The Challenge of Regulating Clinical Decision Support Software after 21st Century Cures

Barbara Evans[†] & Pilar Ossorio[††]

I.   INTRODUCTION

The U.S. Food and Drug Administration ("FDA") for many years has regulated "software in a medical device"[1]—software embedded in traditional devices like pacemakers, drug infusion pumps, and *in vitro* diagnostic ("IVD") test kits where the software affects the safety and effectiveness of the device as a whole.[2] In 2013, as the use of software grew pervasive in health care, FDA worked with its counterpart agencies from other nations to define "software as a medical device" ("SaMD"): stand-alone medical software, designed to run on diverse platforms such as personal computers, smartphones, or in the cloud, that constitutes a medical device in its own right.[3] Alarmed that FDA might be embarking on a broad program to regulate stand-alone medical software, the software industry pressed Congress for clarification. Section 3060 of the 21st Century Cures Act (hereinafter, the "Cures Act")[4] was Congress's response.

Section 3060(a) removes five categories of software from FDA's jurisdiction by excluding them from the Food, Drug, and Cosmetic Act's definition of a "medical device."[5] The five exclusions appear at 21 U.S.C. § 360j(o)[6] and encompass: (A) healthcare business software such as systems for billing, scheduling appointments, tracking population health and the cost-effectiveness of care, and managing laboratory workflows;[7]; (B) software for encouraging wellness or a healthy lifestyle, provided it does not cross the line into "diagnosis, cure, mitigation,

[1] *See Software as a Medical Device (SaMD): Key Definitions*, INT'L MEDICAL DEVICE REGULATOR'S FORUM (Dec. 9, 2013), http://www.imdrf.org/docs/imdrf/final/technical/imdrf-tech-131209-samd-key-definitions-140901.pdf (distinguishing software in a device from software as a device).
[2] *Id.*; *see also What are examples of Software as a Medical Device?* U.S. FOOD & DRUG ADMIN. (last updated Dec. 6, 2017), https://www.fda.gov/medicaldevices/digitalhealth/softwareasamedicaldevice/ucm587924.htm.
[3] *Id.*
[4] *See* 21st Century Cures Act, Pub. L. No. 114-255, § 3060(a), 130 Stat. 1033 (2016).
[5] *See* 21 U.S.C. § 321(h) (codifying § 201(h) of the Food, Drug, and Cosmetic Act, which defines the medical devices that Congress has authorized FDA to regulate).
[6] See 21 U.S.C. § 360j(o)(1) (excluding five categories of software from the definition of a medical device).
[7] *Id.* at § 360j(o)(1)(A).

prevention, or treatment of a disease or condition"[8]; (C) Electronic Health Record ("EHR") software, provided it does not interpret or analyze the data for the purpose of "diagnosis, cure, mitigation, prevention, or treatment of a disease or condition"[9]; (D) Medical Device Data System ("MDDS") software that transfers, stores, converts formats, or displays clinical laboratory test results or other health information, once again provided that it does not cross the line into interpreting and analyzing the data;[10] and (E) a subset of Clinical Decision Support ("CDS") software,[11] the focus of this article.

Some of these software categories, arguably, were never subject to FDA regulation in the first place, either because they do not fit within the device definition[12] or because of underlying First Amendment constraints on the regulation of software.[13] Confusion persisted, however, and Section 3060(a) reflects Congress's attempt to clarify matters by expressly stripping FDA of jurisdiction to regulate certain types of medical software. Even before the Cures Act, FDA disclaimed its intent to regulate several software categories[14] and, for these, the Cures Act merely extends FDA's prior policies, making them mandatory rather than an exercise of enforcement discretion. The fifth exclusion, for CDS software,[15] addresses questions that remain less settled.

FDA acknowledges that the term CDS "is used broadly and in different ways, depending on the context,"[16] but the Cures Act excludes only a subset of CDS software—not all of it—from

---

[8] *Id.* at § 360j(o)(1)(B)

[9] *Id.* at § 360j(o)(1)(C)(iii).

[10] *Id.* at § 360j(o)(1)(D).

[11] *Id.* at § 360j(o)(1)(E).

[12] *See* Jeffrey K. Shapiro & Jennifer D. Newberger, *Highlights of Medical Device Related Provision of 21ˢᵗ Century Cures Act*, FDA L. BLOG (Dec. 8, 2016), http://www.fdalawblog.net/2016/12/highlights-of-medical-device-related-provision-in-the-21st-century-cures-act/ [https://perma.cc/XMN3-VPX6] (noting that administrative support software arguably never was a medical device in the first place).

[13] *See* R. Polk Wagner, *The Medium is the Mistake: The Law of Software for the First Amendment*, 51 STAN. L. REV. 387 (1999).

[14] See Shapiro & Newberger, *supra* note 12 (noting that the Cures Act's exclusion for software for maintaining and encouraging a healthy lifestyle is "non-controversial and is consistent with FDA's General Wellness and Mobile Medical Applications guidance documents" and noting that the exclusion for electronic health record systems "codifies policy already implemented by FDA as a matter of enforcement discretion" and that its exclusion for software for transferring, storing, formatting, or displaying data merely codifies a prior FDA exemption for medical device data systems (MDDS)).

[15] 21 U.S.C. § 360j(o)(1)(E).

