it as possible to hold a machine or computer (like AI) legally responsible, as they are “not legal persons.”⁷¹ Chung argues that, by way of analogy comparing AI to medical students, “there is enough overlap between the two [i.e., AI and medical students] in terms of level of authority, tasks and level of oversight for the court to make such a determination [i.e. a determination of personhood status for AI].”⁷² David Vladeck notes that “[c]onferring ‘personhood’ on these machines would resolve the agency question; the machines would become principals in their own right, and along with new legal status would come new legal burdens, including the burden of self-insurance.”⁷³ Of course, if we look to the personhood solution for AI, Vladeck points out that it is likely that the standard of care applied to AI and that applied to humans would be different, and autonomous machines would then have a standard of care unique to themselves.⁷⁴ Additionally, Chinen notes that conferring “personhood” on AI may not be such a far leap, as “giving legal personhood to things is not new. Ships and corporations enjoy status as legal persons and assume liabilities.”⁷⁵

B. COMMON ENTERPRISE LIABILITY

As discussed earlier, a key problem with AI technology is that there are many different developers involved in its creation. Scherer notes that to some extent, tort law already has a method to apportion liability, explaining that “[t]he discreteness of AI is

71. Chung & Zink, *supra* note 7, at 51.

72. Chung, *supra* note 5, at 39.

73. David. C. Vladeck, *Machines Without Principles: Liability Rules and Artificial Intelligence*, 89 WASH. L. REV. 117, 150 (2014); *see also* Allain, *supra* note 49, at 1079 (explaining that such self-insurance could be “a mandatory malpractice insurance policy” that is paid for by the AI’s owner).

74. Vladeck, *supra* note 73, at 130, 132. (“[T]he court in the first driver-less car case [using driver-care technology as the example AI] will likely ask whether the car involved in the accident performed up to the standards achievable by the majority of the other driver-less cars, as well as the performance specification set by the car’s manufacturer.”).

75. Chinen, *supra* note 60, at 387–88 (noting that conferring personhood on AI may also come with a “combination of pragmatic and ethical quandaries,” such as “whether machines should be granted legal rights in addition to duties.”).

also shared by many other modern and non-so-modern technologies.”⁷⁶ American jurisprudence already has a body of common law regarding tort apportionment of liability,⁷⁷ and this may work to solve some issues. However, because of the nature of some AI and the intricate aspect of its development between multiple parties, being able to apportion liability under traditional scenarios may not always be possible.

Vladeck explains that one could, in theory, apply a strict liability model that would simply function to hold only the manufacturer or developer of an AI automatically liable.⁷⁸ However such a model is not ideal⁷⁹ as there are other companies and entities that help produce the final complex product, not simply the main manufacturer.⁸⁰ As a solution, Vladeck proposes “common enterprise” liability,⁸¹ where “each entity within a set of interrelated companies may be held jointly and severally liable for the actions of other entities that are part of the group.”⁸² This theory of liability is advantageous when applied to AI, as it “would not require that the companies function jointly; all that would be required is that they work to a common end—to design, program, and manufacture” an AI product and “its various component parts.”⁸³ As Vladeck further explains, common enterprise liability “permits the law to impose joint liability without having to

76. Scherer, *supra* note 17, at 374 (noting an example with regards to automobile technology, explaining that “[a]utomobiles have long been manufactured using components from multiple companies and courts long ago developed rules for apportioning liability when harm is caused by defects in multiple such components.”).

77. See RESTATEMENT (THIRD) OF TORTS: APPORTIONMENT OF LIABILITY §§ 10–17 (Am. L. Inst. 2013).

78. See Vladeck, *supra* note 73, at 146 (“My proposal is to construct a system of strict liability, completely uncoupled from notions of fault for this select group of cases.”).

79. See Bathaee, *supra* note 36, at 894. While not explored in great detail in this article, strict liability could be another potential solution. However, as Bathaee points out, it is “a poor solution for the problem because if one cannot foresee the solutions an AI may reach or the effects it may have, one also cannot engage in conduct that strict liability is designed to incentivize, such as taking necessary precautions or calibrating the level of financial risk one is willing to tolerate.” *Id.*

80. See Vladeck, *supra* note 73, at 148 (“From a cost-spreading standpoint, it is far from clear that the manufacturer should absorb the costs when parts and computer code supplied by other companies may be the root cause.”).

