Federal Oversight Of AI In Health Care Is Taking Shape

By **Jeremy Sherer, Amy Joseph and Monica Massaro** (May 4, 2023)

Jeremy Sherer

Amy Joseph

Monica Massaro

On April 18, the Office of the National Coordinator of Healthcare Technology within the U.S. Department of Health and Human Services issued a significant proposed rule which begins to outline one means by which the federal government is looking to oversee the use of artificial intelligence and machine learning in health care.[1]

The Health Data, Technology and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing proposed rule, or HTI-1, has significant implications for many health care industry stakeholders, including developers of health care technology solutions, health care providers and provider organizations, and patients.

This article outlines some early observations for health care stakeholders.

Further Oversight Is Coming

Since the passage of the 21st Century Cures Act in 2016, federal health care agencies have worked to promote interoperability, information sharing, optimization of digital health tools and equity.

Focused efforts on AI have attempted to balance optimization while addressing potential risks including fostering trust and confidence in AI, protecting privacy, promoting equity and rooting out bias in the design of such technologies.[2] And as utilization of such tools has grown, there has been greater attention paid and calls for regulatory oversight.[3]

Previously collecting information through requests for information, HHS is now setting its sights on a formal timeline toward adoption of standards for utilization of AI in health settings. HTI-1 is just the beginning of this process.

At a high level, the proposed rule outlines a number of new and updated criteria which must be met for a health information technology developer to have health IT certified under the Office of the National Coordinator Health IT Certification Program, which is currently a voluntary program.

Several of those criteria would adopt significant new standards for AI and machine learning tools in health care, referred to as predictive decision support interventions, or DSIs.

The ONC describes its goal as being "to assist in addressing the gaps between the promise and peril of AI in health."

Despite ONC's insistence that it is not regulating this technology, the proposed rule appears to lay the groundwork for a type of indirect regulation the federal government has utilized in other areas of health care, whether by effectively setting a standard for the industry without establishing a mandate, or by blazing the path for further regulation to come.

At the very least, if the HTI-1 proposed rule is adopted as drafted, it will set new baseline

expectations for developers of health IT incorporating DSIs, although the ONC certification is voluntary. For example, it would not be surprising to see health care providers or others contracting with vendors for the use of such tools to contractually require that the tools are certified or at least meet the ONC certification requirements.

It is also possible that other state or federal agencies will look to the ONC certification when adopting their own standards, either by incorporating by cross reference or adopting similar requirements.

For example, Medicare remains the largest payor in the U.S. representing approximately 20% of patients.

As a result, for many provider stakeholders, being able to treat Medicare beneficiaries is necessary for survival, meaning they must, for example, satisfy the Centers for Medicare & Medicaid Services' conditions of participation, in the case of hospital facilities, or become users of certified electronic health record technology.

In the future, CMS could include standards for technology in its conditions of participation for hospitals as technology plays a larger role in health care, alongside other conditions of participation, such as those addressing medical staff and the physical environment of the facility.

Furthermore, the proposed rule expressly acknowledges the interests of many other federal agencies in the development, implementation and use of artificial intelligence in health care.

Earlier this month, the U.S. Food and Drug Administration published new guidance on marketing submissions for AI and machine learning.[4]

The Office for Civil Rights at HHS, which administers the Health Insurance Portability and Accountability Act; CMS; and the Federal Trade Commission are acknowledged as having significant interests in this area at different points throughout the proposed rule as well.

Even within the ONC itself, two other proposed rules were included in the Biden administration's unified agenda, one establishing disincentives to information blocking for health care providers and the other on information sharing and public health interoperability.

As the agency departments work in parallel on their various rulemakings, there is likely opportunity for pieces of this and future regulations on AI to be tied in to other administration proposals including annual payment rules, and other digital health rules such as those on emerging technologies or electronic prior authorization.

Suffice it to say, while the HTI-1 proposed rule is the most recent comprehensive rulemaking that will affect the utilization of AI in health care, it will not be the last.