[16] *Clinical and Patient Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff*, FOOD & DRUG ADMIN. (Dec. 8, 2017), https://www.fda.gov/downloads/MedicalDevices/DeviceRegulationandGuidance/GuidanceDocuments/UCM587819.pdf; *see also Clinical Decision Support (CDS)*, HEALTHIT.GOV (last updated Jan. 15, 2013), https://www.healthit.gov/policy-researchers-implementers/clinical-decision-support-cds [https://perma.cc/JWV8-YUGQ] (noting that CDS,

> provides clinicians, staff, patients or other individuals with knowledge and person-specific information, intelligently filtered or presented at appropriate times, to enhance health and health care. CDS encompasses a variety of tools to enhance decision-making in the clinical workflow. These tools include computerized alerts and reminders to care providers and patients; clinical guidelines; condition-specific order sets; focused patient data reports and summaries; documentation templates; diagnostic support, and contextually relevant reference information, among other tools.).

FDA oversight.[17] The exclusion is subject to a limiting clause that lets FDA continue to regulate software functions "intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system."[18] An example would be software that enhances diagnostic images, such as mammograms or dental X-rays, to highlight areas suspicious for disease.[19] Such software has long been, and after the Cures Act still is, an FDA-regulated medical device.

The Cures Act focuses on CDS software that harnesses patient-specific information (such as "diagnosis, treatments, allergies"[20] or "demographic information, symptoms, and test results"[21]) and uses this information, often in combination with external (non-patient-specific) sources of medical knowledge,[22] for the purpose of "supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition."[23] Simple software that bases its recommendations strictly on patient-specific information—for example, "Do not prescribe this drug to Mary because she is allergic to it"—seemingly lies outside this definition and is more in the nature of EHR software, if the purpose is merely to shield physicians from liability for failing to heed information already in the patient's chart.

The Cures Act singles out CDS software that recommends diagnoses or actions to treat or prevent disease. It defines a standard for deciding when such software can be excluded from FDA regulation. Congress excludes CDS software from FDA regulation if the software is intended to enable the "health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient."[24] To escape FDA regulation, the software vendor/manufacturer must intend for the software to make it possible for health care professionals to override its recommendations by explaining its rationale in terms that a clinician could understand, interrogate, and possibly reject. Whether CDS software is subject to FDA regulation potentially turns on the software's ability to answer the quintessential epistemological[25] question: How do we know?

This article explores whether the line Congress drew in Section 3060(a) is a workable standard: Will FDA be able to tell when CDS software explains its recommendations in a way that physicians can understand and critique? More to the point, will FDA be able to tell when the manufacturer *intended* for its software to do so? If not, what problems may clinicians face in using CDS software?

---

[17] *See* 21 U.S.C. § 360j(o)(1)(E).

[18] *Id.*

[19] Bradley Merrill Thompson, *Learning from Experience: FDA's Treatment of Machine Learning*, MOBILE HEALTH NEWS (Aug. 23, 2017), http://www.mobihealthnews.com/content/learning-experience-fda%E2%80%99s-treatment-machine-learning [https://perma.cc/Q95C-9R22].

[20] *Draft Guidance*, *supra* note 16, at 8.

[21] *Id.* at 7.

[22] *See* 21 U.S.C. § 360j(o)(1)(E)(i) (describing software that uses patient specific information "or" general medical information). *But see* FDA, *Draft Guidance*, *supra* note 16, at 7 (sensibly reading this "or" as meaning "and/or").

[23] *Id.* at § 360j(o)(1)(E)(ii).

[24] *Id.* at § 360j(o)(1)(E)(iii)

[25] *See* STANFORD DICTIONARY OF PHILOSOPHY (Dec. 14, 2005), https://plato.stanford.edu/entries/epistemology/ [https://perma.cc/M2KE-22NJ] (defining epistemology as "the study of knowledge and justified belief. As the study of knowledge, epistemology is concerned with the following questions: What are the necessary and sufficient conditions of knowledge? What are its sources? What is its structure, and what are its limits?").

II.   CDS SOFTWARE AND THE ROLE OF THE PHYSICIAN

Some of the most promising CDS software applications are those that combine patient-specific data with external sources of medical knowledge, highlighting the relevance of that knowledge to the patient's particular circumstances. This article sorts such software into two broad categories, depending on whether the external knowledge is already well-established or is inferred by the CDS software itself. They reflect two different approaches for improving the quality of clinical care.

The first type of CDS software relies on existing sources of medical information such as clinical practice guidelines, published medical compendia, established clinical practices, information from FDA-approved drug labeling, peer-reviewed literature, or lists of genetic variants that have a well-established clinical significance. An example would be CDS software that combines patient-specific information about the medicines a person is already taking with external knowledge gleaned from FDA-approved drug labeling, with the goal of alerting healthcare workers who are about to administer a new medicine that is incompatible with the patient's prior drug regimen. Such software seeks to improve the quality of healthcare by enhancing the application of existing medical knowledge—an approach that offers considerable prospects for quality improvement in a healthcare system where physicians widely disregard warnings in drug labeling,[26] and compliance with clinical practice guidelines is often incomplete and long-delayed.[27]

In the second type of CDS software, the external source of knowledge might itself be an algorithm, possibly incorporating machine learning or other artificial intelligence ("AI") techniques,[28] to glean real-world insights from clinical experience. An algorithm is simply a procedure for transforming input data into a desired output based on specified calculations. A machine-learning algorithm is one that, over time, improves its ability to perform some task, such as classifying patients according to whether they have or do not have a particular disease, or

---

[26] *See, e.g.*, Walter Smalley et al., *Contraindicated Use of Cisapride: Impact of Food and Drug Administration Regulatory Action*, 284 JAMA 3036, 3038 (2000) (finding that labeling revisions and efforts to communicate contraindications of the drug cisapride (Propulsid) had little impact on prescribing behavior); Raymond L. Woosley & Glenn Rice, *A New System for Moving Drugs to the Market*, 21 ISSUES SCI. & TECH. ONLINE, 63, 64 (2005), (listing six drugs that had to be withdrawn from the market because of physician non-compliance with warnings and contraindications in labeling).

