

FDA Regulation of Predictive Clinical Decision-Support Tools: What Does It Mean for Hospitals?

Gary E Weissman, MD, MSHP

Palliative and Advanced Illness Research (PAIR) Center, Perelman School of Medicine, University of Pennsylvania, Philadelphia, Pennsylvania; Leonard Davis Institute of Health Economics, University of Pennsylvania, Philadelphia, Pennsylvania.

Recent experiences in the transportation industry highlight the importance of getting right the regulation of decision-support systems in high-stakes environments. Two tragic plane crashes resulted in 346 deaths and were deemed, in part, to be related to a cockpit alert system that overwhelmed pilots with multiple notifications.¹ Similarly, a driverless car struck and killed a pedestrian in the street, in part because the car was not programmed to look for humans outside of a crosswalk.² These two bellwether events offer poignant lessons for healthcare, an industry in which human lives also depend on decision-support systems.

Clinical decision-support (CDS) systems are computerized applications, often embedded in an electronic health record (EHR), that provide information to clinicians to inform care. Although CDS systems have been used for many years,³ they have never been subjected to any enforcement of formal testing requirements. However, a draft guidance document released in 2019 from the US Food and Drug Administration (FDA) outlined new directions for the regulation of CDS systems.⁴ Although the FDA has thus far focused regulatory efforts on predictive systems developed by private manufacturers,^{5,6} this new document provides examples of software that would require regulation for CDS systems that hospitals are already using. Thus, this new guidance raises critical questions—will hospitals themselves be evaluated like private manufacturers, be exempted from federal regulation, or require their own specialized regulation? The FDA has not yet clarified its approach to hospitals or hospital-developed CDS systems, which leaves open numerous possibilities in a rapidly evolving regulatory environment.

Although the FDA has officially regulated CDS systems under section 201(h) of the Federal Food, Drug, and Cosmetic Act (1938), only recently has the FDA begun to sketch the shape of its regulatory efforts. This trend to actually regulate CDS systems began with the 21st Century Cures Act (2016) that amended the definition of software systems that qualify as medical devices and outlined criteria under which a system may be exempt from FDA oversight. For example, regulation would not apply to systems that support “population health”

or a “healthy lifestyle” or to ones that qualify as “electronic patient records” as long as they do not “interpret or analyze” data within them.⁷ Following the rapid proliferation of many machine learning and other predictive technologies with medical applications, the FDA began the voluntary Digital Health Software Precertification (Pre-Cert) Program in 2017. Through this program, the FDA selected nine companies from more than 100 applicants and certified them across five domains of excellence. Notably, the Pre-Cert Program currently allows for certification of software manufacturers themselves and does not approve or test actual software devices directly. This regulatory pathway will eventually allow manufacturers to apply under a modified premarket review process for individual software as a medical device (SaMD) that use artificial intelligence (AI) and machine learning. In the meantime, however, many hospitals have developed and deployed their own predictive CDS systems that cross the boundaries into the FDA’s purview and, indeed, do “interpret or analyze” data for real-time EHR alerts, population health management, and other applications.

Regulatory oversight for hospitals could provide quality or safety standards where currently there are none. However, such regulations could also interfere with existing local care practices, hinder rapid development of new CDS systems, and may be perceived as interfering in hospital operations. With the current enthusiasm for AI-based technologies and the concurrent lack of evidence to suggest their effectiveness in practice, regulation could also prompt necessary scrutiny of potential harms of CDS systems, an area with even less evidence. At the same time, CDS developers—private or hospital based—may be able to avoid regulation for some devices with well-placed disclaimers about the intended use of the CDS, one of the FDA criteria for determining the degree of oversight. If the FDA were to regulate hospitals or hospital-developed CDS systems, there are several unanswered questions to consider so that such regulations have their intended impact.

First, does the FDA intend to regulate hospitals and hospital-developed software at all? The framework for determining whether a CDS system will be regulated depends on the severity of the clinical scenario, the ability to independently evaluate the model output, and the intended user (Table). Notably, many types of CDS systems that would require regulation under this framework are already commonplace. For example, the FDA intends to regulate software that “identifies patients who may exhibit signs of opioid addiction,” a scenario similar to prediction models already developed at academic hospitals.⁸ The FDA also plans to regulate a software device

Corresponding Author: Gary E Weissman, MD, MSHP; Email: gary.weissman@pennmedicine.upenn.edu; Telephone: 215-746-2887; Twitter: @garyweissman.

Published online first August 19, 2020.

