2017

Artificial Intelligence in Health Care: Applications and Legal Implications

W. Nicholson Price II
University of Michigan Law School, wnp@umich.edu

ARTIFICIAL INTELLIGENCE IN HEALTH CARE
APPLICATIONS AND LEGAL ISSUES

BY W. NICHOLSON PRICE II

Artificial intelligence (AI) is rapidly moving to change the healthcare system. Driven by the juxtaposition of big data and powerful machine learning techniques—terms I will explain momentarily—innovators have begun to develop tools to improve the process of clinical care, to advance medical research, and to improve efficiency. These tools rely on algorithms, programs created from healthcare data that can make predictions or recommendations. However, the algorithms themselves are often too complex for their reasoning to be understood or even stated explicitly. Such algorithms may be best described as "black-box."[1] This article briefly describes the concept of AI in medicine, including several possible applications, then considers its legal implications in four areas of law: regulation, tort, intellectual property, and privacy.

AI in Medicine

Medicine, like many other fields, is experiencing a confluence of two recent developments: the rise of big data, and the growth of sophisticated machine learning/AI techniques that can be used to find complex patterns in those data. Big data as a phenomenon is characterized by the "three Vs" of *volume* (large quantities of data), *variety* (heterogeneity in the data), and *velocity* (fast access to the data). In medicine, the data come from many sources: electronic health records, medical literature, clinical trials, insurance claims data, pharmacy records, and even information entered by patients into their smartphones or recorded on fitness trackers. Machine learning techniques, a subset of AI, use simple learning rules and iterative techniques to find and use patterns in these vast amounts of data. The resulting algorithms can make predictions and group sets—how long is a patient expected to live given his collection of symptoms, and does that picture of a patch of skin look like a benign or a cancerous lesion?—but typically, these techniques cannot explain *why* or *how* they reach the conclusion they do. Either they cannot explain it at all, or they can give explanations that are accurate but meaningless in terms of medical understanding.[2] Because of this inherent opacity (which might or might not be augmented with deliberate secrecy about how the algorithms were developed and validated), I describe this field as to "black-box medicine," though it has also been referred to as AI in medicine or "predictive analytics."[3] To add to the complexity, when more data are available for the machine learning algorithms, those data can be incorporated to refine future predictions, as well as to change the algorithms themselves. The algorithms at the heart of black-box medicine, then, are not only opaque but also likely to change over time.

Black-box medicine has tremendous potential for use throughout the healthcare system, including in prognostics, diagnostics, image analysis, resource allocation, and treatment recommendations. Machine learning is most familiar in the context of image recognition, and an algorithm has already been developed that can identify skin cancer by analyzing images of skin lesions; the algorithm performs as well as board-certified dermatologists.[4] A recent *New England Journal of Medicine* article suggests that such algorithms could soon enter widespread use in image analysis, aiding or displacing much of the work of anatomical pathologists or radiologists within the span of years.[5] Another current algorithm can predict which trauma victims are likely to hemorrhage by constantly analyzing vital signs and can in turn call for intervention to forestall catastrophe; such prognostic algorithms could come into use in a similarly short time frame.[6] A bit farther off, black-box algorithms could be used for diagnosis more generally, to recommend off-label uses for existing drugs, to allocate scarce resources to patients most likely to benefit from them, to detect fraud or problematic medical behavior, or to guide research into new diseases or conditions. In fact, black-box algorithms are already in use today in smartphone apps that aim to identify developmental disorders in infants based on facial features[7] or autism in young children based on eye movement tracking.[8] The potential for benefit from such black-box medicine is substantial, but it comes with its own challenges: scientific and medical, certainly, but also legal. How do we ensure that black-box medicine is safe and effective, how do we ensure its efficient development and deployment, and how do we protect patients and patient privacy throughout the process?

Regulation

The first question to ask is perhaps the most fundamental: How do we ensure that black-box algorithms are high quality—that is, that they do what they say, and that they do it well and safely? New and emerging medical technologies and devices are typically regulated for safety and efficacy by the Food and Drug Administration (FDA). Whether the FDA actually has statutory authority over free-standing algorithms used to make medical decisions (or to help make them) depends on the relatively complex question of what is a "medical device." The FDA's regulation

***W. Nicholson Price II**, PhD (wnp@umich.edu) is an assistant professor of law at the University of Michigan Law School. His work focuses on innovation in the life sciences, with a significant emphasis on the use of big data and artificial intelligence in health care.*

of black-box medical algorithms may also conflict with its long-standing statement that it does not regulate the practice of medicine.[9] Elsewhere, I argue that the FDA has this authority, probably over algorithms standing alone and almost certainly in the context of linked technology that may more readily be called a "medical device," but disputes may arise over this point.[10] Industry dynamics may also play a role here: Silicon Valley, the hub of much of the innovation in AI generally, traditionally has not worked closely with regulators like the FDA.