[27] Michelle M. Mello, *Of Swords and Shields: The Role of Clinical Practice Guidelines in Medical Malpractice Litigation*, 149 U. PA. L. REV. 645, 681 (2001).

[28] *See What is Machine Learning? A definition*, EXPERT SYS., http://www.expertsystem.com/machine-learning-definition/ [https://perma.cc/S88G-2REA]

> (Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it learn for themselves. The process of learning begins with observations or data, such as examples, direct experience, or instruction, in order to look for patterns in data and make better decisions in the future based on the examples that we provide. The primary aim is to allow the computers learn automatically without human intervention or assistance and adjust actions accordingly.).

See also Maxwell W. Libbrecht & William Stafford Noble, *Machine Learning in Genetics and Genomics*, 16 NATURE REV. GENETICS 321, 321 (2015) (describing machine learning algorithms as those that "improve with experience" at the task of recognizing patterns in data, to "assist humans in making sense of large, complex data sets").

predicting who is most likely to develop a disease or benefit from a particular course of treatment. When an algorithm "learns," [29] it changes something in its programming, such as the variables it considers, the weights, or probabilities. It learns by developing a better computational model that more effectively carries out the prescribed task.

CDS software of this latter type processes large volumes of data, such as outcomes data from real-world patient experience, to infer new general medical knowledge that can then be combined with patient-specific data to produce recommendations.[30] For example, the CDS software might incorporate predictive analytics software that continuously mines data about other patients treated in a hospital's Intensive Care Unit (ICU), with the aim of identifying attributes that predict which patients are most likely to achieve good health outcomes after ICU care.[31] The software would combine the attributes of a particular patient, Mary, with this inferred knowledge to predict whether admitting Mary to the ICU is likely to help her or be futile, and then recommend a course of action. CDS software of this type seeks to improve the quality of healthcare by pushing past existing medical knowledge. This approach acknowledges that existing medical knowledge may be flawed: e.g., FDA-approved labeling is based on clinical trials that may not reflect real patients in real clinical settings;[32] peer-reviewed literature is subject to publication bias that favors studies in which the treatment worked and clinical practice guidelines sometimes are driven by commercial interests.[33] A well-validated CDS algorithm fueled by real-world outcomes data collected at the point of care has the potential to make better decisions than the application of established knowledge would do.[34]

Cohen et al. note the potential for such software to raise thorny ethical issues, for example, if ICU beds are scarce and the software's recommendations influence resource allocations, such as whether the one available bed should go to Mary or to Sue.[35] In non-medical contexts such as approving loans, machine-learning algorithms have a disturbing potential to achieve discriminatory outcomes by relying of facially neutral criteria (e.g., zip codes) that have underlying correlations with membership in protected classes (e.g., race, ethnicity, or gender).[36] What if CDS software preferentially admits white males to the ICU because outcomes data show "they do better"? Other salient ethical concerns relate to the role of physicians, who are bound to act in the

---

[29] TOM M. MITCHELL, MACHINE LEARNING 2 (1997) (noting that "a computer program is said to *learn* from experience E with respect to some class of tasks T and performance measure P, if its performance at tasks in T, as measured by P, improves with experience E").

[30] *See* I. Glenn Cohen et al., *The Legal and Ethical Concerns That Arise from Using Complex Predictive Analytics in Health Care*, 33 HEALTH AFF. 1139, 1140 (2014).

[31] *See id.* at 1139-47 (discussing examples of predictive analytics software used in health care settings).

[32] See INST. OF MED., THE FUTURE OF DRUG SAFETY 39, 122 (Alina Baciu et al. eds., 2007), http://books.nap.edu/openbook.php?record_id=11750 [https://perma.cc/TD5P-7SAN] (noting the prevalence and lack of evidence for off-label uses and noting the limitations of clinical trial data to inform decisions about the real-world clinical impacts even for a drug's on-label, indicated uses).

[33] INST. OF MED., CONFLICT OF INTEREST IN MEDICAL RESEARCH, EDUCATION, AND PRACTICE 189-215 (Bernard Lo & Marilyn J. Field eds., 2009).

[34] *See* C.A. Longhurst et al., *A 'Green Button' for Using Aggregate Patient Data at the Point of Care*, 33 HEALTH AFF. 1229, 1229 (2014).

[35] Cohen et al., *supra* note 30, at 1139.

[36] *See generally* Danielle Keats Citron & Frank Pasquale, *The Scored Society: Due Process for Automated Predictions*, 89 WASH. L. REV. 1, 14-15 (2014).

best interest of individual patients, while the default rules driving CDS software may give weight to the healthcare institution's broader population health and financial objectives.[37]

Institutions implementing CDS software must carefully weigh whether, and when, to allow physicians to override the default recommendations that a CDS algorithm provides. Whether a physician override is good or bad depends on many factors, such as the purpose of the software, how well validated it is, the context, and the physician's level of knowledge and experience in specialty areas that may be relevant to a particular patient-care decision, such as medical genetics. A physician override may be counterproductive, if the physician is ignoring relevant, established knowledge (served up by a simple CDS software) or is relying on established knowledge that is biased, corrupted, or unreflective of the patient's circumstances (as revealed by a more sophisticated CDS algorithm that incorporates outcomes data).[38] On the other hand, the physician may have access to facts that the algorithm was not programmed to consider.[39]

In passing the Cures Act, Congress determined that a physician override adds a significant layer of patient safety to a CDS algorithm—and, we suspect, most patients would probably agree and want their physician kept in the loop. This determination is implicit in Section 3060(a), which excludes CDS software from FDA regulation if the software is intended to explain its recommendations in a transparent way that enables the physician to understand, critique, and perhaps override them.[40]

III.   DID THE CURES ACT SET A WORKABLE STANDARD?

To be excluded from FDA oversight, CDS software is not required to *succeed* in explaining its recommendations to physicians in a way that enables independent physician review. The software manufacturer merely needs to *intend* for the software to explain its recommendations transparently.[41] Is it even possible to assess this?

