



Office of the National Coordinator
for Health Information Technology

Health Data, Technology, and Interoperability: Decision Support Intervention

HTI-1 Proposed Rule Subgroup 2 – DSI Full Deck

Presented by Kathryn Marchesini, Jordan Everson, and Jeffery Smith



Disclaimer and Public Comment Guidance

- The materials contained in this presentation are based on the proposals in the "Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing" proposed rule. While every effort has been made to ensure the accuracy of this restatement of those proposals, this presentation is not a legal document. The official proposals are contained in the proposed rule.
- ONC must protect the rulemaking process and comply with the Administrative Procedure Act. During the rulemaking process, ONC can only present the information that is in the proposed rule as it is contained in the proposed rule. ONC cannot interpret that information, nor clarify or provide any further guidance.
- ONC cannot address any comments made by anyone attending the presentation or consider any such comments in the rulemaking process, unless submitted through the formal comment submission process as specified in the Federal Register.
- This communication is produced and disseminated at U.S. taxpayer expense



Agenda

1. HITAC TF Meeting #1 – April 25, 2023
 - Policy Overview
 - Context and Background
 - Proposed Revisions and Criterion Mechanics
2. HITAC TF Meeting #2 – April 26, 2023
 - Source Attributes
3. HITAC TF Meeting #3 – May 3, 2023
 - Intervention Risk Management
 - Oversight & Implementation






Policy Overview

Decision Support Interventions (DSI) Proposals

ONC proposes to revise the existing CDS criterion § 170.315(a)(9) to reflect an array of contemporary and emerging software functionalities that aid user decision-making in health care, including artificial intelligence (AI) and machine learning (ML).

This revision includes:

- A definition for “predictive decision support intervention”
- Updating the Base EHR definition to include the proposed revised DSI criterion in § 170.315(b)(11)
- Requirements for Health IT Modules that enable or interface with predictive DSIs to provide relevant technical and performance information to users
- Requirements for certified health IT developers to employ or engage in risk management practices related to predictive DSIs
- Requirements for certified health IT developers with Health IT Modules certified to DSI criterion to participate in Real World Testing



DSI Proposals – Benefits

ONC proposes these revisions to optimize the use of predictive and other DSIs types in health care. These baseline requirements for transparency aim to improve the trustworthiness of predictive algorithms and support their widespread use in health care.

Other intended outcomes include:

- **Improving transparency** regarding how a predictive DSI is designed, developed, trained, evaluated, and should be used
- **Enhancing trustworthiness** through transparency on how certified health IT developers manage potential risks and govern predictive DSIs that their certified Health IT Modules enable or interface with
- **Supporting consistency** in the availability of predictive DSI information to users, so that users may determine the DSI's quality and whether its recommendations are fair, appropriate, valid, effective, and safe (FAVES)
- **Advancing Health Equity by Design** by addressing bias and health disparities potentially propagated by predictive DSIs to expand the use of these technologies in safer, more appropriate, and more equitable ways



