

MACHINE VISION, MEDICAL AI, AND MALPRACTICE

*Zach Harned, M.A. M.S.**
Matthew P. Lungren, M.D. M.P.H.†
Pranav Rajpurkar, M.S.‡

Published March 13, 2019
Original link: <https://jolt.law.harvard.edu/digest/machine-vision-medical-ai-and-malpractice>

Recommended Citation

Zach Harned, Matthew P. Lungren & Pranav Rajpurkar, Comment, *Machine Vision, Medical AI, and Malpractice*, HARV. J.L. & TECH. DIG. (2019), <https://jolt.law.harvard.edu/digest/machine-vision-medical-ai-and-malpractice>.

Read more about JOLT Digest at <https://jolt.law.harvard.edu/digest>

*Zach Harned, Stanford Law School, Stanford, CA

†Dr. Matthew P. Lungren, Department of Radiology, Stanford School of Medicine, Stanford, CA

‡Pranav Rajpurkar, Stanford Department of Computer Science, Stanford, CA

ABSTRACT

The introduction of novel medical technology into clinical practice gives rise to novel questions of legal liability when something goes wrong. The complexity of the technology is often paralleled by the complexity of the liability analysis, which is why questions of malpractice involving medical artificial intelligence are so vexing. There are myriad medical use cases for artificial intelligence (AI), but some of the most promising applications involve the use of machine vision for imaging diagnostics.

However, these machine vision applications involve complicated software models, the operation of which can be opaque at times even to its designers. This introduces concerns from physicians over whether they can trust a machine they do not fully understand or rely on its judgements. This can also arouse fear over the possibility of malpractice claims.

Some of the recent advances in machine learning technology make its results easier to interpret, allowing medical professionals to feel more confident in using the technology. This article illustrates how such innovations are likely to impact the legal system and malpractice suits. We conclude that the unique capabilities and functions of AI and machine vision, especially when conjoined with the aforementioned advances in their interpretability, create an opportunity to argue that the technology actually minimizes physician liability.

These advances in machine vision interpretability also change the legal landscape for the manufacturers of this technology. We examine impacts to products liability, focusing specifically on the issue of whether such technology would (or will soon) be considered a “product,” and how this might affect manufacturers’ product development and marketing strategies. We also consider how the learned intermediary defense might be deployed in failure-to-warn cases involving medical machine vision, again looking to how the legal doctrine is likely to impact manufacturer behavior in the design and deployment of such technologies.

TABLE OF CONTENTS

OVERVIEW 3

I ACCURACY 3

II INTERPRETABILITY 4

III MEDICAL MALPRACTICE 6

IV PRODUCTS LIABILITY 7

V LEARNED INTERMEDIARY DOCTRINE 9

CONCLUSION 10

OVERVIEW

Medical use cases for artificial intelligence (AI) are rapidly expanding, and the most promising early applications involve the use of machine vision for medical imaging. However, the introduction of AI technology into clinical practice gives rise to challenging legal questions of liability given the absence of pertinent case law. Consideration of the accuracy and interpretability of these machine vision systems will assist with these challenging malpractice concerns. Additionally, looking to machine vision applications in other industries will help illuminate the various strategies for minimizing exposure to liability.

I. ACCURACY

One of the reasons for the great interest in medical machine vision applications is the impressive accuracy these algorithms exhibit. This has resulted in expert performance in a variety of clinical diagnostic tasks across numerous image-centric specialties, including radiology,¹

¹ See Pranav Rajpurkar et al., *Deep Learning for Chest Radiograph Diagnosis: A Retrospective Comparison of the CheXNeXt Algo-*

dermatology,² pathology,³ and ophthalmology.⁴ Algorithms designed to automate the performance of clinical tasks will likely be categorized by the US Food and Drug Administration (FDA) as medical devices, by either being embodied in traditional medical devices, or classified under the Software as a Medical Device guidance.⁵ The FDA has recently approved numerous AI medical tools, the majority of which are machine vision imaging applications,⁶ thus evincing the early popularity of this technology in medicine.

II. INTERPRETABILITY

It is obvious that we want AI applications in radiology to be as accurate as human experts, if not more so. But why is it so critical for these models to also be interpretable? There are some engineering benefits of interpretability, including the ability to use the feedback to make targeted adjustments to the model. Yet, interpretability is particularly important to radiologists because it provides assurance that the model is behaving as intended (e.g. focusing on the relevant aspects of the image). Additionally, interpretability provides a warning mechanism for the clinician regarding potential biases in the machine vision system.