Received: January 16, 2020; **Revised:** April 9, 2020; **Accepted:** April 21, 2020

© 2021 Society of Hospital Medicine DOI 10.12788/jhm.3450

TABLE. FDA Framework for Determining Device Status and Regulatory Focus for CDS Systems*

Severity of clinical state	Possible to independently review CDS recommendations?	Intended user	
		Clinician	Patient or caregiver
Critical	Yes	Not a device, so no oversight	A device subject to oversight
	No	A device subject to oversight	A device subject to oversight
Serious	Yes	Not a device, so no oversight	A device subject to oversight
	No	A device subject to oversight	A device subject to oversight
Nonserious	Yes	Not a device, so no oversight	Qualifies as a device, but no plans for oversight
	No	Qualifies as a device, but no plans for oversight	A device subject to oversight

*Adapted from Table 3 of the FDA draft guidance “Clinical Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff.”⁴ The FDA intends to determine its oversight efforts for CDS systems based on the clinical severity of the situation in which the CDS system will be used, whether the end user can independently review the recommendations of the CDS, and whether the intended user is a clinician or a patient/caregiver. This framework applies only to CDS systems that are intended to “inform” clinical management rather than “drive” management or to “treat or diagnose.” These two latter categories are considered devices and would be subject to FDA oversight.

Abbreviations: CDS, clinical decision support; FDA, Food and Drug Administration.

even if it is not a CDS system if it is “intended to generate an alarm or an alert to notify a caregiver of a life-threatening condition, such as stroke, and the caregiver relies primarily on this alarm or alert to make a treatment decision.” Although there are no published reports of stroke-specific early warning systems in use, analogous nonspecific and sepsis-specific early warning systems to prompt urgent clinical care have been deployed by hospitals directly⁹ and developed for embedding in commercial EHRs.¹⁰ Hospitals need clarification on the FDA’s regulatory intentions for such CDS systems. FDA regulation of hospitals and hospital-developed CDS systems would fill a critical oversight need and potentially strengthen processes to improve safety and effectiveness. But burdensome regulations may also restrain hospitals from tackling complex problems in medicine for which they are uniquely suited.

Such a regulatory environment may be especially prohibitive for safety-net hospitals that could find themselves at a disadvantage in developing their own CDS systems relative to large academic medical centers that are typically endowed with greater resources. Additionally, CDS systems developed at academic medical centers may not generalize well to populations in the community setting, which could further deepen disparities in access to cutting-edge technologies. For example, racial bias in treatment and referral patterns could bias training labels for CDS systems focused on population health management.¹¹ Similarly, the composition of patient skin color in one population may distort predictions of a model in another with a different distribution of skin color, even when the primary outcome of a prediction model is gender.¹² Additional regulatory steps may apply for models that are adapted to new populations or recalibrated across locations and time.¹³ Until there are more data on the clinical impact of such CDS systems, it is unknown how potential differences in evaluation and approval would actually affect clinical outcomes.

Second, would hospitals be eligible for the Pre-Cert program, and if so, would they be held to the same standards as a private technology manufacturer? The domains of excellence required for precertification approval such as “patient

safety,” “clinical responsibility,” and “proactive culture” are aligned with the efforts of hospitals that are already overseen and accredited by organizations like the Joint Commission on Accreditation of Healthcare Organizations and the American Nurses Credentialing Center. There is limited motivation for the FDA to be in the business of regulating these aspects of hospital functions. However, while domains like “product quality” and “cybersecurity” may be less familiar to some hospitals, these existing credentialing bodies may be better suited than the FDA to set and enforce standards for hospitals. In contrast, private manufacturers may have deep expertise in these latter domains. Therefore, as with public-private partnerships for the development of predictive radiology applications,¹⁴ synergies between hospitals and manufacturers may also prove useful for obtaining approvals in a competitive marketplace. Simultaneously, such collaborations would continue to raise questions about conflicts of interest and data privacy.

Finally, regardless of how the FDA will regulate hospitals, what will become of predictive CDS systems that fall outside of the FDA’s scope? Hospitals will continue to find themselves in the position of self-regulation without clear guidance. Although the FDA suggests that developers of unregulated CDS systems still follow best practices for software validation and cybersecurity, existing guidance documents in these domains do not cover the full range of concerns relevant to the development, deployment, and oversight of AI-based CDS systems in the clinical domain. Nor do most hospitals have the infrastructure or expertise to oversee their own CDS systems. Disparate recommendations for development, training, and oversight of AI-based medical systems have emerged but have yet to be endorsed by a federal regulatory body or become part of the hospital accreditation process.¹⁵ Optimal local oversight would require a collaboration between clinical experts, hospital operations leaders, statisticians, data scientists, and ethics experts to ensure effectiveness, safety, and fairness.

