POLICY FORUM

What Should Oversight of Clinical Decision Support Systems Look Like?

Emily L. Evans, PhD, MPH and Danielle Whicher, PhD, MHS

Abstract

A learning health system provides opportunities to leverage data generated in the course of standard clinical care to improve clinical practice. One such opportunity includes a clinical decision support structure that would allow clinicians to query electronic health records (EHRs) such that responses from the EHRs could inform treatment recommendations. We argue that though using a clinical decision support system does not necessarily constitute a research activity subject to the Common Rule, it requires more ethical and regulatory oversight than activities of clinical practice are generally subjected to. In particular, we argue that the development and use of clinical decision support systems should be governed by a framework that (1) articulates appropriate conditions for their use, (2) includes processes for monitoring data quality and developing and validating algorithms, and (3) sufficiently protects patients' data.

Learning Health Systems and Patient-Centered Care

The increasing adoption of electronic health records (EHRs) and other technological advances allowing for routine collection of patient-generated data contributes to the infrastructure needed to transform health systems within the United States and abroad into learning health systems. A learning health system has been defined by the National Academy of Medicine (NAM) as one "in which science, informatics, incentives, and culture are aligned for continuous improvement and innovation, with best practices seamlessly embedded in the delivery process and new knowledge captured as an integral by-product of the delivery experience."[1]

The ability to leverage routinely collected data, both within and across health systems, holds promise for improving the organization and quality of care delivered to patients and for informing diagnostic, treatment, and other decisions based on patients' needs and individual characteristics. For example, some have argued that clinical decision support systems leveraging data aggregated from patients with similar clinical presentations could be designed to provide real-time, point-of-care feedback to help inform personalized treatment choices.[2] However, a number of scientific, ethical, and regulatory

questions remain regarding development and use of such clinical decision support systems for the purpose of making treatment recommendations.

In this paper, we focus on the question, What constitutes appropriate regulatory oversight of clinical decision support systems? We argue that while use of these systems does not necessarily constitute a research activity subject to the Common Rule,[3] development and implementation of these systems requires a greater level of ethical and regulatory oversight than is generally applied to activities of clinical practice or other health systems-level decisions about care delivery. In particular, ethical and regulatory oversight should ensure that (1) conditions for use of these systems (including adherence to evidence-based approaches) and the basis for the recommendations they generate are appropriately articulated, (2) systems rely on validated algorithms and address issues of data quality, and (3) sufficient privacy protections exist for patients whose data are used.

Applicability of Regulatory Oversight Requirements for Human Subjects Research to Clinical Decision Support Systems

Within the United States, the Common Rule provides the primary framework for ethical and regulatory oversight of federally funded biomedical and behavioral research involving human subjects.[3] The Common Rule defines types of research activities subject to regulations; these definitions draw heavily on the distinction between research and treatment articulated in the Belmont Report.[4] Under the Common Rule, research is defined as "a systematic investigation, including research development, testing, and evaluation, designed to develop or contribute to generalizable knowledge."[3] The tenability and utility of this distinction—and the resulting research and clinical oversight practices—are increasingly challenged by activities within learning health care systems, which cannot be classified exclusively as either research or practice.[5] However, much discussion around the need for a new ethical and regulatory framework has focused on reducing barriers to research activities rather than on strengthening oversight of clinical practice activities.[4] Amid debates about how to address "overprotection" of research participants, concerns persist about comparable lack of oversight for clinical practice and risks to which patients are exposed.[5,6]

A clinical decision support system that allows clinicians to query EHRs to inform individual point of care treatment recommendations (hereafter referred to as CDS-EHR) would generally not be considered research under the Common Rule's definition and would, therefore, not be subject to federal regulation. The objective of a CDS-EHR is not to produce generalizable knowledge but rather to provide a specific recommendation to a clinician and patient regarding appropriate treatment options. The use of this system has more in common with traditional "static" clinical prediction models (CPMs) and other decision tools currently used in clinical practice to inform treatment recommendations (eg, the Framingham Risk Score). A CDS-EHR simply applies previously developed

algorithms to existing data for the purpose of generating one-off estimates of potential risks and benefits of interventions under consideration by a clinician and a patient at the point of care. Development and validation of such a CPM might constitute research under the Common Rule, as might the evaluation of outcomes for patients whose treatment choices were informed by the use of these systems.

A CDS-EHR is better understood as what Faden et al. describe as a "learning activity," namely, an activity that involves integration of clinical care delivery with the objective of learning how to improve clinical practice or health care delivery.[7] Oversight of learning activities requires a framework that establishes conditions and governance for ensuring their sound scientific and ethical conduct.

