Special Review

How the FDA Regulates AI

H. Benjamin Harvey, MD, JD, Vrushab Gowda, BS

Recent years have seen digital technologies increasingly leveraged to multiply conventional imaging modalities' diagnostic power. Artificial intelligence (AI) is most prominent among these in the radiology space, touted as the "stethoscope of the 21st century" for its potential to revolutionize diagnostic precision, provider workflow, and healthcare expenditure. Partially owing to AI's unique characteristics, and partially due to its novelty, existing regulatory paradigms are not well suited to balancing patient safety with furthering the growth of this new sector. The current review examines the historic, current, and proposed regulatory treatment of AI-empowered medical devices by the US Food and Drug Administration (FDA). An innovative framework proposed by the FDA seeks to address these issues by looking to current good manufacturing practices (cGMP) and adopting a total product lifecycle (TPLC) approach. If brought into force, this may reduce the regulatory burden incumbent on developers, while holding them to rigorous quality standards, maximizing safety, and permitting the field to mature.

Key Words: Artificial intelligence; FDA; Medical Device; Regulation; Policy; Radiology.

INTRODUCTION

Artificial intelligence (AI) has found integration across a wide swathe of applications, running the gamut from digital advertising to self-driving cars and unmanned aerial vehicles. Healthcare is no exception. It offers substantial promise in addressing a number of critical issues – among others, expenditure growth, inefficient resource allocation, wait time, physician burnout, and missed diagnoses – so much so that it has been dubbed the "stethoscope of the 21st century" (1). Last April, the U.S. Food and Drug Administration (FDA or Agency) approved IDx-Dr, a software program assessing the progression of diabetic retinopathy based on fundoscopic images; this represents the first approval of a device authorized to provide a screening decision independent of physician confirmation (2). AI, however, is no panacea.

Like other novel technologies, it is not without its risks and unique implementation challenges. Indeed, in one of his final statements in office, former FDA Commissioner Scott Gottlieb acknowledged AI's transformative potential while underscoring the need to develop a wholly new regulatory paradigm (3). AI technologies naturally lend themselves to radiology, a data-driven field with extensive scope for pattern recognition, spatial modeling, algorithmic clinical scoring, and long-term surveillance (4). It comes as no surprise that more AI applications have been approved for use in radiology than for any other medical specialty and continue to proliferate (5). This piece looks beyond the hype to explore the regulatory hurdles particular to AI devices and critically evaluate their implications for radiologist practice.

EXISTING REGULATORY FRAMEWORKS: FROM PREMARKET REVIEW TO THE 21ST CENTURY CURES ACT

Rather than taxonomizing digital therapeutics based on technical characteristics, the FDA primarily assesses functionality with a view to risk-benefit balancing and intended clinical use. There is no separate review process for AI-based medical technologies as such at the time of writing. Depending on their function, they may alternatively receive approval as a "device" or exemption altogether from further FDA oversight.

Put briefly, the federal Food, Drug & Cosmetic Act defines device in relevant part as "an instrument...intended for use in the diagnosis...or in the cure, mitigation, treatment, or prevention of disease...which does not achieve its primary intended purposes through chemical action" (6). The FDA offers three avenues of premarket review for devices, each with differing data submission requirements based on the degree of patient risk exposure (7). New devices may require a De Novo request if they entail low to moderate risk, while higher-risk novel devices necessitate the more stringent Premarket Approval . Those representing modifications of existing devices may proceed through the abridged Premarket Notification process, commonly referred to as 510(k), pending a demonstration of "substantial equivalence" to a previously approved predicate.

Acad Radiol 2020; 27:58–61

From the Harvard Medical School, 175 Cambridge Street, Suite 200, Boston, MA 02114 (H.B.H.); Massachusetts General Hospital Department of Radiology, Boston, MA (H.B.H.); Massachusetts General Hospital Institute for Technology Assessment, Boston, MA (H.B.H.); Harvard Law School, Cambridge, Massachusetts (V.G.). Received September 7, 2019. accepted September 14, 2019. Financial Disclosure: Authors have no relevant financial disclosures. **Address correspondence to:** H. Benjamin Harvey, MD, JD, 175 Cambridge St., Suite 200, Boston, MA 02114. e-mail: harlanharvey@gmail.com

https://doi.org/10.1016/j.acra.2019.09.017

Apart from these conventional approaches, the FDA has outlined a novel pathway specifically for software as a medical device (SaMD) technologies in its Digital Health Innovation Action Plan, offering a streamlined route towards precertification (8). In the vein of good manufacturing practices, it looks to the characteristics of the developer, rather than those of the product itself (9). The Pre-Cert 1.0 pilot program is currently in its testing phase; participants were selected on the basis of "a robust culture of quality and organizational excellence, and who are committed to monitoring real-world performance" (10).