FDA has regulations describing the evidence it considers when assessing a manufacturer's intended use of a drug or a device.[42] Revisions to these regulations in 2017 were postponed multiple times before being indefinitely delayed.[43] Under the current regulation at 21 C.F.R. § 801.4, FDA considers objective evidence of the software manufacturer's intent.[44] This includes

---

[37] Cohen et al., *supra* note 30, at 1146-47.

[38] See Jason B. Liu et al., *Defining the Intrinsic Cardiac Risks of Operations to Improve Preoperative Cardiac Risk Assessments,* 128 ANESTHESIOLOGY 283, 285 (2018) (finding that an established index for predicting the risk of cardiac arrest during surgery was less accurate than an algorithm that incorporates data from over 3 million actual operations).

[39] See J.F. Peterson et al., *Physician Response to Implementation of Genotype-Tailored Antiplatelet Therapy*, 100 CLINICAL PHARMACOLOGY & THERAPEUTICS 67, 72 (2016) (finding various reasons for instances of physician non-compliance with recommendations by a CDS system incorporating pharmogenetic data).

[40] 21 U.S.C. § 360j(o)(1)(E)(iii).

[41] *See id.*

[42] 21 C.F.R. § 201.128 (2018) (for drugs); 21 C.F.R. § 801.4 (2017) (for devices).

[43] *See* Clarification of When Products Made or Derived from Tobacco Are Regulated as Drugs, Devices, or Combination Products, 83 Fed. Reg. 2092, 2093 (proposed Jan. 16, 2018) (to be codified at 21 C.F.R. pts. 201, 801, 1100).

[44] 21 C.F.R. § 801.4 (2018) (defining "intended use" as the "objective intent of the persons legally responsible for the labeling of devices").

direct evidence—such as claims the manufacturer or its representatives make in labeling, advertising, or oral and written statements—and circumstantial evidence that the software is being used, with the manufacturer's knowledge, "for a purpose for which it is neither labeled nor advertised."[45]

In practice, the direct evidence—statements of intent made by the manufacturer and its representatives—tend to be dispositive. FDA has had little success ascribing intent to manufacturers, based on circumstantial evidence of a known misuse.[46] If a software manufacturer states that its CDS software is *intended* to enable independent review of its recommendations by a physician and is *not intended* for use as the primary basis for physician decision-making, these statements are likely to carry the day. Evidence that physicians do not actually understand the basis of the software's recommendations, or that they actually misuse the software as their primary basis for decision-making, would not tend to alter this conclusion, even if the manufacturer is aware of the misuse.

This suggests that the Section 3060(a) exclusion criterion may be unworkable, and that a manufacturer can invoke this exclusion—and escape FDA regulation of its software—merely by asserting the requisite intent. There is, we believe, a crucial nuance: Evidence that a device *could not possibly* be used as intended constitutes relevant circumstantial evidence that FDA could use to negate a manufacturer's stated intent. If a manufacturer intends its device (e.g., a pig) to be used in a way it could not possibly be used (e.g., to fly), then FDA need not credit the manufacturer's claims that the intended use of the pig is to fly. Circumstantial evidence negating intent may be relevant, even when circumstantial evidence establishing intent is not. If true, the standard Congress enunciated at 21 U.S.C. § 360j(o)(1)(E)(iii) may be workable after all. FDA can negate a software manufacturer's intent to enable a physician override, if FDA can establish that CDS software is so opaque that the healthcare providers who are its intended users are unlikely to be able to understand the basis of its recommendations.

IV.   CAN CDS SOFTWARE EXPLAIN ITSELF IN TERMS A DOCTOR CAN UNDERSTAND?

Readers who follow European privacy law may already have noticed that Section 3060(a) presents some of the same problems seen with the "right to explanation" under Article 22 of the 2016 EU General Data Protection Regulation ("GDPR"),[47] which takes effect in April 2018. Article 22, entitled "Automated individual decision-making, including profiling," grants data subjects (the people that data describe) a right "not to be subject to a decision based solely on automated processing" if the decision produces "significant effects on him or her."[48] It has

---

[45] *Id.*

[46] *See*, *e.g.*, National Nutritional Foods Association v. Matthews, 557 F.2d 325 (2d Cir. 1977), *but see* U.S. v. Travia, 180 F. Supp. 2d 115 (D.D.C. 2001) (finding intent based on actual use in the absence of any claims indicating the manufacturer intended such use, but doing so in a factual setting so unusual—unlawful sale of balloons full of laughing gas at a rock concert—that the precedential value of this holding can be questioned).

[47] *See* Regulation 2016/679, of the European Parliament and of the Council of 27 April 2016 on the Protection of Natural Persons with Regard to the Processing of Personal Data and on the Free Movement of Such Data, and Repealing Directive 95/46/EC (General Data Protection Regulation), (L 119/1) 1. (hereinafter, "EU-GDPR")).