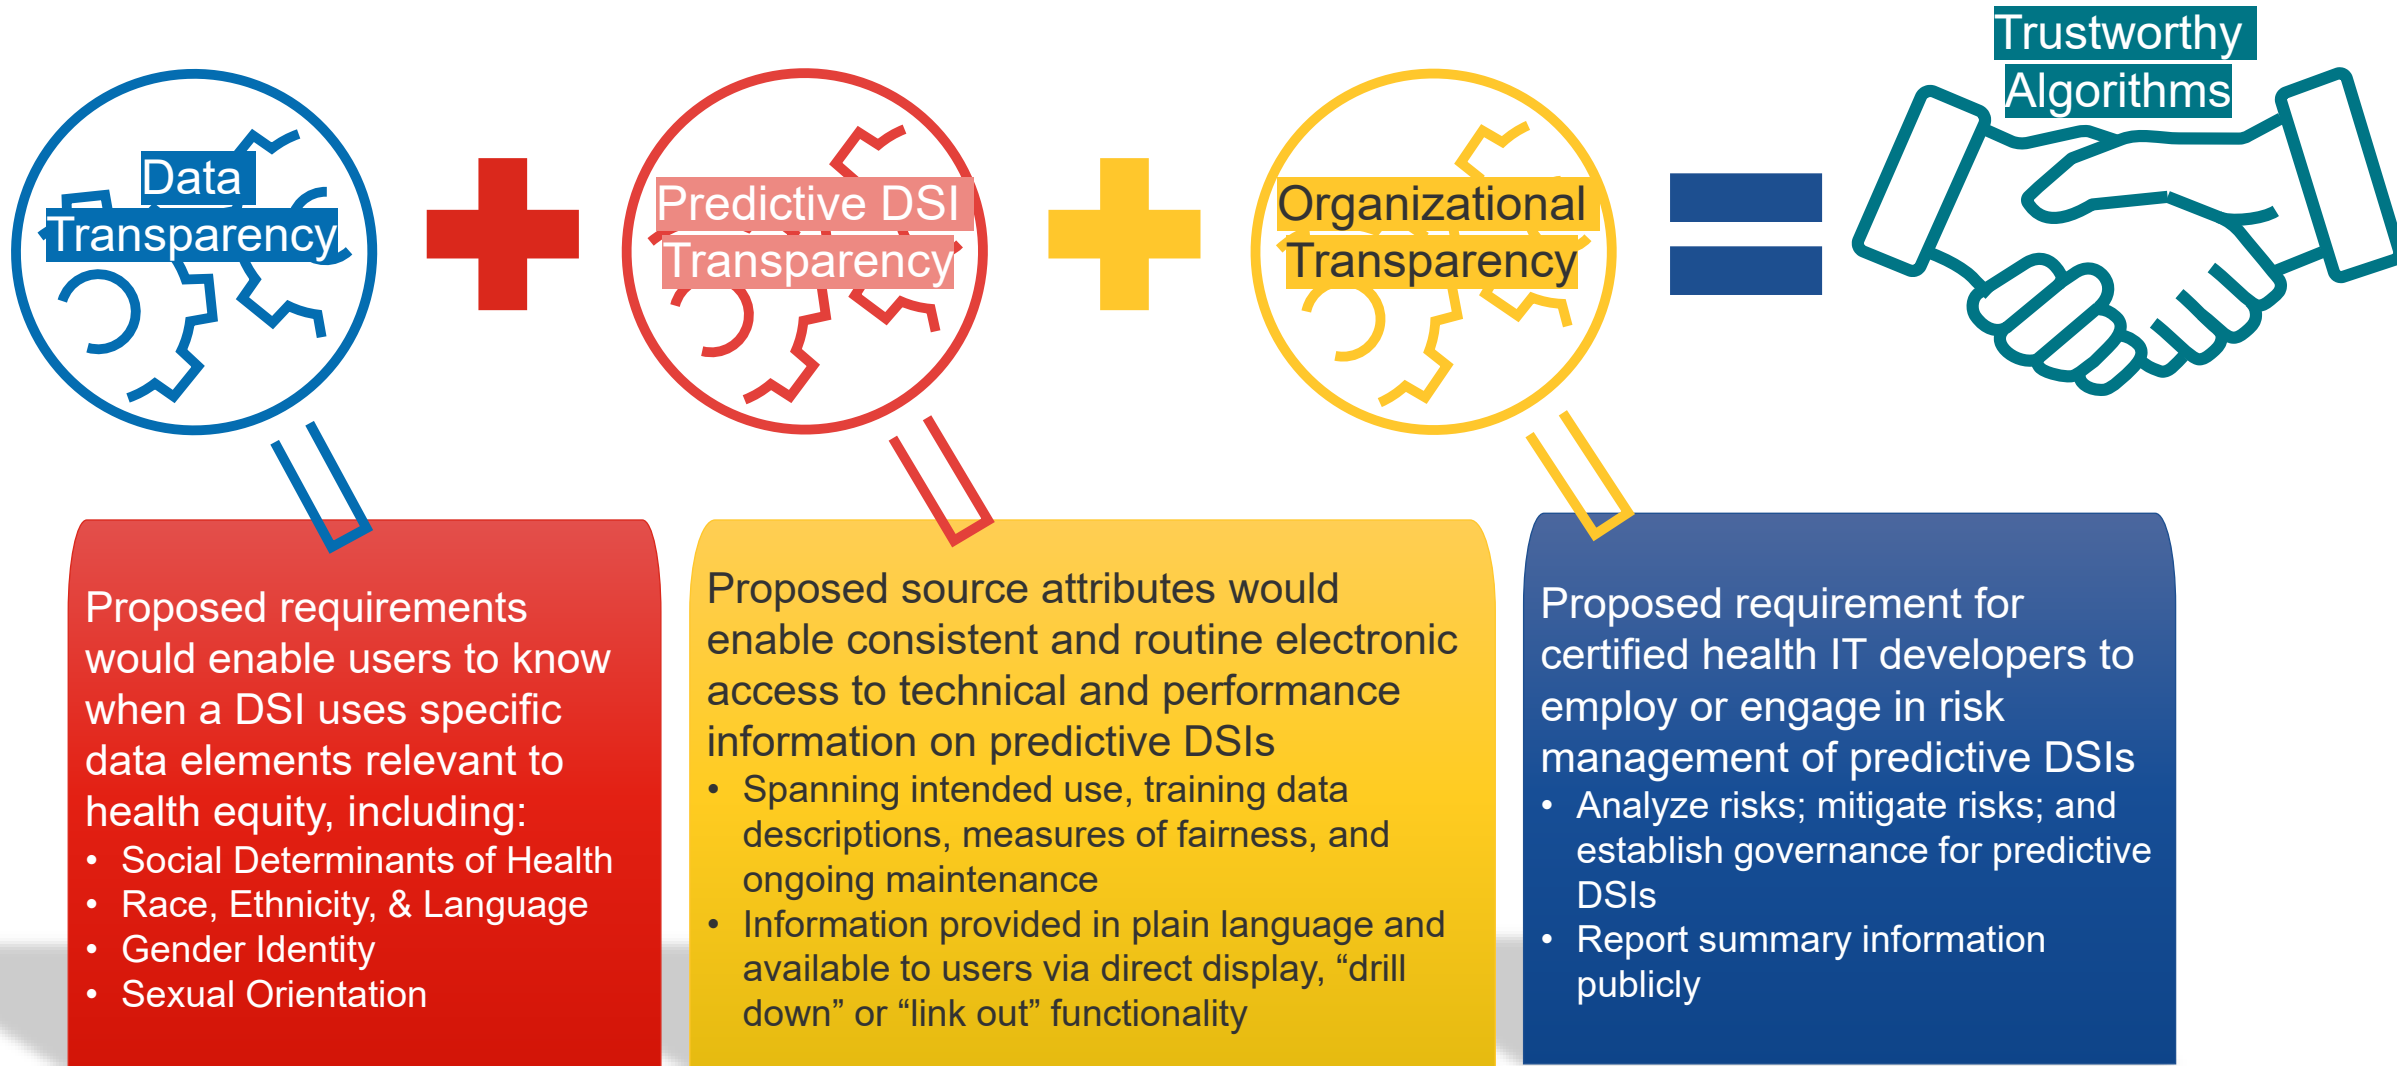
Proposed Definition: “Predictive Decision Support Intervention”