Transparency and FAVES

The ONC articulates that its proposals

> are not aimed at approving or guaranteeing the quality of predictive DSIs or the models they are based on. Instead, [the] proposals are intended to provide users and the public greater information … on whether predictive DSIs are fair, appropriate, valid, effective, and safe.[5]

That objective is referenced throughout the proposed rule as assessing whether DSIs adhere to fair, appropriate, valid, effective and safe, or FAVES, principles.

To increase transparency and encourage adherence to FAVES principles, the ONC proposes to introduce a range of requirements that DSIs must satisfy to obtain ONC certification.

These broadly fall into three categories: (1) providing technical and performance information to users of DSIs; (2) requiring developers of DSIs to follow a range of risk management practices; and (3) requiring developers of DSIs to participate in real-world testing.

Providing Technical and Performance Information to Users of DSIs

The goal of the ONC's certification criteria regarding source attributes is to enable users to "make informed decisions about whether and how to use predictive DSIs."[6]

The proposed criteria require developers to make a plain-language description of source attribution information directly available to users. For health IT modules that interface with predictive DSIs, source attributes must be available regarding intervention details, intervention development, quantitative measures of intervention performance, and ongoing maintenance of the intervention, its implementation and use.

The ONC does not attempt to require specific measures or thresholds to demonstrate sufficient FAVES. Even without those metrics, though, the proposed criteria clearly signal that obtaining ONC certification will require enabling a product's users to meaningfully understand, through plain English explanations, how the technology was developed and how it functions.

Risk Management Practices

Beyond requirements for the DSIs themselves, the ONC's proposals also require developers of health IT modules to engage in specific types of risk management practices, and make information about those practices available to the public.

Specifically, risk management practices would include risk analysis, risk mitigation and governance. These criteria specifically focus on developers of certified health IT as organizations, and creating transparency into sociotechnical dimensions of the predictive DSI.

Meeting the FAVES standard will require developers to establish ongoing risk assessment programs to effectively harness the technology they build.

The risk-analysis criterion will require organizations developing certified health IT to analyze "risks related to the lack or failure of validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy."[7]

Risk mitigation efforts, meanwhile, are expected to include practices to rank risks based on their potential impact, minimize identified potential risks, change control plans, processes to intervene and harness DSIs that are malfunctioning, and seeking expert input when assessing performance of a DSI in a particular setting.[8]

The governance criterion envision establishing policies and controls for predictive DSIs. As ONC states in the proposed rule,

> We propose that a health IT developer of a certified Health IT Module that enables or interfaces with a predictive DSI must establish policies and implement controls for how data are acquired, managed, and used for said predictive DSI.[9]

The risk of bias associated with predictive DSIs is well documented and features prominently in the HTI-1 proposed rule. The risk management practices discussed herein are one of the main ways the federal government seeks to "promote equity in science and root out bias in the design and use of new technologies, such as artificial intelligence."[10]

For companies developing technology and intending to obtain ONC certification, adopting formal policies to ensure they meaningfully assess risk on an ongoing basis to combat bias and other threats will be crucial.

These provisions of the proposed rule are likely to receive significant comment, given the volume of analysis already released by digital health thought leaders, provider groups and trade associations about proper ways to eliminate bias.

Real-World Testing

The third category of transparency the ONC addresses in the proposed rule is achieved through real-world testing.

Real-world testing is already required under the ONC's existing clinical decision support criterion, which is being replaced by DSIs,[11] and the ONC is proposing to apply the same standards to DSIs.

Real-world testing increases transparency by requiring developers of certified health IT to submit "real world testing plans and corresponding real world testing results … demonstrating the real world use of each DSI the developer of certified health IT supports."[12]

Interactions With the Existing Health Care System

DSIs, and AI more broadly, appear poised to initiate real change in health care delivery in the U.S.

Whenever a new modality, program, or model emerges in the American health care system, it is critical to consider how the new development will interact with the intricacies and complexities of the existing system. A few of the ways DSIs interact with the existing system are outlined below.