Passed in late 2016, the landmark 21st Century Cures Act (Cures Act) amends the Food, Drug & Cosmetic Act to include a new section, exempting a number of software functions from the statutory definition of device and hence from FDA oversight. Its §3060(a) expressly includes products (1) intended for administrative support, (2) those employed in the maintenance of a healthy lifestyle, (3) electronic patient records, (4) storage or display of clinical data, and of particular salience to radiology — (5) "unless the function is intended to acquire, process, or analyze a medical image. . .for the purpose of (i) displaying, analyzing, or printing medical information. . .(ii) supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment and...(iii) enabling such health care professional to independently review (emphasis added) the basis for such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient" (11).

This was followed by a pair of draft guidance documents, nonbinding policy statements issued by the FDA aimed at delineating the scope of its regulatory authority (12). The first, Clinical and Patient Decision Support Software (CDS Guidance), offers additional clarity regarding the aforementioned fifth prong of §3060(a) (13). It defines such products as CDS software, identifying "independent review" as the key operative provision. In order to bypass designation as a device and the approval process it entails, healthcare professionals must "be able to reach the same recommendation. . .without relying primarily on the software function." Specifically, the FDA holds CDS developers to explain the "purpose or intended use of the software function, the intended user (e.g., ultrasound technicians, vascular surgeons), the inputs used to generate the recommendation (e.g., patient age and gender), and the rationale or support for the recommendation." Of special relevance to AI products, it is unclear at this time how the FDA will interpret the latter provision in particular.

"Rationale or support" may at least partially depend on the familiarity of the healthcare provider with the software's data processing algorithm. Product developers, however, may be reluctant to share this information for fear of disclosing trade secrets in a highly competitive market. Those seeking to withhold proprietary details could find their products falling outside the scope of the §3060(a) exceptions, rendering them subject to comprehensive regulation as devices (14). In this sense, opacity proves a double-edged sword; it may confer a competitive edge while posing a number of regulatory obstacles that ultimately erect barriers to market entry. It remains to be determined how industry will respond to these countervailing forces. Moreover, in the case of products incorporating sophisticated machine learning or "black-box" algorithms, characterization of the underlying mechanism may prove technically infeasible. This would present a number of validation challenges that expose the limits of traditional clinical trial design, further complicating the approval process (15).

The second guidance document, Changes to Existing Medical Software Policies (Changes Guidance), seeks to risk-stratify products captured by the remaining four §3060(a) exceptions and concomitantly establishes principles of enforcement discretion. Namely, the Agency intends to prioritize regulation of tools displaying patient-specific medical device data and those flagging patient results based on specific clinical parameters, based on the extent to which they "alert a caregiver to take immediate clinical action" (16).

FUTURE REGULATORY DIRECTION: TOWARDS A CONSENSUS STANDARD

Owing to the novel concerns implicated by AI SaMD, the FDA has sought to develop an entirely new regulatory infrastructure to govern their approval, monitoring, and post-market modification. To this end, the Agency published Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning-Based Software as a Medical Device (Discussion Paper) in April 2019, outlining a new approach and soliciting public comment to inform future policymaking (17). It draws attention to the difficulty of regulating AI SaMD through conventional frameworks. For example, would a product incorporating a continuously learning AI algorithm require a new 510(k) application each time it "learned" a new function? When would these functions rise to the level of clinical significance? How would this be measured over time and at what point would a program cease to share "substantial equivalence" with its earlier, originally approved configuration?

The Discussion Paper offers a new paradigm for risk stratification incorporating three axes, to include the significance of information the program provides to the healthcare decision, the state of the healthcare condition, and the algorithm type — whether or not it is "locked" as opposed to "adaptive" and continuously learning in real-time. Flowing from this, modifications to AI SaMD may be classified into one or more of three broad categories necessitating premarket review: (1) a performance modification without changes to inputs or intended use (e.g. training the device with new data sets), (2) an input modification without change to intended use (e.g. incorporating new data types, enhancing interoperability), and (3) a modification to intended use (e.g. targeting a new patient population or indication).