Predictive Decision Support Intervention Means:

“Technology intended to support decision-making based on algorithms or models that derive relationships from training or example data and then are used to produce an output or outputs related to, but not limited to, prediction, classification, recommendation, evaluation, or analysis.”

- Technology estimates a value based on relationships ‘learned’ in training data
- Agnostic to specific purposes or intended uses
- Does not convey or consider a level of risk associated with its use
- Not dependent on who developed the algorithm or model (can be someone other than a certified health IT developer)
- Not limited by the specific nature of the data to be processed (includes models that analyze text or images are included)
- Examples include:
 - Simple statistics and regression models used in a risk calculator (e.g., such as the widely used ASCVD model, which predicts heart events, and APACHE IV model, which predicts in-hospital death for ICU patients)
 - Machine learning models of various complexity, including neural networks and gradient boosted machines (used, for example, to predict hospital readmission, sepsis onset, and patient no-shows) and large language models including generative pre-trained transformers (e.g., ChatGPT)
- Outputs of predictive model may be presented in a broad array of forms that DSIs can take (e.g., alerts, order sets, flowsheets, etc.)

Transparency Is A Prerequisite For Trustworthy AI




F.A.V.E.S. as an Intended Policy Outcome

- We posit that transparency regarding (1) the technical and performance dimensions of predictive DSIs and (2) the organizational and socio-technical competencies employed by those who develop predictive DSIs is foundational for users to determine whether their predictive DSI is fair, appropriate, valid, effective, and safe, or FAVES
- Proposals for source attribute and intervention risk management information would provide essential information to users determining whether and how to apply a predictive output to medical decisions at the time and place of care
- Numerous and parallel efforts across industry, academia, and the public sector are developing means to communicate measures of FAVES including through
- Model cards, Model nutrition labels, Datasheets, Data cards, Algorithmic audits, impact assessments, etc.
 - Proposals would provide a foundation for these efforts meant to shed light on the quality of predictive models in health care.
- While not guaranteeing whether a predictive DSI is FAVES, these proposals would promote transparency necessary for a dynamic marketplace of high-quality predictive models to support decision-making in health care
 - We believe such transparency would also foster confidence and trust among interested parties that the technical and organization processes used in designing and developing the predictive DSI were FAVES and high-quality.
 - We believe that transparency can increase public trust and confidence in technology.