[48] Council Regulation 2016/679, art. 22, 2018 (L 119/1).

exceptions that allow automated decision-making,[49] but only if "suitable measures" are in place to protect the data subjects' rights.[50] These measures must include "at least" the rights for data subjects to obtain human intervention, to express their own point of view, and to contest the decision.[51] Article 22 specifically prohibits algorithmic decision-making based on special categories of personal data—including, among other things, genetic and health data[52]—unless these protections are in place.[53]

      Some scholars read Article 22 as creating "a 'right to explanation,' whereby a user [i.e., the data subject] can ask for an explanation of an algorithmic decision that was made about them."[54] Other scholars disagree that Article 22 implies such a right.[55] What scholars do seem to agree on is that a right of explanation, if it does exist, will be challenging to implement.[56] The difficulty is that machine-learning algorithms, even when they make an accurate prediction, may not be able to explain the basis of their predictions in terms intelligible to a human being. They search data for associations and correlations that support accurate predictions, without addressing why.[57] For some kinds of machine learning algorithms, nobody can determine which factors the algorithm used to make its predictions. This problem is particularly acute for deep learning neural

---

[49] *Id.* at ¶ 2 (allowing decisions to be based solely on automated processing in narrow circumstances, e.g., if authorized by law or if the data subject consents).

[50] *Id.* at paras. 2(b) and (3).

[51] *Id.* at para. 3.

[52] Council Regulation 2016/679, art. 9, 2018 (L 119/1) (defining the special categories of data).

[53] Council Regulation, *supra* note 48, at para. 4.

[54] Bryce Goodman & Seth Flaxman, *European Union Regulations on Algorithmic Decision-Making and a "Right To Explanation",* ARXIV (June 28, 2016), https://arxiv.org/abs/1606.08813 [https://perma.cc/UT2E-8D76].

[55] *See* Andrew Burt, *Is There a Right to Explanation for Machine Learning in the GDPR?*, INT'L ASS'N OF PRIVACY PROF'LS (June 1, 2017), https://iapp.org/news/a/is-there-a-right-to-explanation-for-machine-learning-in-the-gdpr/ [https://perma.cc/3879-JEJA] (noting that the existence of a right to explanation "has become a controversial subject. Some scholars, for example, have spoken out vehemently against the mere possibility that such a right exists. Others, such as the UK's own Information Commissioner's Office, seem to think the right is pretty clearly self-evident"). *See also* Sandra Wachter, Brent Mittelstadt & Luciano Floridi, *Why a Right to Explanation of Automated Decision-Making Does Not Exist in The General Data Protection Regulation*, INT'L DATA PRIVACY LAW (2017), https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2903469 [https://perma.cc/VP4K-5R9S].

[56] *See* Goodman & Flaxman, *supra* note 54, at 6 – 7 (stating, in connection with neural networks, "what hope is there of explaining the weights learned in a multilayer neural net with a complex architecture?"); Burt, *supra* note 55 (stating that "if you're a privacy professional, you're going to find it difficult to implement these requirements in practice").

[57] See Goodman & Flaxman, *supra* note 54, at 6 (noting that,
    [s]tandard supervised machine learning algorithms for regression or classification are inherently
    based on discovering reliable associations / correlations to aid an accurate out-of-sample
    prediction, with no concern for causal reasoning or "explanation" beyond the statistical sense in
    which it is possible to measure the amount of variance explained by a predictor.).

networks,[58] which may internally represent data in a manner that has no "direct physical meaning."[59]

This problem was aired in the popular press recently, after researchers trained a machine-learning algorithm by letting it scan over 200,000 facial photos people had posted on dating web sites, along with statements the people had posted about their personal and political views and sexual orientation.[60] Controversially, the algorithm taught itself to predict, merely by inspecting a person's facial photograph, whether the person self-identifies as straight or gay. The machine's predictions reportedly have 91% accuracy for males and 83% accuracy for females, whereas human beings can only make such predictions with 60% accuracy—i.e., humans barely beat the 50/50 accuracy of random guessing.[61] The algorithm outperforms human judgment, but analysts are unsure how it does so. Has the software spotted a minor nuance of facial expression that humans never noticed, or is it assigning weight to factors such as whether the person wears a baseball cap or their preferred styles of posing for photographs, or other factors?[62]

A machine-learning CDS algorithm might, hypothetically, notice that past patients who drive green cars tend to survive cancer better than those with red cars, and apply this knowledge to make accurate prognoses for future patients, without there being any plausible biological mechanism to "explain" this result in terms a physician would find useful. Moreover, in real-world data environments like genomics and health care, algorithmic decisions may rest on many factors, each exerting a small influence and with complex interplays among the factors. There may not be a compact set of factors that can be said to "explain" a decision even from a statistical standpoint. Human minds can only wield a few factors at once. Explanations that invoke too many factors are not intelligible to us. As one observer put it, there is a "mismatch between the mathematical optimization in high-dimensionality characteristic of machine learning and the demands of human-scale reasoning."[63] Persons with the requisite technical expertise could inspect the code and the datasets used to train an algorithm and possibly come to trust it after doing so, but for most people, reading computer code "explains" nothing.

---

[58] *See, e.g.*, Rich Caruana et al., *Intelligible Models for Healthcare: Predicting Pneumonia Risk and Hospital 30-day Readmission*, ASS'N FOR COMPUTING MACHINERY KDD (2015), http://dx.doi.org/10.1145/2783258.2788613 [https://perma.cc/9LT9-TXFB]; A. Fischer, *How to Determine the Unique Contributions of Input-Variables to the Nonlinear Regression Function of a Multilayer Perception*, ECOLOGICAL MODELLING 309, 310, 60, 63 (2015).

[59] Tim Albrecht et al., *Deep Learning for Single-Molecule Science*, 28 NANOTECHNOLOGY 1, 9 (2017), https://doi.org/10.1088/1361-6528/aa8334.

[60] Cliff Kuang, *Can AI be Taught to Explain Itself?*, N.Y. TIMES (Nov. 21, 2017), https://www.nytimes.com/2017/11/21/magazine/can-ai-be-taught-to-explain-itself.html; *See also*, Tim Collins, *Outrage over AI that "identifies gay faces' as Google experts say the machine relies on patterns in how people take selfies and NOT on their facial features*, DAILY MAIL (Jan. 15, 2018), http://www.dailymail.co.uk/sciencetech/article-5270365/Google-experts-debunk-sexuality-detecting-AI.html [https://perma.cc/UH28-KE6F].