It furthermore advocates for a total product lifecycle (TPLC) approach to AI SaMD, bearing echoes of the Pre-Cert program. The Discussion Paper's organization-centric focus looks to culture, process, and methodological rigor in its evaluation of "good machine learning practices,"

analogous to the current good manufacturing practicesthe FDA requires of traditional medical device producers, and explicitly calls upon manufacturers to develop algorithm change protocols in anticipation of future modifications. Furthermore, TPLC facilitates continuous monitoring of the SaMD product across the span of its development through its post-market performance, mandating the periodic disclosure of performance data to the FDA. In so doing, the Agency attempts to strike a balance between reducing the overall regulatory burden incumbent on AI SaMD developers while achieving its regulatory objectives.

IMPLICATIONS FOR RADIOLOGY: A BITE AT THE FEEDING HAND?

Radiology lies on the leading edge of SaMD innovation. Imaging departments have not only been among the first to adopt AI products in clinical use, but in many cases developed these technologies themselves (18). Their applications are manifold and span cardiac MR segmentations (Arterys) and neuroimaging analytics (icometrix) to cloud-based body CT critical feature flagging (Aidoc) and pneumothorax alerts (Zebra Medical Vision) (19–22). Reflective of their functional heterogeneity, AI-based radiology tools fall under virtually all of the regulatory classifications outlined previously.

The FDA has made its intention clear to regulate as devices software products used in the *interpretation* of diagnostic imaging. This is evident from the plain language of the Cures Act, and indeed the CDS Guidance expressly lists a number of radiological interpretive aids as examples of "devices [upon which] FDA intends to focus its regulatory oversight." However, this should not be taken to exclude AI products employed in workload reduction or logistical capacities within a radiology context. These would presumably bypass device designation per §3060(a) and find their way into the market without requiring FDA approval.

FDA has additionally sought to draw a distinction between computer-assisted detection devices (CADe) as opposed to the more stringently regulated computer-assisted diagnosis devices; the former are intended to "identify, mark, highlight, or in any other manner direct attention" to imaging features, rather than autonomously diagnose, stage, or triage pathology (23). Indeed, a number of imaging AI products offering CADe functions have received approval as CDS, including a CT-based stroke notification program (24). In a move to further ease the regulatory onus on SaMD developers, the FDA is considering reclassifying CADe used in the visualization of breast lesions, lung nodules, and dental caries to a lower-risk category, to require 510(k) submission rather than premarket approval (25).

The spread of AI technologies through the imaging space has understandably raised fears of radiologist employment security and compensation. While valid concerns, it is more appropriate to view these tools as adjuncts, rather than substitutes for an interpreting physician. Both existing frameworks as well as the proposed TPLC approach present CADe as the regulatory path of least resistance. SaMD developers will likely continue to market their products as aids to radiologists, who stand to realize efficiency gains and benefit from higher-value workflow changes (26,27). Moreover, the liability hazard associated with autonomous diagnostics ventures into legal *terra incognita* with attendant ill-defined legal risks, which developers may be reluctant to undertake absent further direction (28).

CONCLUSION

The integration of man and machine to drive outcomes is not a concept particularly new to radiology. However, AI poses unique regulatory issues which set it apart from other advances in imaging technology. Unlike the case for the majority of pharmaceutical products, devices, and foods, the FDA has indicated its preference to regulate AI software based on function, rather than technical components or indicated use. Consequently, medical products incorporating AI will likely straddle the boundaries delineated in decent guidance documents and find incorporation into both CDS and regulated devices necessitating conventional premarket review. As the market grows, the FDA will likely promulgate new regulations with a greater degree of specificity than those currently in existence, particularly in the realms of data security and privacy. This is critical to cloud-based systems susceptible to cyberattack (29,30). Further research in these and other areas will dynamically inform policymaking as the field matures.

REFERENCES

1. Artificial intelligence is the stethoscope of the 21st century. Available at: https://medicalfuturist.com/ibm-watson-is-the-stethoscope-of-the-21st-century. Accessed on July 25, 2019.
2. Abràmoff MD, Lavin PT, Birch M, et al. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. NPJ Digit Med 2018; 28:1–39. https://doi.org/10.1038/s41746-018-0040-6.
3. Statement from FDA Commissioner Scott Gottlieb, M.D. on steps toward a new, tailored review framework for artificial intelligence-based medical devices. Available at: https://www.fda.gov/news-events/press-announcements/statement-fda-commissioner-scott-gottlieb-md-steps-toward-new-tailored-review-framework-artificial. Accessed on July 25, 2019.
4. Kohli A, Mahajan V, Seals K, Kohli A, Jha S. Concepts in U.S. food and drug administration regulation of artificial intelligence for medical imaging. AJR Am J Roentgenol 2019; 5:1–3. https://doi.org/10.2214/AJR.18.20410.
5. FDA approvals for smart algorithms in medicine in one giant infographic. Available at: https://medicalfuturist.com/fda-approvals-for-algorithms-in-medicine. Accessed on July 25, 2019.
6. 21 U.S.C. §321(h).
7. Overview of medical device classification and reclassification. Available at: https://www.fda.gov/about-fda/cdrh-transparency/overview-medical-device-classification-and-reclassification. Accessed on July 25, 2019.
8. Software as a medical device (SAMD): clinical evaluation. Available at: https://www.fda.gov/media/100714/download. Accessed on July 25, 2019.
9. Digital health innovation action plan. Available at: https://www.fda.gov/media/106331/download. Accessed on July 25, 2019.