Typically, machine vision models are black boxes, making it difficult to see *why* the model made a particular decision or diagnosis. One can argue that the human decision-making process is equally opaque.

rithm to Practicing Radiologists, PUB. LIBR. OF SCI. (Nov. 20, 2018),

<https://journals.plos.org/plosmedicine/article?id=10.1371/journal.pmed.1002686>.

² See Andre Esteva et al., *Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks*, 542 NATURE 115, 115–18 (2017).

³ See Kun-Hsing Yu et al., *Predicting Non-Small Cell Lung Cancer Prognosis by Fully Automated Microscopic Pathology Image Features*, NATURE COMM. (Aug. 16, 2016), <https://www.nature.com/articles/ncomms12474>.

⁴ See Jonathan Krause et al., *Grader Variability and the Importance of Reference Standards for Evaluating Machine Learning Models for Diabetic Retinopathy*, 125 OPHTHALMOLOGY 1264, 1264–1272 (July 3, 2018).

⁵ See *generally Software as a Medical Device (SAMM): Clinical Evaluation - Guidance for Industry and Food and Drug Administration Staff*, FDA (Dec. 8, 2017), <https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm524904.pdf>.

⁶ See Dave Muoio, *Roundup: 12 Healthcare Algorithms Cleared by the FDA*, MOBIHEALTHNEWS (Nov. 15, 2018), <https://www.mobihealthnews.com/content/roundup-12-healthcare-algorithms-cleared-fda>.

However after giving a diagnosis, the physician can be asked to justify her decision, which she can often do via simple ostension, pointing to the area of the scan that is most relevant for making a particular diagnosis.

There are now various technical methods that allow machine vision models to justify their decisions as well. One such technique involves overlaying a heatmap onto the medical image, which shows that the machine vision model is indeed looking at the relevant parts of the image when coming to its decision. These heatmaps can be created using various machine learning techniques, including class activation mapping (CAM)⁷ and saliency mapping.⁸

In addition to pointing to the scan, there are other ways in which a radiologist's diagnosis is justified. For instance, the radiologist might highlight the most pertinent features of the patient's clinical presentation or prior history from the electronic health record. Similarly, machine learning researchers have developed a counterfactual explanation technique that performs a comparable function in the medical context, listing the top diagnostic, procedural, medication, encounter, and demographic factors that contributed to its decision.⁹

Another strategy commonly used by physicians is explanation by example. That is, a doctor may give a certain diagnosis because the patient image and presentation are similar to paradigmatic cases the doctor has encountered before. An analogous machine learning technique uses case-based learning methods, which allows for the identification of the cases most similar to the case needing to be explained.¹⁰ By using these methods, physicians can ensure that the machine is not only making the same decisions a human physician would, but that it is doing so based on the *same reasons* as a human physician.

⁷ See Rajpurkar, *supra* note 1.

⁸ See Esteva, *supra* note 2.

⁹ See generally Anand Avati et al., *Improving Palliative Care with Deep Learning*, 2018 BMC MED. INFORMATICS & DECISION MAKING 122 (2018), available at <https://bmcmmedinformdecismak.biomedcentral.com/articles/10.1186/s12911-018-0677-8>.

¹⁰ See Rich Caruana et al., *Case-Based Explanation of Non-Case-Based Learning Methods*, AMIA ANN. SYMP., Feb. 1999, at 212, 212–15.

III. MEDICAL MALPRACTICE

Machine vision's ability to match not only human-level accuracy in radiology but also human-level interpretability is particularly important in making determinations of legal liability for physician medical malpractice. The legal standard most frequently employed for malpractice is whether the physician failed to comply with customary medical practice.¹¹ Because a physician's usage of nearly any new medical technology runs the risk of failing to comport with custom, the method in which these machine vision models are deployed is important.

AI-powered machine vision is just starting to be explored in clinical practice, and its prevalence is bound to increase as patients and medical professionals grow more comfortable with the technology. Early on, machine vision will likely be deployed as a triage tool for patient images, or serve as a computer aided-detection (CAD) product. Both of these applications could be used to shield a physician from liability in a lawsuit.¹² The CAD paradigm functions as an "over-read" or "second read" to identify pathology that would have otherwise been missed. But such technology could also be used as a preliminary read, like that performed by radiology residents before the attending physician signs off.