Hospitals will remain at the forefront of developing and implementing predictive CDS systems. The proposed FDA regulatory framework would mark an important step toward realizing

benefit from such systems, but the FDA needs to clarify the requirements for hospitals and hospital-developed CDS systems to ensure reasonable standards that account for their differences from private software manufacturers. Should the FDA choose to focus regulation on private manufacturers only, hospitals leaders may both feel more empowered to develop their own local CDS tools and feel more comfortable buying CDS systems from vendors that have been precertified. This strategy would provide an optimal balance of assurance and flexibility while maintaining quality standards that ultimately improve patient care.

Disclosure: Dr Weissman holds a National Institutes of Health grant award (K23HL141639). He has no conflicts of interest to disclose.

References

- Sumwalt RL III, Landsbert B, Homendy J. *Assumptions Used in the Safety Assessment Process and the Effects of Multiple Alerts and Indications on Pilot Performance*. National Transportation Safety Board; 2019. <https://www.nts.gov/investigations/AccidentReports/Reports/ASR1901.pdf>
- Becic E, Zych N, Ivarsson J. *Vehicle Automation Report*. National Transportation Safety Board; 2019. <https://dms.nts.gov/public/62500-62999/62978/629713.pdf>
- Sutton RT, Pincock D, Baumgart DC, Sadowski DC, Fedorak RN, Kroeker KI. An overview of clinical decision support systems: benefits, risks, and strategies for success. *NPJ Digit Med*. 2020;3:17. <https://doi.org/10.1038/s41746-020-0221-y>
- Clinical Decision Support Software: Draft Guidance for Industry and Food and Drug Administration Staff. Food and Drug Administration. September 27, 2019. Accessed October 15, 2019. <https://www.fda.gov/media/109618/download>
- Gulshan V, Peng L, Coram M, et al. Development and validation of a deep learning algorithm for detection of diabetic retinopathy in retinal fundus photographs. *JAMA*. 2016;316(22):2402-2410. <https://doi.org/10.1001/jama.2016.17216>
- Abràmoff MD, Lavin PT, Birch M, Shah N, Folk JC. Pivotal trial of an autonomous AI-based diagnostic system for detection of diabetic retinopathy in primary care offices. *NPJ Digital Medicine*. 2018;1(1):39. <https://doi.org/10.1038/s41746-018-0040-6>
- Changes to Existing Medical Software Policies Resulting from Section 3060 of the 21st Century Cures Act: Guidance for Industry and Food and Drug Administration Staff. Food and Drug Administration. September 27, 2019. Accessed March 18, 2020. <https://www.fda.gov/media/109622/download>
- Lo-Ciganic W-H, Huang JL, Zhang HH, et al. Evaluation of machine-learning algorithms for predicting opioid overdose risk among Medicare beneficiaries with opioid prescriptions. *JAMA Netw Open*. 2019;2(3):e190968. <https://doi.org/10.1001/jamanetworkopen.2019.0968>
- Smith MEB, Chiovaro JC, O'Neil M, et al. Early warning system scores for clinical deterioration in hospitalized patients: a systematic review. *Ann Am Thorac Soc*. 2014;11(9):1454-1465. <https://doi.org/10.1513/annalsats.201403-102oc>
- WAVE Clinical Platform 510(k) Premarket Notification. Food and Drug Administration. January 4, 2018. Accessed March 3, 2020. <https://www.accessdata.fda.gov/scripts/cdrh/cfdocs/cfpmn/pmn.cfm?ID=K171056>
- Obermeyer Z, Powers B, Vogeli C, Mullainathan S. Dissecting racial bias in an algorithm used to manage the health of populations. *Science*. 2019;366(6464):447-453. <https://doi.org/10.1126/science.aax2342>
- Buolamwini J, Gebru T. Gender shades: intersectional accuracy disparities in commercial gender classification. *Proc Machine Learning Res*. 2018;81:1-15.
- Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD). Food and Drug Administration. April 2, 2019. Accessed April 6, 2020. <https://www.regulations.gov/contentStream?documentId=FDA-2019-N-1185-0001&attachmentNumber=1&contentType=pdf>
- Allen B. The role of the FDA in ensuring the safety and efficacy of artificial intelligence software and devices. *J Am Coll Radiol*. 2019;16(2):208-210. <https://doi.org/10.1016/j.jacr.2018.09.007>
- Reddy S, Allan S, Coghlan S, Cooper P. A governance model for the application of AI in health care. *J Am Med Inform Assoc*. 2019. <https://doi.org/10.1093/jamia/ocz192>