[61] Kuang, *supra* note 60.

[62] *Id.*

[63] Goodman & Flaxman, *supra* note 54 (quoting J. Burrell, *How the machine "thinks": Understanding opacity in machine-learning algorithms*, 3 BIG DATA AND SOCIETY (2016)).

V. FOR THE CURES ACT TO WORK, FDA NEEDS TO ENUNCIATE TRANSPARENCY STANDARDS

A recent CDS draft guidance[64] discusses how FDA plans to apply 21 U.S.C. § 360j(o)(1)(E), the statute that excludes certain CDS systems from the definition of a medical device. As guidance documents go, this one is disappointing.[65] It avoids discussing the most glaring question the statute presents: How does FDA intend to distinguish machine-learning software that can explain its recommendations to a physician from software that cannot do so? Under the Cures Act, this distinction has regulatory significance: it affects a CDS system's eligibility to be excluded from FDA oversight. Unless FDA can wield this distinction, the agency has no way to negate a manufacturer's claim that it has the requisite intent to qualify for an exclusion from FDA regulation.

There are various approaches the draft guidance might have taken. Significant research is underway to develop "explainable artificial intelligence" (dubbed "XAI").[66] Much of this research is in the military sphere where—a little chillingly—there is growing recognition that "future warfighters" will need to be able to "understand, appropriately trust, and effectively manage an emerging generation of artificially intelligent machine partners."[67] (Yes, please manage them!) One possible approach would be for FDA to summon the best minds on XAI and human psychology to develop a list of features FDA should look for when deciding whether CDS software is transparent or inscrutable. For example, to be understandable, does an algorithm need to explain its decisions in terms of fewer than 50 variables, 10 variables, 5 variables? Do its associations need to be backed by a plausible causal mechanism, or will well-described statistical associations suffice if clearly explained? Just what does intelligibility mean?

An alternative approach would be for FDA to require empirical evidence that CDS recommendations are understandable to doctors. For example, the CDS draft guidance could have announced that to qualify for the exclusion at §360j(o)(1)(E), CDS system developers must empanel a group of doctors, hand them output from CDS systems, and ask, "Does this make any sense to you?" This is not an exhaustive list of possible approaches.

The draft guidance does not discuss machine-learning algorithms at all. It confirms that very simple CDS systems—those that incorporate existing knowledge that is "publicly available (e.g., clinical practice guidelines, published literature)"[68]—meet the §360j(o)(1)(E)(iii) "explainability" standard and thus qualify to escape FDA regulation.[69] The draft guidance offers examples of excluded CDS systems that rely on "FDA-approved drug labeling," "established

---

[64] *Draft Guidance*, *supra* note 16.

[65] *See*, *e.g*., Evan Sweeney, *After a 6-year wait, FDA's Clinical Decision Support Guidelines Get a Mixed Reaction*, FIERCE HEALTHCARE (Dec. 7, 2017), https://www.fiercehealthcare.com/regulatory/fda-clinical-decision-support-bradley-merrill-thompson-bethany-hills-ai-21st-century [https://perma.cc/BT9T-TK7J].

[66] David Gunning, *Explainable Artificial Intelligence (XAI)*, DARPA, https://www.darpa.mil/program/explainable-artificial-intelligence [https://perma.cc/97MG-CSH9].

[67] *Id.* (noting that "the effectiveness of [AI] systems is limited by the machine's current inability to explain their decisions and actions to human users" and describing XAI as new or modified machine-learning techniques that "will produce more explainable models" with "the ability to explain their rationale, characterize their strengths and weaknesses, and convey an understanding of how they will behave in the future").

[68] *Draft Guidance*, *supra* note 16, at 8.

[69] *Id.*

guidelines," and "generally accepted clinical practice."[70] CDS systems of this sort, in effect, serve as electronic libraries that store and organize already-existing knowledge; some of them also synthesize conclusions based on that knowledge. With or without the Cures Act, libraries are not within the definition of medical devices that FDA can regulate, because the First Amendment constrains federal agencies' power to regulate libraries.[71] The draft guidance takes an important 21st-century statute and interprets it as saying no more than the Founding Fathers already told us in the Bill of Rights: It is a bad idea for the federal government to regulate libraries. As for the function of synthesizing original conclusions from the content of a library, that activity, too, receives significant First Amendment protection. If the conclusions are irresponsible and medical in nature, this seemingly is a matter for state medical practice regulators rather than FDA. FDA seemingly agreed, and excluded such software from FDA oversight.

The draft guidance does not, of course, rule out the possibility that other, more sophisticated CDS systems (e.g., those that infer general knowledge from real-world data) might also meet the § 360j(o)(1)(E)(iii) "explainability" criterion. It merely offers no insight on how to make that happen, or how FDA plans to assess whether it *has* happened or, perhaps, could not possibly happen. This last omission is a crucial one. Unless FDA can establish criteria—or at least a process—for rejecting claims that software is intended to be transparent to physicians, the standard in Section 3060(a) of the Cures Act is unenforceable. Manufacturers can escape FDA oversight merely by asserting the requisite intent.