It is natural to expect machine vision technologies to eventually assume functions of expert specialist clinicians, particularly as accuracy continues to improve and familiarity with routine clinical use becomes more widespread. Yet because a lawsuit involving medical machine vision would be a matter of first impression with no clear precedent, it is challenging to predict exactly how a court would handle physician liability. Interestingly, the interpretability of the machine vision device might provide an interesting legal analogue to comporting with custom: if the machine vision device were merely an accurate black box, using it as a second opinion clearly does not comport with any custom. In other words, there is no medical custom for consulting an "oracle." However, if the machine vision device is accurate *and interpretable*, it seems that consulting with it could possibly be construed as customary medical

¹¹ See Ben A. Rich, *Medical Custom and Medical Ethics: Rethinking the Standard of Care*, 14 CAMBRIDGE Q. HEALTHCARE ETHICS, Feb. 2005, at 27.

¹² Cf. *Darden v. Driscoll*, No. 1:12-cv-01001-EPG (PC), 2016 U.S. Dist. LEXIS 174510, at *13–16 (E.D. Cal. Dec. 16, 2016).

behavior, similar to “curbside” consultation with an expert specialist clinician.

For example, consider the elements that constitute a legitimate consultation with a radiologist to diagnose pneumonia: (1) the consulting radiologist has received proper training, (2) she can describe how the findings and patient history contribute to the diagnosis of pneumonia, (3) she can identify similar cases, and (4) she can point to the scan to show what she thinks is relevant. Analogously, a physician consulting an accurate and interpretable machine vision radiology device is in dialogue with a consultant that (1) has received proper training, albeit a somewhat different kind than a human radiologist receives, (2) can list the supporting clinical history and findings that led to the diagnosis of pneumonia (using counterfactual explanation), (3) can identify similar cases (using case-based learning methods), and (4) can “point” to the parts of the scan that are relevant (using heatmaps). Given these similarities, a physician consulting a machine vision radiology device could be considered to be complying with medical custom. This is because the four elements legitimating a radiology consultation are present both in the case of consulting an expert clinician and consulting an accurate and interpretable machine vision device.

As medical AI improves, however, its use might itself become customary or even necessary, especially if the algorithm has a sufficiently long track record of outperforming human physicians in diagnosing a disease. Indeed, we may even reach a point where the customary standard of care requires the use of such an algorithm, just as custom has changed to adopt other diagnostic technologies such as MRIs or CT scans.

IV. PRODUCTS LIABILITY

The physician is not the only player in the machine vision ecosystem concerned with questions of liability. Another similarly concerned party is the technology manufacturer. A manufacturer can typically be found liable for its product causing a harm under the doctrine of products liability. However, given that products are defined as *tangible* personal

property,¹³ judges have been loath to apply products liability doctrine to software on the grounds that software is not truly a product.¹⁴

This may be changing, as software increasingly becomes embodied in tangible systems capable of causing harm. One of the driving forces behind this development is the advent of autonomous vehicles, wherein AI software is embodied in a machine capable of causing great harm to passengers and pedestrians. Similarly, medical machine vision software may be embodied in a tangible machine involving a camera or various sensors, possibly opening it up to products liability if there is a resultant harm.

Medical machine vision manufacturers may attempt to solve this issue in a manner similar to the autonomous vehicle manufacturers, that is, by appealing to their conformity with industry standards, or assisting in the creation of such standards when they do not yet exist.¹⁵ But perhaps this development will instead push innovators in this space to manufacture software-only medical machine vision technologies to try and stave off products liability. In this case, it would then be up to the courts to expand the products liability doctrine to include pure software and hold such manufacturers responsible. But some legal scholars have noted that this change may be arriving in the not-too-distant future, given that we have already seen a few cases chipping away at this doctrine,¹⁶ and the question of computer code constituting a product has not been aggressively litigated.¹⁷

¹³ See RESTATEMENT (THIRD) OF TORTS: PRODUCTS LIABILITY § 2 (1998).

¹⁴ See Seldon J. Childers, *Don't Stop the Music: No Strict Products Liability for Embedded Software*, 19 U. FLA. J.L. & PUB. POL'Y 125, 142 (2008).