Background & Context for Proposals

Current View of Artificial Intelligence in Health Care



ONC Blog Series on Artificial Intelligence and Machine Learning in Health Care



“ ONC’s policy goal has been and continues to be centered on ensuring that certified health IT can support broad categories of decision support intervention types, while being agnostic as to the intended purpose of such decision support. This approach has led to a dynamic and flourishing landscape of decision support technologies, varied in purpose and scope, ranging from patient safety and clinical management to administrative and documentation functions.

“ These technologies have the potential to drive innovation, increase market competition, and vastly improve care for patients and populations. However, like any new health IT, it requires examination and inquiry to establish an evidence-base of benefits and risks.

“ While predictive decision support interventions (DSIs) have enormous potential to improve many aspects of health care, they also present several potential risks that could lead to adverse impacts or outcomes. These risks may be magnified because of their potential to “learn” rapidly and produce predictions across many hundreds or thousands of patients.

Common Terminology Around Key Concepts

Term	Description
Predictive decision support (<i>Model</i>)	Technology intended to support decision-making based on algorithms that derive relationships from training or example data and then are used to produce an output or outputs.
Transparency	Sufficient information provided on the model, including input data, validation of performance, and intended use.
Trustworthiness	Model risks identified, mitigated, managed, and evaluated to provide confidence in the positive impact of using the model, and information about steps taken to govern the model and address negative impacts and/or reduce bias or harm are documented.
Fair (<i>Unbiased, Equitable</i>)	Model does not exhibit prejudice or favoritism toward an individual or group based on their inherent or acquired characteristics. The impact of using the model is similar across same or different populations or groups.
Appropriate	Model is well matched to specific contexts and populations to which it is applied.
Valid	Model has been shown to estimate targeted values accurately and as expected in both internal and external data.
Effective	Model has demonstrated benefit in real-world conditions.
Safe	Model is free from any unacceptable risks and for which the probable benefits outweigh any probable risk.

What's Hindering the Use of AI/ML In Health Care?

During a June 2022 Health Information Technology Advisory Committee hearing on “health equity by design,” **we heard that clinicians have unmet needs for information and transparency, and that until these needs are met, they are unlikely to use ML-driven tools or risk misapplying them to their patients.**

- Clinicians need to know that an AI product has been evaluated in their setting of care, that the technology was trained on data that reflects their practice population, and that the product will be continuously monitored.
- Clinicians want to be able to communicate back to developers of such AI products when a predictive recommendation did not work well for a patient.
- General concern that ML-driven technology does not create or recreate systemic inequalities that come with the lack of access to quality health insurance and quality care.

HealthTbuzz 
THE LATEST ON HEALTH IT FROM ONC

Minimizing Risks and Maximizing Rewards from Machine Learning

Kathryn Marchesini; Jeff Smith and Jordan Everson | SEPTEMBER 7, 2022

Blog 4: Information Asymmetry & Predictive Models

“ We see [a] lack of consistent information availability (or information asymmetry) as a defining challenge inhibiting the optimization of predictive decision support interventions (DSIs) in health care. For students of economics, this type of insufficient information or “quality uncertainty” is one of the most famous forms of market failures, often colloquially called a “market for lemons”—as in the ancient slang for malfunctioning used cars.

HealthTbuzz 
THE LATEST ON HEALTH IT FROM ONC

Back to the Future: What Predictive Decision Support Can Learn from DeLoreans and The Big Short