If FDA deems all but the simplest CDS systems to be "unexplainable," this may have detrimental impacts on innovation and on patients. Simple CDS systems that merely compile libraries of established clinical practice guidelines, FDA-approved labeling, and peer-reviewed literature would qualify for the exclusion and be sped to market. More sophisticated CDS algorithms that enable a learning healthcare system, informed by real-world evidence, could face long regulatory delays. Such delays could deny patients the potential benefits of having their treatment decisions informed by clinical experience, which is a hoped-for antidote to defects in the evidentiary basis of today's evidence-based medicine. Another problem is that manufacturers of sophisticated CDS systems will have little incentive to invest research dollars in making their algorithms more "explainable," if FDA plans to deem all but the simplest systems to be unexplainable.

Also of concern are the impacts on physicians using CDS software. If FDA has no objective method to assess whether a CDS system is "explainable," the agency may be left in the position of taking manufacturers' word for it. Manufacturers of machine-learning algorithms may tend to label them for "use only under the supervision of a physician," to invoke the exclusion from FDA regulation. Product labeling will exhort physicians to "independently verify the recommendations" and warn them not to rely on the algorithm's recommendations as a primary basis for medical decision-making. Manufacturers may label their products this way even when the explanation that the algorithm provides is so opaque that a reasonable physician cannot actually critique the recommendations (remember, FDA has provided no objective standard or process for establishing when such claims are untrue). CDS systems labeled this way will enter the market without

---

[70] *Id.* at 8-9.

[71] *See, generally*, Susan Nevelow Mart, *The Right to Receive Information*, 95 LAW LIBR. J. 175 (2003) (discussing the permissible scope of governmental regulation of libraries under the First Amendment).

regulatory oversight. Once on the market, physicians may use such systems off-label: relying on the algorithm's recommendations as the primary or sole basis for medical decisions is an off-label use, when labeling calls for use subject to independent review by physicians. Off-label use of a medical device is legal. Section 1006 of the Food, Drug, and Cosmetic Act expressly limits FDA's power to interfere with off-label use of medical devices.[72] FDA cannot stop physicians from relying on CDS software as their primary basis for decision-making, even when the software was excluded from FDA oversight on the premise that it would not be so used. When medical devices are used off-label, liability for patient injuries generally falls on the physician rather than on the manufacturer. Clinicians face a daunting future, applying tools that may have received no regulatory review, unsure how the tools produce their recommendations, and liable for any injuries that result.

This prospect, in fact, may have motivated FDA to take the very conservative approach reflected in its CDS draft guidance. That guidance merely confirms that the Cures Act's regulatory exclusion applies to the simplest CDS systems and remains silent about how it applies to more complex systems. For now, FDA's plan may be to subject all CDS systems that incorporate machine-learning algorithms to regulatory scrutiny. This was not necessarily the balance of innovation and safety that the Cures Act aimed to strike. Moreover, it is not without practical challenges of its own: If an algorithm cannot explain its reasoning to a physician, how can it explain its reasoning to FDA so that FDA can regulate it?

Algorithmic transparency is not the only factor FDA needs to consider, when assessing whether software manufacturers have displayed the requisite intent to be excluded from FDA oversight. Data transparency and transparent business practices also matter. If clinical decisions incorporate real-world data, there has to be a way for physicians to satisfy themselves about the quality of those data. The CDS draft guidance addresses this point, stating that practitioners "would be unable to independently evaluate the basis of a recommendation if the recommendation were based on non-public information…".[73] Murky, black-box CDS algorithms that mine trade-secret-protected, proprietary data sets will not be eligible for the §360j(o)(1)(E) exclusion from FDA regulation. This provides a strong nudge for CDS system developers to be transparent, at least with physicians who use their products. Some authors have expressed concern that providing such access would require changes to the Health Insurance Portability and Accountability Act ("HIPAA") Privacy Rule.[74] This concern, fortunately, is unfounded: the Privacy Rule's treatment exception allows physician access to data for use in treating patients, and this includes accessing data from patients other than the person being treated.[75] There is no HIPAA impediment to letting physicians who use CDS software in a treatment setting inspect the HIPAA-protected datasets on which CDS decisions rely.

---

[72] See 21 U.S.C. § 396 (2012) (providing that FDA is not authorized to "limit or interfere with the authority of a health care practitioner to prescribe or administer any legally marketed device to a patient for any condition or disease within a legitimate health care practitioner-patient relationship").

[73] *Draft Guidance*, *supra* note 16, *at* 8.

[74] *See, e.g.*, Longhurst et al., *supra* note 34, at 1233.

[75] *See* Barbara J. Evans & Gail P. Jarvik, *Impact of HIPAA's Minimum Necessary Standard on Genomic Data Sharing*, GENETICS IN MED. (2017) (discussing HIPAA's broad exceptions to individual authorization and to the minimum necessary standard for data provided to healthcare providers for use in clinical care).

For EHR systems, there are ongoing concerns that vendors' contracts with user healthcare organizations often include terms that block free and open exchange of information about system performance, safety problems and adverse events, and advantages and disadvantages of particular systems.[76] EHR contract terms commonly include gag clauses, very broad confidentiality clauses, high switching costs that prevent dissatisfied users from commenting "with their feet" by replacing an unsatisfactory system, and dispute resolution provisions that block public disclosure of problems. FDA should take steps to ensure CDS software manufacturers renounce these sorts of nontransparent contracting practices, in order to qualify for exclusion from FDA oversight. Excluding CDS software from FDA regulation places it outside FDA's usual adverse-event reporting mechanisms, so physicians' ability to communicate frankly with one another about problems is crucial to patient safety.