Assuming that the FDA can and will regulate AI in the healthcare system (and the agency has asserted this ability and intent),[11] typically two tools help ensure safety and efficacy of new medical technology: scientific understanding and clinical trials. Unfortunately, these two tools do not work well in the context of black-box medicine. Understanding does not work for obvious reasons—we do not understand how a black-box algorithm makes decisions, because the machine learning techniques generally cannot tell us their reasoning, and even when they can, the results are often too complex to understand. Using clinical trials for testing safety, efficacy, and validity might work for some algorithms, but will not work for many others. For algorithms that divide patients into groups and suggest a particular treatment, clinical trials could be used to test their efficacy. But some algorithms will make highly personalized treatment predictions or recommendations, so that the use of clinical trials would be infeasible. And even for algorithms that are amenable to trials, the benefits of black-box medicine—quick, cheap shortcuts to otherwise inaccessible medical knowledge—would be seriously delayed or even curtailed due to the slow, ponderous, expensive enterprise of clinical trials. For algorithms that change as they incorporate more data, the challenges are even more pronounced. In short, in black-box medicine, traditional methods of testing new medical technologies and devices are likely not to work at all in some instances, and to slow or stifle innovation in others.

So how should the FDA tackle this challenge? The most fruitful path, I argue, will likely be more flexible than rigid, involving somewhat lighter premarket scrutiny (focused on procedural safeguards like the quality of the data used, the development techniques, and the validation procedures) coupled with robust post-market oversight as these algorithms enter into clinical care. The FDA has recently expressed interest in this approach.[12] Of course, this is easier said than done; the parallel case of post-market surveillance for drugs is notoriously troublesome to implement. One attractive possibility would be for the FDA to enable oversight help from other sophisticated healthcare entities by collaborating with them and, crucially, enabling ways to get them important and useful information. Hospitals, insurance companies, and physician specialty associations all have an interest in ensuring that black-box algorithms actually work to help patients (and, potentially, their bottom lines). Rival developers may also have an interest, especially in finding problems with existing algorithms. In addition, these sophisticated entities may have the capacity to perform evaluations, especially as they are used in clinical practice, and to generate performance data. Nevertheless, performing this type of collaborative governance role requires information, and many algorithm developers are reluctant to share that kind of information with any other entities. Potentially the FDA could serve as a centralized information-sharing role to allow those other entities to play their part in regulating black-box medicine. However, exactly how this idea might become a reality is very much an unresolved question.

Tort

What do we do when black-box medicine goes awry? The law of tort interacts with black-box medicine in a few different contexts. First, if there are flaws built into the algorithms themselves, or if regulation fails to ensure that algorithms are high quality, then the developers of algorithms (or technologies that rely on them) might become liable under tort law. However, courts have been reluctant to extend or apply product liability theories to software developers, and even more reluctant in the context of healthcare software.[13]

Part of that reluctance has come from the fact that healthcare software to date has been characterized primarily as technology that helps healthcare providers make decisions by providing them with information or analysis, with the final decision always resting in the hands of the provider. Black-box medicine turns that notion on its head, or at least it can. Can and should healthcare providers be fully responsible for decisions suggested or made by black-box algorithms that they do not, or cannot, understand?

This raises a second set of questions. What must healthcare providers and healthcare institutions—doctors, nurses, hospitals, managed-care organizations, and the like—do to fulfill their duties of care to patients in a healthcare world with black-box algorithms? Must providers themselves evaluate the quality of black-box algorithms, based on procedural measures (validation undertaken, performance statistics, etc.) before relying on those algorithms in the course of providing care? And should healthcare institutions perform similar evaluations before implementing black-box software? I argue elsewhere that they should, but currently the information necessary for that type of evaluation is largely unavailable—just as in the parallel regulatory context mentioned above.[14] Similarly, if an algorithm suggests an intervention that seems mundane but unhelpful, useless and expensive, or dangerous, should the provider second-guess the recommendation? On the one hand, the answer seems an obvious "yes"—providers are trained to care for patients—but on the other hand, if providers only implement those decisions they would have reached on their own, they will leave on the table much of the benefit that black-box medicine promises to extract from otherwise inaccessible patterns in big data. This would not leave *everything* on the table—algorithms can still potentially perform the usual analyses more quickly and cheaply[15]—but excessive caution is not costless. Courts have not tackled these issues yet, but they will need to in the near future.