Oversight Approaches for Clinical Decision Support Systems
To date, there have been several efforts to clarify appropriate oversight mechanisms for clinical decision support tools. For example, the 21st Century Cures Act identifies 4 conditions that must be met for clinical decision support software to *not* be defined as a device regulated by the Federal Food and Drug Administration (FDA)[8]; the FDA has also issued draft guidance to clarify the agency's interpretation of those 4 conditions.[9] Broadly, a CDS-EHR would be excluded from FDA regulation as a device if (1) it is intended for displaying, analyzing, or printing medical information, including information about a patient (eg, test results); (2) it is intended to support or provide recommendations to health care professionals about prevention, diagnoses, or treatment of medical conditions; (3) health care professionals are able to independently review the basis for such recommendations and do not rely primarily on the CDS-EHR in making treatment recommendations for an individual patient; and (4) it does not acquire, process, or analyze information from diagnostic devices.

Regardless of the status of CDS-EHRs with respect to current (or pending) regulatory requirements or proposed voluntary guidance,[10] we argue that transparent and responsible use of CDS-EHRs requires adherence to a set of baseline requirements.

First, CDS-EHRs are largely meant to aid in the decision-making process and should not be the sole source of information used to inform a clinical decision. Therefore, it is imperative that clinicians understand the basis of recommendations generated and the appropriate conditions for using the software, including that recommendations generated are not meant to replace existing guidelines. In particular, CDS-EHR software should be transparent about sources of patient-specific information and sources of clinical information or decision rules (eg, guidelines) used to generate recommendations.[9] When possible, the CDS-EHR software should also describe levels of certainty or reliability of recommendations and their clinical rationale.[10] Finally, clinicians using CDS-EHRs should have sufficient expertise to make the clinical decisions in question without the software and adequate time to consider generated recommendations in the context

of other clinically relevant information, including guidelines. Clinicians should also inform patients of how they arrived at their recommendations and how patients' data might be used in generating recommendations for other patients.[10]

Second, oversight is also needed to ensure that development, validation, and upkeep of a CDS-EHR adhere to best methodological practices. The need to adhere to best methodological practices, and not simply standard software validation practices, reflects pervasive concerns about how clinical prediction models are developed and validated.[11,12] There are several issues that should be addressed as part of this oversight. First, the quality of data used to develop a model must be carefully examined. For example, despite opportunities presented by access to increasing amounts of EHR data, such data can be incomplete, inaccurate, or otherwise unfit for use in research, including in developing CDS-EHRs.[12] Second, all methods used to develop models, no matter how sophisticated, have limitations; their potential impact on the validity and reliability of the models and of resulting treatment recommendations should be assessed. Third, models developed using data from one particular population might not produce valid and reliable recommendations when used in different patient populations without appropriate recalibration procedures.[12] Fourth, a CDS-EHR is not self-sufficient; once implemented, oversight is required to ensure its upkeep and evaluation.

Finally, insofar as a CDS-EHR actively queries an EHR system at the point of care, the recommendations generated can reveal protected health information (PHI) to clinicians, particularly in cases of rare diseases or small patient populations. (Even when recommendations are based on aggregate data, the relevant comparison group might be so small as to result in inadvertent disclosure of PHI.) Depending on the level of detail a clinician shares with a patient regarding a treatment recommendation's underlying rationale, such information could also be inadvertently revealed to a patient being treated. As with other activities in a learning health system, oversight of CDS-EHRs is needed to ensure sufficient privacy protections for patients whose data are used as part of the system.

Oversight of CDS-EHRs with respect to these baseline requirements requires identifying (or creating) the institution(s) with the appropriate independence, expertise, and enforcement capabilities, especially for systems not defined as devices subject to FDA regulations. A national independent body is needed to ensure appropriate use of these systems by clinicians and to establish and monitor adherence to standards for data quality and model validation. Whether an existing federal agency is adequately positioned to assume this role or whether a nongovernmental agency should be charged with this responsibility requires further deliberation. Local hospitals or health care systems, however, might be sufficient for ensuring patient privacy.

Improving Oversight to Ensure Patient-Centered Care

While a CDS-EHR is not research as defined by the Common Rule, we argue that such systems constitute learning activities that should be subject to appropriate oversight. Although CDS-EHRs hold great promise for informing patients' and clinicians' point-of-care decision making by leveraging large amounts of routinely generated data, they could result in increased risk of harm to patients (eg, inappropriate treatment recommendations, privacy breaches). Therefore, it is imperative that CDS-EHRs be developed using high-quality data and valid and reliable models. It is also imperative that clinicians are informed about appropriate use of CDS-EHRs, that they sufficiently understand the recommendations generated, and that the privacy of patients' data used by such systems is adequately protected.