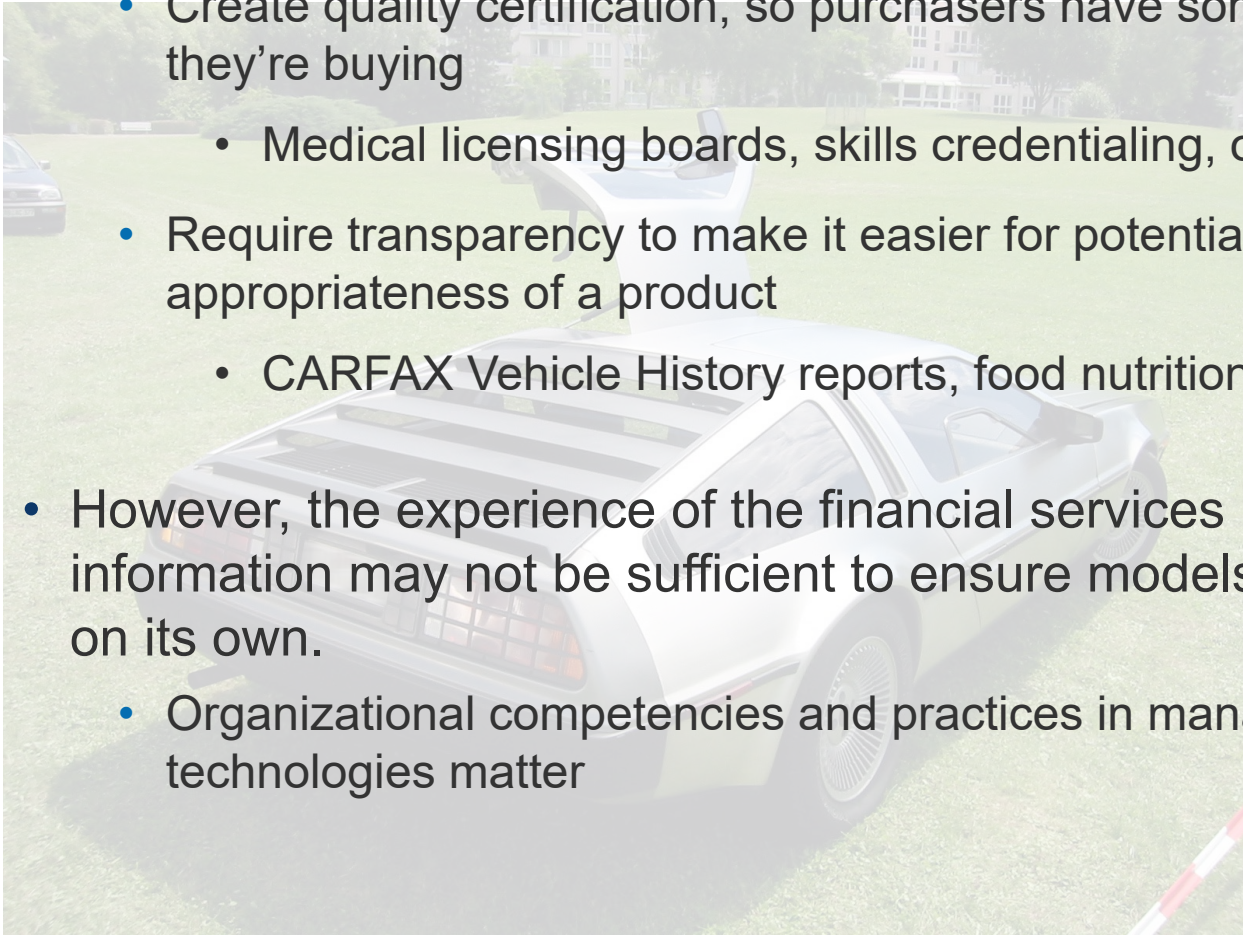


Signs that you have a market for lemons

1. Purchasers or Users Complain About Real Lemons
2. Information Asymmetry Abounds, Leading to An Inability to Determine the “Good” from the “Bad”
3. Producers or Suppliers of Quality Products Exit the Market

DeLoreans and The Big Short

- There are two classic things that can be done about a “market for lemons”
 - Create quality certification, so purchasers have some trust in the underlying quality of what they’re buying
 - Medical licensing boards, skills credentialing, organization accreditation, product approvals
 - Require transparency to make it easier for potential users to ascertain the quality or appropriateness of a product
 - CARFAX Vehicle History reports, food nutrition labels, drug facts labels
- However, the experience of the financial services industry highlights that validation information may not be sufficient to ensure models are high quality and used appropriately on its own.
 - Organizational competencies and practices in managing risk for models and AI/ML-related technologies matter





Decision Support And Certified Health IT

- Since 2010, the Program has maintained a CDS certification criterion, consistent with the “qualified electronic health record” definition in section 3000(13) of the PHSa,
 - An electronic record of health-related information on an individual that has the capacity to “provide clinical decision support” (42 U.S.C. § 300jj(13)(B)(i)).
- The initial CDS criterion required that a Health IT Module could:
 - Implement rules, “according to specialty or clinical priorities;”
 - “Automatically and electronically generate and indicate in real-time, alerts and care suggestions based upon clinical decision support rules and evidence grade;” and
 - Track, record, and generate reports on the number of alerts responded to by a user (75 FR 2046)
- HITPC recommendations in 2012 provided the framework for our current CDS criterion, including requirements that Health IT Modules support CDS that :
 - Displays source or citation of CDS
 - Is configurable based on patient context (e.g., inpatient, outpatient, problems, meds, allergies, lab results)
 - Is presented at a relevant point in clinical workflow
 - Include alerts presented to users who can act on alerts (e.g., licensed professionals);
 - Is integrated with the EHR (i.e., not standalone)