¹⁵ See Daniel A. Crane et al., *A Survey of Legal Issues Arising from the Deployment of Autonomous and Connected Vehicles*, 23 MICH. TELECOMM. TECH. L. REV. 190, 272 (2017).

¹⁶ See, e.g., *Winter v. G.P. Putnam's Sons*, 938 F.2d 1033, 1036 (stating in dicta that software could be viewed as a "product"); see also Susan M. Gilles, "Poisonous" Publications and Other False Speech Physical Harm Cases, 37 WAKE FOREST L. REV. 1073, 1076 n.11 (2002) (collecting cases).

¹⁷ See 3 RAYMOND T. NIMMER, THE LAW OF COMPUTER TECHNOLOGY § 12:31, at 12-78 (4th ed. 2013).

V. LEARNED INTERMEDIARY DOCTRINE

Under products liability, the manufacturer also has a duty to adequately warn the consumer—or in the medical context, the patient—about the risks of the product. Given how onerous this could be in the medical setting, many courts accept the learned intermediary rule, which asserts that once the physician has received adequate warning of the products risks, it is then her duty to convey that warning to the patient.¹⁸

However, there are two exceptions to this rule. The first is that if the manufacturer engages in direct consumer marketing, then the learned intermediary rule cannot be used as a defense.¹⁹ The second—and more interesting of the two—is that if the physician is not playing an active role with regard to the product and patient, then the manufacturer cannot make use of the learned intermediary defense.²⁰

This second exception is interesting because it may very well bear on the design and development of diagnostic machine vision systems used by physicians. For instance, if such a diagnostic system is designed to take the scan, read it, make the diagnosis, and then present it to the physician who acts merely as a messenger between the system and the patient, then it would seem that the physician is playing a relatively passive role in this provision of treatment. If a faulty diagnosis is made resulting in patient harm, and the patient was not adequately warned regarding this risk, the physician's relatively passive role could end up insulating her from liability, because it could eliminate the manufacturer's ability to utilize the learned intermediary defense. Manufacturers might therefore be more likely to design their medical machine vision systems to actively engage the physician not only because they believe it will increase the likelihood of physician and hospital adoption, but also because it may provide the manufacturer shelter from liability by enabling them to use the learned intermediary defense.

¹⁸ See RESTATEMENT (THIRD) OF TORTS, PRODUCTS LIABILITY § 6(d) (1998).

¹⁹ See *Perez v. Wyeth Labs. Inc.*, 734 A.2d 1245, 1258 (N.J. 1999).

²⁰ See *MacDonald v. Ortho Pharm. Corp.*, 475 N.E.2d 65, 69 (Mass. 1985) (citing the “relatively passive role” the physician plays in prescribing oral contraceptives to young women as one of the reasons the learned intermediary doctrine did not apply).

CONCLUSION

Although the history of medicine is rife with instances of technological innovation, AI-based machine vision medical devices appear different in kind, not merely degree. Most new medical devices aim to augment the physician in some way, whereas these machine vision applications are designed to emulate—and in some circumstances, completely assume—specific tasks of physician specialists and subspecialists. This raises ethical, financial, and regulatory questions, all of which involve significant legal concerns. Properly trained physicians are proficient at making accurate diagnoses and providing explanations and reasons for their decisions; machine vision devices are able to match and sometimes exceed that accuracy, and further technical tools provide the devices with the ability to explain their decisions as well. The more that these AI devices behave like physicians—accurately and interpretably—the more likely using such technology could be seen as comporting with typical medical practice, thereby helping minimize liability exposure for physicians. Manufacturers will also be looking to minimize their liability profile for machine vision medical devices, perhaps by using the software exemption under products liability,²¹ or by having physicians take an active role when using these machine vision systems in order to provide the manufacturers with the option of using the learned intermediary defense. Careful consideration is needed to determine how the legal system should handle machine vision and medical AI more broadly, in order to ensure that incentives are properly aligned, and that liability falls where it should.²² ■

²¹ *But cf.* Winter, *supra* note 16; Gilles, *supra* note 16; NIMMER, *supra* note 17 (describing the erosion of this exemption).

²² For a discussion of how the novelty of 3-D printing seems to impede the current legal system from placing liability on any of the traditional tortfeasors, see generally Nora Freeman Engstrom, *3-D Printing and Product Liability: Identifying the Obstacles*, 162 U. PA. L. REV. ONLINE 35 (2013).