Efforts to improve oversight of CDS-EHRs should also consider where oversight authority should be situated. Regardless of which agencies are charged with oversight, we believe that compliance with oversight policies or regulations should be required rather than voluntary, especially since individuals and organizations developing these systems can be influenced by incentives not always consistent with improving outcomes for patients. Systems developers, methodologists, clinicians, patients, and other health care stakeholders should be involved in efforts to inform development of appropriate policies and methodological standards for developing, validating, and maintaining CDS-EHRs to ensure high-quality, patient-centered care.

References

1. National Academy of Medicine. The learning health system series. https://nam.edu/programs/value-science-driven-health-care/learning-health-system-series/. Accessed March 21, 2018.
2. Longhurst CA, Harrington RA, Shah NHA. A "green button" for using aggregate patient data at the point of care. *Health Aff (Millwood)*. 2014;33(7):1229-1235.
3. Federal policy for the protection of human subjects; final rule. *Fed Regist*. 2017;82(12):7149-7274. https://www.gpo.gov/fdsys/pkg/FR-2017-01-19/pdf/2017-01058.pdf. Accessed April 6, 2018.
4. National Commission for the Protection of Human Subjects of Biomedical and Behavioral Research. The Belmont Report: ethical principles and guidelines for the protection of human subjects of research. https://www.hhs.gov/ohrp/regulations-and-policy/belmont-report/read-the-belmont-report/index.html. Published April 18, 1979. Accessed April 6, 2018.
5. Kass NE, Faden RR, Goodman SN, Pronovost P, Tunis S, Beauchamp TL. The research-treatment distinction: a problematic approach for determining which activities should have ethical oversight. *Hastings Cent Rep*. 2013;43(suppl 1):S4-S15.
6. Beauchamp TL, Saghai Y. The historical foundations of the research-practice distinction in bioethics. *Theor Med Bioeth*. 2012;33(1):45-56.

7. Faden RR, Kass NE, Goodman SN, Pronovost P, Tunis S, Beauchamp TL. An ethics framework for a learning health care system: a departure from traditional research ethics and clinical ethics. *Hastings Cent Rep*. 2013;43(suppl 1):S16-S27.
8. 21st Century Cures Act, HR 34, 114th Cong, 2nd Sess (2016).
9. US Food and Drug Administration. Clinical and patient decision support software: draft guidance for industry and Food and Drug Administration staff. https://www.fda.gov/downloads/medicaldevices/deviceregulationandguidance/guidancedocuments/ucm587819.pdf. Accessed March 18, 2018.
10. Clinical Decision Support Coalition. Introductory memorandum. http://cdscoalition.org/wp-content/uploads/2017/08/CDS-3060-Guidelines-Final-2.pdf. Published August 30, 2017. Accessed March 18, 2018.
11. Collins GS, Reitsma JB, Altman DG, Moons KG. Transparent reporting of a multivariable prediction model for individual prognosis or diagnosis (TRIPOD): the TRIPOD statement. *Ann Intern Med*. 2015;162(1):55-63.
12. Hersh WR, Weiner MG, Embi PJ, et al. Caveats for the use of operational electronic health record data in comparative effectiveness research. *Med Care*. 2013;51(8)(suppl 3):S30-S37.

Emily L. Evans, PhD, MPH is a program officer for the Clinical Effectiveness and Decision Science program at the Patient-Centered Outcomes Research Institute (PCORI) in Washington, DC. Prior to joining PCORI, she served as a consultant to the Institute of Medicine Committee on Ethical and Scientific Issues in Studying the Safety of Approved Drugs and as a member of research teams at the Johns Hopkins School of Medicine. She earned a BS in ethics, history, and public policy from Carnegie Mellon University, an MPH from Johns Hopkins Bloomberg School of Public Health, and a PhD in philosophy from Georgetown University.

Danielle Whicher, PhD, MHS is a senior program officer for the Leadership Consortium for a Value and Science-Driven Health System at the National Academy of Medicine in Washington, DC. She was previously a program officer for the Clinical Effectiveness and Decision Science program at the Patient-Centered Outcomes Research Institute, a project coordinator at the Johns Hopkins Berman Institute for Bioethics, and a guest lecturer at Johns Hopkins Bloomberg School of Public Health. She earned a BA in molecular biology from Colgate University as well as an MHS in health policy and management and a PhD in health policy and management and bioethics from the Johns Hopkins Bloomberg School of Public Health.