The Landscape for CDS has evolved since 2012

- Predictive models are increasingly being used and relied upon to inform an array of decision-makers, including clinicians, payers, researchers, and individuals
- Certified health IT is a central component and data source of these predictive models
 - Power the training and real-world use of predictive models as testing data or real-time inputs into deployed predictive models
 - Create and deploy predictive algorithms or models for use in production environments through their Health IT Modules
 - Enable other parties, including third-party developers and customers of the developer of certified health IT, to create and deploy predictive models through the developer's Health IT Modules
 - Are often the vehicle or delivery mechanism for predictive model outputs to reach users, such as clinicians, through decision support



Proposed Revisions

CDS versus DSI

- Clinical Decision Support (CDS)
 - 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
- Decision Support Intervention (DSI)
 - DSI expand on the concept of CDS to reflect the various and expanding forms of decision support that certified Health IT Modules enable or interface with
 - Increasingly, DSIs include use cases or are intended to support decision-making across all areas of health care, not just clinical workflow, including early detection of disease, automating billing procedures, facilitating scheduling, supporting public health disease surveillance, and other uses beyond traditional CDS
 - We intend for the DSI criterion to be inclusive of the wide variety of use cases that Health IT Modules may support moving forward
 - DSIs covered by our proposed requirements include existing DSI types, evidence-based and linked referential, and the proposed predictive DSI type





§ 170.315(a)(9) to § 170.315(b)(11)

- Much of the proposed structure and requirements are duplicated across the CDS § 170.315(a)(9) criterion for the proposed DSI § 170.315(b)(11) criterion and reflect the capabilities included in the CDS criterion
 - Health IT Modules must enable Evidence-based DSIs and Linked referential DSIs based on a defined set of data elements
 - Problems, medications, allergies and intolerances, demographics, laboratory, vital signs
 - **NEW**: Procedures and Unique Identifier(s)
 - Health IT Modules are required to enable a user to review “source attributes” information
 - Bibliographic citation of the intervention
 - Developer of the intervention
 - Funding source of the intervention
 - Release, and if applicable, revision date(s) of the intervention
 - **NEW**: Use in the intervention of specific demographic data
 - **NEW**: Use of social determinants of health data
 - **NEW**: Use of health status/assessment data



Decision Support And Certified Health IT

- Since 2010, the Program has maintained a CDS certification criterion, consistent with the “qualified electronic health record” definition in section 3000(13) of the PHSa,
 - An electronic record of health-related information on an individual that has the capacity to “provide clinical decision support” (42 U.S.C. § 300jj(13)(B)(i)).
- The initial CDS criterion required that a Health IT Module could:
 - Implement rules, “according to specialty or clinical priorities;”
 - “Automatically and electronically generate and indicate in real-time, alerts and care suggestions based upon clinical decision support rules and evidence grade;” and
 - Track, record, and generate reports on the number of alerts responded to by a user (75 FR 2046)
- HITPC recommendations in 2012 provided the framework for our current CDS criterion, including requirements that Health IT Modules support CDS that :
 - Displays source or citation of CDS
 - Is configurable based on patient context (e.g., inpatient, outpatient, problems, meds, allergies, lab results)
 - Is presented at a relevant point in clinical workflow
 - Include alerts presented to users who can act on alerts (e.g., licensed professionals);
 - Is integrated with the EHR (i.e., not standalone)



Scope of Certification and Decision Support Criterion

- The decision support intervention does not get certified, the Health IT Module supporting decision support does
- Current CDS criterion for Health IT Modules is part of the “Base EHR” definition
 - The “Base EHR” is referenced in CMS payment policy
 - We propose to update the “Base EHR” definition to include the new DSI criterion
- Current requirements are for Health IT Modules to:
 - Enable interventions based on (1) specific data elements and (2) when meds, allergies, and problems are incorporated from a transition of care/referral summary received
 - Enable “evidence-based decision support interventions” based on a set of data elements
 - Identify for a user diagnostic or therapeutic reference information based on set of data elements
 - Enable a user to review “source attributes”
 - Bibliographic citation, developer details, funding source, release/revision information



Proposed New Requirements for All Health IT Modules Certified to the DSI Criterion