Attestation Requirement

Under the proposed rule, health IT developers would need to attest "yes" or "no" as to whether their health IT module enables or interfaces with one or more predictive DSIs.

Providing an attestation may assist in providing comfort to users of the tools, but could also be a means to invite audits of such tools.

The ONC notes that there are numerous efforts being led by industry groups to develop methods to evaluate predictive DSIs, including algorithmic audits, which may be conducted by independent or adversarial parties, and ONC supports such monitoring — referred to as algorithmovigilance.[13]

For example, health systems or other health care providers seeking to contract with a health IT developer that has attested "yes" might require audit rights as part of the terms of its contract, similar to the request for audit rights often seen with respect to a vendor's use of a health care provider's data in a services agreement.

Regulators Already Pondering Potential Kickback Implications

The ONC acknowledges that AI in health care has significant potential value, including efficient allocation of resources, improved accuracy and reducing clinician burnout.[14]

That said, the ONC also addresses potential risk, including risk of violation of the federal Anti-Kickback Statute.

As addressed above, a significant risk with such tools generally is the risk of bias, where the model performs differently among certain patients, populations and communities, which could worsen disparities and access to care.

Though in most cases it seems likely that any such bias would not be intentionally included in the model, the ONC expressly raises a concern regarding a developer deliberately introducing bias given their financial interest in the result.

The ONC notes that where a third party provides remuneration to a health IT developer "for integrating or enabling DSI where one purpose is to increase sales of the third-party's products or services," the AKS is potentially implicated.

Pharmaceutical manufacturers and clinical laboratories are specifically called out as entities that could have a financial interest in the outputs of such tools, which may inform a health care provider's ultimate decision as to ordering items or services for a patient.

Time will tell how significant of a risk this is in practice, but one can imagine a new trend of health care fraud cases — with or without merit — focused on circumstances where a health IT developer licenses their software with DSI for a fee to a third party that stands to gain financially from provision of items or services reimbursable by a federal health care program that were recommended by the tool.

Though there is nothing inherently problematic with the use of DSIs under the AKS, it could be a new means by which the same old fraud schemes are perpetuated.

Conclusion

Plenty remains to be seen about how AI will be regulated in the health care space. The ONC's HTI-1 proposed rule offers an early overview of the health care technology regulator's objectives, including increased transparency and promoting equity, as well as its concerns around bias and other issues.

Jeremy Sherer is a partner and co-chair of the digital health practice at Hooper Lundy & Bookman PC.

Amy Joseph is a partner and co-chair of the academic medical center working group at the firm.

Monica Massaro is director of government relations and public policy at the firm.

[1] Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing, 88 Fed. Reg. 23746 (Apr. 18, 2023).

[2] 88 Fed. Reg. 23775 (May 23, 2018).

[3] https://www.democrats.senate.gov/newsroom/press-releases/schumer-launches-major-effort-to-get-ahead-of-artificial-intelligence.

[4] Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning (AI/ML)-Enabled Device Software Functions -Draft Guidance for Industry and Food and Drug Administration Staff (April 3, 2023).

[5] 88 Fed. Reg. 23780.

[6] 88 Fed. Reg. 23780.

[7] 88 Fed. Reg. 23799.

[8] 88 Fed. Reg. 23802.

[9] 88 Fed. Reg. 23802.

[10] See Executive Order 14091 (Feb. 16, 2023).

[11] ONC explains that it is revising the name of the CDS criterion to DSIs "to reflect the various and expanding forms of decision support that Health IT Modules enable or interface with. Increasingly, DSIs include use cases or are intended to support decision-making across all areas of healthcare, including early detection of disease, automating billing procedures, facilitating scheduling, supporting publish health disease surveillance, and other uses beyond traditional CDS." 88 Fed. Reg. at 23781.

[12] 88 Fed. Reg. 23783.

[13] 88 Fed. Reg. 23780, 23794-95.

[14] 88 Fed. Reg. 23776.