Intellectual Property

Intellectual property protection creates another set of challenges for the development of black-box medicine.[16] When firms invest in developing black-box algorithms, how can they protect that investment? Developing black-box algorithms can involve considerable expense. Developers must generate, assemble, or acquire the tremendous data sets needed to train their algorithms; they must assemble the expertise and resources to actually develop those algorithms; and they must validate them to make sure they work. Normally, we might expect intellectual property to provide some measure of protection for the information goods created by such expenditures, so that firms are willing to invest the necessary funds for their development without fear that resulting inventions will be appropriated by others.[17] However, intellectual property fits relatively poorly for black-box medicine.

Patents are a natural choice to protect technological innovation, but patents do not provide strong incentives for black-box medicine. A string of recent decisions by the U.S. Supreme Court interpreting section 101 of the Patent Act, which governs patentable subject matter, has made it very difficult to patent black-box algorithms.[18] In *Mayo Collaborative Services v. Prometheus Laboratories, Inc.*, the Supreme Court repeated its long-standing statement that laws of nature cannot be patented.[19] However, the Court applied that rule to a diagnostic test that used the measurement of a metabolite level in a patient's blood to adjust the dosage of a drug, which many, including the Federal Circuit below, had thought to be a patentable application of such a law. The Supreme Court used very broad language to invalidate the patent: "[W]ell-understood, routine, conventional activity previously engaged in by scientists who work in the field . . . is normally not sufficient to transform an unpatentable law of nature into a patent-eligible application of such a law."[20] Where underlying information about the biological world is the heart of the invention, merely using that information to guide medical treatment is unpatentable (as is the information itself). But this describes most black-box algorithms quite well and suggests that those algorithms are unlikely to be patentable subject matter. Further patent problems might arise under section 112, which requires a "written description" of the invention. Although this issue has not been tested in the courts, it is at least debatable how well one can describe an algorithm that is opaque, and how broad the resulting protection would be.[21]

Trade secrecy—or secrecy in general—seems an obvious solution but comes with its own problems. Trade secret law protects from appropriation information that is kept secret and gets commercial value from its secrecy. What better way than secrecy to protect an algorithm that is already opaque and cannot be understood? The data on which an algorithm is generated, the method by which the algorithm was developed, and the process of its validation can all be kept secret by firms looking to protect their investment in the algorithm's development. And indeed, firms that are developing black-box algorithms seem to be relying on just such secrecy. But while secrecy may be an effective intellectual property strategy, it runs headlong into the concerns raised above about safety, malpractice, and regulation. How willing will doctors, patients, and insurers be to accept medical algorithms where not only is the working of the algorithm a mystery, but also the way the algorithm was made and tested, along with the data underlying its development? And if third parties are indeed to be actively involved in ensuring algorithmic quality and validity, as I suggest above, how can they conduct such evaluations without the underlying information? The reliance of algorithm developers on trade secrecy echoes other past situations where information relevant to public health has been kept secret, and these experiences suggest that there may be similar fights over access to algorithmic information.[22]

However, if intellectual property incentives are unavailable to help protect investments in black-box medicine, will firms invest sufficiently? How can the government help drive this form of innovation while ensuring that it is safe and effective? These questions are and will remain pressing for the development of AI in health care.

Privacy

Finally, privacy concerns run through the development and deployment of black-box medicine.[23] Privacy is important in at least two areas: gathering immense amounts of healthcare data to develop

algorithms, and sharing such data to oversee them. Algorithm developers need to assemble data from multiple sources to train machine learning algorithms. Those data—as well as data about how the algorithms perform in practice—may then be shared with other entities in the healthcare system for the purpose of evaluation and validation, as described above. In each case, patient-oriented data privacy is a concern, most notably as mandated under the Health Insurance Portability and Accountability Act's (HIPAA's) Privacy Rule. The Privacy Rule governs and restricts both disclosure and use of "protected health information" (that is, most individually identifiable health information) by "covered entities" (mostly, healthcare providers, health insurers, health information clearinghouses, and business associates of the same).[24] HIPAA creates a relatively complex set of permitted and restricted uses of protected health information. Notably, de-identified information is not governed by the Privacy Rule (though it raises its own concerns about data aggregation and the possibility of re-identification), and neither is information collected by noncovered entities like Google, Apple, or other aggregators of big data.[25] Navigating the HIPAA Privacy Rule—and otherwise managing and addressing the privacy concerns of those whose data is used throughout black-box medicine—creates yet another ongoing set of potential legal concerns.