10. Digital health software precertification (pre-cert) program. Available at: https://www.fda.gov/medical-devices/digital-health/digital-health-software-precertification-pre-cert-program. Accessed on July 25, 2019.
11. 21 U.S.C. §360j(o).
12. Statement from FDA Commissioner Scott Gottlieb, M.D., on advancing new digital health policies to encourage innovation, bring efficiency and modernization to regulation. Available at: https://www.fda.gov/news-events/press-announcements/statement-fda-commissioner-scott-gottlieb-md-advancing-new-digital-health-policies-encourage. Accessed on July 25, 2019.
13. Clinical and patient decision support software: draft guidance for industry and food and drug administration staff. Available at: https://www.fda.gov/media/109618/download. Accessed on July 25, 2019.
14. Price WN. Big data and black-box medical algorithms. Sci Transl Med 2018; 12(471):10. https://doi.org/10.1126/scitranslmed.aao5333.
15. Price II WN. Black-box medicine. Harv J L & Tech 2015; 28(2):419.
16. Changes to existing medical software policies resulting from Section 3060 of the 21st Century Cures Act. Available at: https://www.fda.gov/media/109622/download. Accessed on July 25, 2019.
17. Proposed regulatory framework for modifications to artificial intelligence/machine learning (AI/ML)-based software as a medical device (SaMD) - discussion paper and request for feedback. Available at: https://www.fda.gov/media/122535/download. Accessed on July 25, 2019.
18. Topol EJ. High-performance medicine: the convergence of human and artificial intelligence. Nat Med 2019; 25:44–56. https://doi.org/10.1038/s41591-018-0300-7.
19. Arterys. Available at: https://www.arterys.com/cardio-ai/. Accessed on July 25, 2019.
20. Icometrix. Available at: https://icometrix.com/. Accessed on July 25, 2019.
21. Aidoc. Available at: https://www.aidoc.com/. Accessed on July 25, 2019.
22. Zebra Medical Vision. Available at: https://www.zebra-med.com. Accessed on July 25, 2019.
23. Computer-assisted detection devices applied to radiology images and radiology device data - premarket notification [510(k)] submissions. Available at: https://www.fda.gov/media/77635/download. Accessed on July 25, 2019.
24. FDA permits marketing of clinical decision support software for alerting providers of a potential stroke in patients. Available at: https://www.fda.gov/news-events/press-announcements/fda-permits-marketing-clinical-decision-support-software-alerting-providers-potential-stroke. Accessed on July 25, 2019.
25. 83 FR 25598, Radiology devices; reclassification of medical image analyzers. Available at: https://www.govinfo.gov/content/pkg/FR-2018-06-04/pdf/2018-11880.pdf. Accessed on July 25, 2019.
26. Recht M, Bryan RN. Artificial intelligence: threat or boon to radiologists? J Am Coll Radiol 2017; 14:1476–1480. https://doi.org/10.1016/j.jacr.2017.07.007.
27. Hosny A, Parmar C, Quackenbush J, et al. Artificial intelligence in radiology. Nat Rev Cancer 2018; 18:500–510. https://doi.org/10.1038/s41568-018-0016-5.
28. Sullivan HR, Schweikart SJ. Are current tort liability doctrines adequate for addressing injury caused by AI? AMA J Ethics 2019; 21:E160–E166. https://doi.org/10.1001/amajethics.2019.160.
29. Pesapane F, Volonté C, Codari M, et al. Artificial intelligence as a medical device in radiology: ethical and regulatory issues in Europe and the United States. Insights Imaging 2018; 9:745–753. https://doi.org/10.1007/s13244-018-0645-y.
30. Charles Q. Choi. Medical imaging AI software is vulnerable to covert attacks. Available at: https://spectrum.ieee.org/the-human-os/biomedical/imaging/medical-imaging-ai-software-vulnerable-to-covert-attacks. Accessed on July 25, 2019.