- Source Attributes must be available as a “plain language description” to users “via direct display, drill down, or link out from a Health IT Module”
 - This would make a historic expectation explicitly required
- If DSI is developed by a developer of certified health IT, all attributes are required, unless otherwise noted as “if available”
- For DSIs that are developed by other parties, clearly indicate when any attribute is not available for the user to review
 - Other parties include health systems, third-party software developers, medical education publishers, etc.
- Health IT Modules must enable users to “author and revise source attributes and information” beyond source attributes listed
 - This would provide flexibility for users to design DSI information unique to their circumstances
- Enable end users to provide feedback regarding the intervention and make available such feedback data for export, in a computable format, including the intervention, action taken, user feedback provided (if applicable), user, date, and location
 - This would support quality improvement for all DSIs

Predictive DSI Definition and Related Request for Comment

Predictive Decision Support Intervention Means:

“Technology intended to support decision-making based on algorithms or models that derive relationships from training or example data and then are used to produce an output or outputs related to, but not limited to, prediction, classification, recommendation, evaluation, or analysis.”

- Request for comment:
 - Predictive DSI definition would not include
 - Simulation models that use modeler-provided parameters rather than training data
 - Unsupervised machine learning techniques that do not predict an unknown value among other technologies.
 - Are there prominent models (e.g., simulation models, unsupervised learning models) used to support decision-making in healthcare that are not effectively captured under the proposed definition of a predictive DSI?
 - If so, is it is feasible and appropriate to include such models in the scope of this proposed rule?





Proposed Predictive DSI Attestation

- Health IT Modules certified to § 170.315(b)(11) **are not required to enable or interface** with predictive DSIs, but developers of certified health IT must make one of the following attestations:
 - Yes – the Health IT Module enables or interfaces with a predictive decision support intervention(s) based on any of the data expressed in the USCDI
 - No – the Health IT Module does not enable or interface with a predictive decision support intervention(s) based on any of the data expressed in the USCDI
- If the developer attests “yes,” to this statement, the developer and its certified Health IT Module are subject to applicable predictive DSI requirements
- If the developer attests “no” to this statement, the developer would be subject to applicable general DSI requirements

Proposed Scope of Covered Technologies

Developers of certified health IT should attest “yes,” if any of the following are true:


- Developer self-develops predictive DSIs for use in their certified Health IT Module; or
- Developer’s Health IT Module enables or interfaces with predictive DSIs developed by its users or customers, such as a health care organization or medical center; or
- Developer’s Health IT Module enables or interfaces with predictive DSIs developed by an “other party,” such as a separate software developer(s)

AND


- Predictive decision support intervention is based on any of the data expressed in the USCDI standards (§ 170.213)

“Enabled by or Interfaced with”

- Enables = The developer of certified health IT has the technical capability to support a predictive model or DSI within the developer’s Health IT Module
 - User-, third-party, and self-developed applications
 - Standalone applications used within or as a part of a Health IT Module
 - For example, if the calculations for a predictive DSI occur within the Health IT Module, either to or through a standalone app used within a Health IT Module or an app developed by a developer of certified health IT for use within a Health IT Module, we would consider this “enabling”
 - Includes instances where predictive DSIs are enabled by default and instances where they can be enabled by users
- Interfaces with = The Health IT Module facilitates either (1) the launch of a predictive model or DSI or (2) the delivery of a predictive model or DSI output(s) to users when such a predictive model or DSI resides outside of the Health IT Module
 - For example, scenarios where the calculations for a predictive DSI occur outside the Health IT Module, and the predicted value or output gets sent to or through a Health IT Module (or to or through an app used within or as part of a Health IT Module) would be considered to “interface with”
 - A Health IT Module would also “interface with,” a predictive DSI in scenarios where an application is launched from a certified Health IT Module, including through the use of a single sign-on functionality



“enables” is about the certified health IT being a container within which a predictive model or DSI can be used (either as an app or as part of the Health IT Module)