Conclusion

Black-box medicine has tremendous potential to reshape health care, and it is moving rapidly to do so. Some healthcare black-box algorithms are already at work in consumer-directed smartphone apps, and others are likely to enter medical practice in the span of years. But the legal issues involved with the development and implementation of AI algorithms, which we do not and cannot understand, are substantial. As described here, regulation, legal causes of action such as medical malpractice and product liability, intellectual property, and patient privacy all have real implications for the way black-box medicine is developed and deployed. In turn, black-box medicine may change the way we approach some of these issues in the context of contemporary health care. Does entity-centered privacy regulation make sense in a world where giant data agglomerations are necessary and useful? Should intellectual property law find new ways to recognize the primacy of health data and the fast-moving nature of algorithms? Must the legal doctrine of the "learned intermediary" bow to the recognition that doctors cannot fully understand all the technologies they use or the choices such technologies help them make when they are not provided the needed and/or necessary information? Should the FDA change how it regulates new medical technology as AI software gains prominence? As black-box medicine develops and evolves, the need to consider these legal issues—and the need for scientifically literate lawyers who can understand them in context—will continue to grow. ◆

Endnotes

1. W. Nicholson Price II, *Black-Box Medicine*, 28 Harv. J.L. & Tech. 419 (2015).

2. Jenna Burrell, *How the Machine "Thinks": Understanding Opacity in Machine Learning Algorithms*, 3 Big Data & Soc'y 1, 5 (2016).

3. I. Glenn Cohen et al., *The Legal and Ethical Concerns That Arise from Using Complex Predictive Analytics in Health Care*, 33 Health Aff. 1139 (2014).

4. Andre Esteva et al., *Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks*, 542 Nature 115 (2017).

5. Ziad Obermeyer & Ezekiel J. Emanuel, *Predicting the Future—Big Data, Machine Learning, and Clinical Medicine*, 375 New Eng. J. Med. 1216 (2016).

6. Nehemiah T. Liu et al., *Development and Validation of a Machine Learning Algorithm and Hybrid System to Predict the Need for Life-Saving Interventions in Trauma Patients*, 52 Med. & Biological Engineering & Computing 193 (2014).

7. Megan Molteni, *Thanks to AI, Computers Can Now See Your Health Problems*, Wired (Jan. 9, 2017), https://www.wired.com/2017/01/computers-can-tell-glance-youve-got-genetic-disorders/.

8. *Autism*, Right Eye, https://www.righteye.com/tests-therapies/autism (last visited Oct. 17, 2017).

9. Patricia J. Zettler, *Toward Coherent Federal Oversight of Medicine*, 52 San Diego L. Rev. 427 (2015).

10. W. Nicholson Price II, *Regulating Black-Box Medicine*, 91 Mich. L. Rev. (forthcoming 2017), https://papers.ssrn.com/abstract_id=2938391.

11. *See* U.S. Food & Drug Admin., Medical Device Accessories—Describing Accessories and Classification Pathway for New Accessory Types: Guidance for Industry and Food and Drug Administration Staff (Jan. 30, 2017).

12. Press Release, FDA, FDA Selects Participants for New Digital Health Software Precertification Pilot Program (Sept. 26, 2017), https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm577480.htm.

13. Randolph A. Miller & Sarah M. Miller, *Legal and Regulatory Issues Related to the Use of Clinical Software in Health Care Delivery*, *in* Clinical Decision Support: The Road Ahead 423, 426 (Robert A. Greenes ed., 2007).

14. W. Nicholson Price II, *Medical Malpractice and Black-Box Medicine*, *in* Big Data, Health Law, and Bioethics (I. Glenn Cohen et al. eds., forthcoming 2018).

15. Megan Molteni, *If You Look at X-Rays or Moles for a Living, AI Is Coming for Your Job*, Wired (Jan. 25, 2017), https://www.wired.com/2017/01/look-x-rays-moles-living-ai-coming-job/.

16. W. Nicholson Price II, *Big Data, Patents, and the Future of Medicine*, 37 Cardozo L. Rev. 1401 (2016).

17. Mark A. Lemley, *Ex Ante versus Ex Post Justifications for Intellectual Property*, 71 U. Chi. L. Rev. 129 (2004).

18. Rebecca S. Eisenberg, *Diagnostics Need Not Apply*, 21 B.U. J. Sci. & Tech. L. 256 (2015).

19. 132 S. Ct. 1289 (2012).

20. *Id.* at 1298.

21. W. Nicholson Price II, *Describing Black-Box Medicine*, 21 B.U. J. Sci. & Tech. L. 347 (2015).

22. David S. Levine, *The People's Trade Secrets?*, 18 Mich. Telecomm. & Tech. L. Rev. 61 (2011).

23. Roger Allan Ford & W. Nicholson Price II, *Privacy and Accountability in Black-Box Medicine*, 23 Mich. Telecomm. & Tech. L. Rev. 1 (2016).

24. 45 C.F.R. pts. 160, 164.

25. Nicolas Terry, *Big Data and Regulatory Arbitrage in Health Care*, *in* Big Data, Health Law, and Bioethics, *supra* note 14.