VI.   THE CHALLENGE OF VALIDATING FUTURE CDS SOFTWARE

Any of us might be willing to take medical advice from an artificially intelligent CDS algorithm that, like the facial-scanning software discussed earlier,[77] is right 91% of the time when humans only achieves 60% accuracy. The problem, of course, lies in establishing what those percentages are—that is, in validating the software. FDA has considerable experience regulating machine-learning algorithms for certain significant- and moderate-risk medical applications, such as software that screens mammograms to highlight anomalies that are suspicious for breast cancer.[78] In this context, it is possible to validate the software by testing its performance using past images in which the presence or absence of cancer has been independently verified by other diagnostic methods. Knowing the "ground truth"[79] enables clinical study designs that compare the performance of two groups of radiologists, one group using the software while the other is working without it.[80]

For many future CDS applications, the ground truth may be unknowable in the premarket period, before the algorithm moves into wide clinical use. Consider CDS software that bases its diagnostic or treatment recommendations on deeply descriptive datasets that incorporate thousands of clinical, genomic, and exposure data points for each individual. For patients who present unique combinations of these data points, there is no ground truth that can be ascertained in advance: nobody like them was ever seen by a doctor before. As our understanding of precision medicine improves, it may become clear that that last sentence describes all of us. When true, the best way to validate CDS software may be through post-marketing observational studies: move it into

---

[76] OFFICE OF THE NAT'L COORDINATOR OF HEALTH INFORMATION TECH., EHR CONTRACTS UNTANGLED 12 (2016); OFFICE OF THE NAT'L COORDINATOR OF HEALTH INFORMATION TECH., REPORT TO CONGRESS ON HEALTH INFORMATION BLOCKING 31-32 (Apr. 2015); Darius Tahir, *Doctors Barred from Discussing Safety Glitches in US-funded Software,* POLITICO (Sept. 11, 2015), http://www.politico.com/story/2015/09/doctors-barred-from-discussing-safety-glitches-in-us-funded-software-213553 [https://perma.cc/F34A-Y7H6].

[77] *See* Kuang, *supra* note 60.

[78] Thompson, *supra* note 19.

[79] *See Ground Truth*, TECHOPEDIA, https://www.techopedia.com/definition/32514/ground-truth [https://perma.cc/8VZR-VBU6] (explaining that the term is borrowed from meteorology where it refers to the process of verifying remote sensing data by conducting on-site field tests).

[80] Thompson, *supra* note 19.

clinical use, let it start making predictions, and then follow the patients to see how the predictions turn out.

FDA's Digital Innovation Action Plan[81] and its Digital Health Software Precertification (Pre-Cert) Program,[82] both announced in 2017, acknowledge this problem:

> FDA's traditional approach to moderate and higher risk hardware-based medical devices is not well suited for the faster iterative design, development, and type of validation used for software-based medical technologies. Traditional implementation of the premarket requirements may impede or delay patient access to critical evolutions of software technology, particularly those presenting a lower risk to patients.[83]

FDA states that the agency is "reimagining its approach to digital health medical devices."[84] The Pre-Cert Program, already in its pilot phase, focuses regulatory scrutiny at the level of the firm that develops the software, rather than on the specific software product.[85] Does the firm "demonstrate a culture of quality and organizational excellence based on objective criteria, for example, that they can and do excel in software design, development, and validation (testing)"?[86] If so, FDA may allow a pre-certified firm to move its lower-risk software to market without any premarket review at all and may provide a more cursory or faster review of the firm's moderate- and higher-risk software.[87] Such firms could "collect real-world data post-market that might be used, for example, to affirm the regulatory status of the product, as well as to support new and evolving product functions."[88]

VII.   CONCLUSION

Among law practitioners and software manufacturers, there is already robust discussion of how the Cures Act may affect the pace of innovation and clinical translation of CDS systems, This discussion will continue as software manufacturers respond to FDA's CDS draft guidance, which is still in its public comment period as this article is written. Less analyzed are issues of potential importance to the physicians using CDS systems. For example, does Section 3060(a), in practical effect, absolve software manufacturers of liability for patient injuries arising during use of CDS systems and assign this liability to physicians? Will CDS software be subject to FDA regulation and validation processes? Will the clinically available CDS technologies be limited to basic

---

[81] *Digital Health Innovation Action Plan*, FOOD & DRUG ADMIN. (2017), https://www.fda.gov/downloads/MedicalDevices/DigitalHealth/UCM568735.pdf; *see also* Scott Gottlieb, *FDA Announces New Steps to Empower Consumers and Advance Digital Healthcare*, FOOD & DRUG ADMIN. (July 27, 2017), https://blogs.fda.gov/fdavoice/index.php/2017/07/fda-announces-new-steps-to-empower-consumers-and-advance-digital-healthcare/ [https://perma.cc/H239-HTP6].

[82] *Digital Health Software Precertification (Pre-Cert) Program*, FOOD & DRUG ADMIN. (last updated Feb. 15, 2018), https://www.fda.gov/MedicalDevices/DigitalHealth/UCM567265 [https://perma.cc/8CSJ-65AM]; *see also FDA Selects Participants for New Digital Health Software Precertification Pilot Program*, FOOD & DRUG ADMIN. (Sept. 26, 2017), https://www.fda.gov/NewsEvents/Newsroom/PressAnnouncements/ucm577480.htm [https://perma.cc/H2U7-LEK2].

[83] *Digital Health Innovation Action Plan*, *supra* note 81, at 2.

[84] *Id.* at 5.

[85] *See generally id.*

[86] *Id.*

[87] *Id.*

[88] *Id.*

systems that harness existing sources of medical knowledge, or will promising technologies that harness real-world evidence also move into clinical use without undue regulatory delays? The answers depend on how, and how well, FDA implements Section 3060. The crux of effective implementation is for FDA to enunciate standards of transparency CDS software must meet, before it will escape FDA regulation. Transparency in this context includes algorithmic transparency, physician access to underlying data that algorithms use, and CDS software vendor contracts that allow open, collegial airing of the strengths and weaknesses physicians encounter as they apply CDS in practice settings.