81. *Id.* at 149.

82. *FTC v. Tax Club, Inc.*, 994 F. Supp. 2d 461, 469 (S.D.N.Y. 2014).

83. Vladeck, *supra* note 73, at 149.

lay bare and grapple with the details of assigning every aspect of wrongdoing to one party or another; it is enough that in pursuit of a common aim the parties in engaged in wrongdoing.”⁸⁴

C. NEW STANDARD OF CARE

Possibly the most natural solution is to have common-law medical malpractice modify its standard of care over time for physicians to reflect the newfound reality of medical AI in clinical practice. Price argues that the standard will likely “develop[] organically,” but that it also could be created via a professional organization as a practice guideline.⁸⁵ Such a guideline might suggest a standard of care that is responsive to the level of risk, where a physician may have to seek greater validation or inquiry before accepting an AI’s determination if the risk and consequences are high.⁸⁶ Greenberg also echoes this notion of validation, explaining that when a physician is exposed to new medical device, the standard of care may require a physician to take extra steps to learn as much as possible, e.g., “by reading all of the clinical trial information on the risks and benefits of a new device, following any pertinent label instructions, and seeking relevant training where appropriate.”⁸⁷ And considering the issues some AI have with being developed with improper data, physicians need to do their research and learn how an AI was developed, what risks may exist in its use, and if its use is appropriate for the physician’s patient population. Just like with any other novel medical device a physician learned to use, the standard of

84. *Id.* at 149 (noting that rather than there being a wrongdoer, there is instead “an inference of liability drawn by operation of law to protect a blameless party [e.g. in the medical AI context, an injured patient] . . .”).

85. See W. Nicholson Price II, *Medical Malpractice and Black-Box Medicine* 9 (U. Mich. L. Sch., Working Paper No. 536, 2017) (“While the standard typically develops organically through the practice of providers (while cyclically shaping those practices), the standard can also be influenced by legislative action, judicial standard-setting, or practice guidelines set by professional organizations.”).

86. *Id.* at 9. (“For minimal-risk interventions, such as otherwise un-indicated testing, increased monitoring, or taking widely used low-side-effect drugs like aspirin, the standard of care might require no particular inquiry of the recommendations of a black-box algorithm. For riskier interventions, such as taking higher doses of a powerful drug—or avoiding such a course when otherwise suggested—providers might require some validation before relying on a black-box algorithm.”).

87. Greenberg, *supra* note 44, at 444.

care will undoubtedly shift to require a physician to perform due diligence⁸⁸ in any AI they use.

V. CONCLUSION

Looking to the future, some mix of the three possibilities outlined above—AI personhood, common enterprise liability, and a new standard of care—is likely to be implemented. All of the solutions serve a common goal: to allow the injured party to have some course of redress. If tort law does not adapt to the reality of medical AI and the problems it creates (e.g., notably, opacity and explainability), many claims will simply not be viable as causation will be impossible to demonstrate. What the exact landscape of tort law will look like in the future will depend on what aspects of the problems associated with AI society prioritizes.⁸⁹ If injured patients are the societal priority, then solutions that allow them to become whole will incorporate liability paradigms that will allow for the developers or physicians to be liable. Assuming that society's priority is to protect injured patients first, the next battle will be waged between AI developers and physicians over who should shoulder more of the liability regarding AI used in medical practice. For example, the solutions of AI personhood and common enterprise liability put most of the liability and risk towards the developers, while a shift in the standard of care can allow for the physicians (as end users of the technology) to have liability in certain scenarios. How this potential tension between AI developers and physicians plays out, and how medical tort law continues to develop in the future in response to AI, will be ever fascinating.

88. See Price II, *supra* note 85, at 14 (“Providers and facilities should evaluate black-box algorithms for hallmarks of careful development, including independent validation of algorithmic results and the qualifications of the developers.”).

89. See Gerke et al., *supra* note 20, 313–14 (reiterating that what society prioritizes in terms of the problem of AI and liability will ultimately shape the legal solutions).