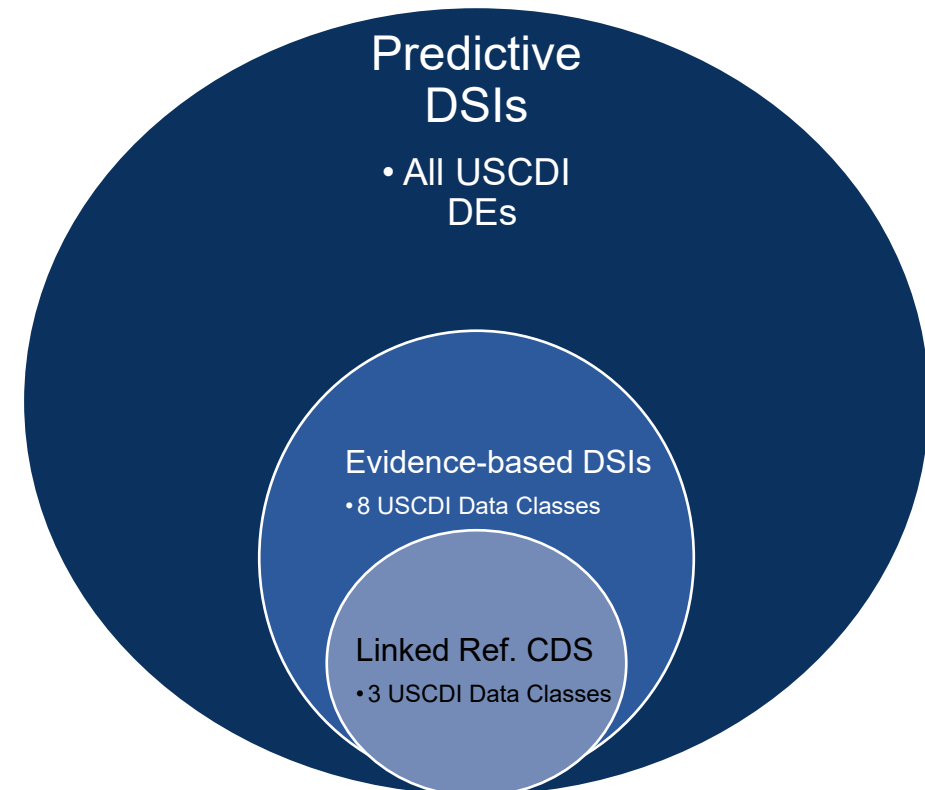


“interface with” is about the certified health IT being a door, through which actions can be taken to launch or deliver a predictive model or DSI

Which DSIs Would Need to Adhere to Relevant ONC Proposed Requirements?

- ***For predictive decision support interventions:***
All DSIs that use any USCDI Data Elements (DEs) at § 170.213
 - This is established in 170.315(b)(11)(v)
 - Yes – the Health IT Module enables, or interfaces with, electronic predictive decision support interventions based on any of the data classes expressed in the standard in § 170.213
- ***For evidence-based decision support interventions:*** All DSIs that use:
 - Problems; Meds; Allergy and intolerances; Demographics; Labs; Vital Signs; and Procedures, according to USCDI at § 170.213
- ***For Linked referential CDS:*** All DSIs that use
 - Problems, Meds, Demographics (no change to current reg)

Scope of relevant DSIs is not based on function/intended use, but on data elements used by the DSIs





Proposed Implementation Timeline and RWT Implications

- Health IT Modules certified to § 170.315(a)(9) would need to update and provide their customers with technology certified to § 170.315(b)(11) and comply with these new requirements by December 31, 2024
 - Health IT Modules may be certified to (a)(9) and/or (b)(11) until December 31, 2024
- Propose to modify the Base EHR definition in § 170.102 to include § 170.315(b)(11)
 - (a)(9) will expire January 1, 2025, and (b)(11) will replace (a)(9) in the Base on and after January 1, 2025
- Developers of certified health IT with Health IT Module(s) certified to § 170.315(b)(11) would be required to submit real world testing plans and corresponding real world testing results, consistent with other “(b)-criteria” in § 170.405(a)
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2023 plans
 - Measures demonstrating conformance to requirements, self-identified by developer
 - Annual cycle of RWT plans and results publicly available via CHPL
- Propose to add (a)(9) to the list of applicable criteria for Real World Testing, effective as of a final rule until it expires



Source Attributes

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Technical & Performance

- Information about how the predictive DSI “works” made available to users, in plain language and via direct display, drill down, or link out:
 - Output and intended use, out of scope use(s), description of training data, external validation, update schedule, etc.
 - Like a “nutrition label”; leverage existing “source attributes” certification requirement
- Supportive of health equity by design:
 - Identification of REL, SOGI, SDOH, & Health Status data elements used
 - Information on validity and fairness of prediction in test and local data (if available)
- Additional enhancements that enable:
 - Authoring and revision capability for users
 - User feedback capabilities and feedback exports for quality improvement of DSIs

Governance

- Public disclosure regarding how certified health IT developer manages risks and govern predictive DSIs:
 - Risk analysis (8 risk types): validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy
 - Risk mitigation of those risks
 - Governance processes, including data management
- Summary documentation must be:
 - Publicly accessible through hyperlink without precondition
 - Reviewed annually for updates
- Detailed documentation must be:
 - Available to ONC upon request from ONC for each predictive DSI the certified health IT enables or interfaces with
 - Reviewed annually for updates

Oversight

- Conformance to proposed new requirements through Real World Testing (RWT) Program:
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2024 plans
 - Annual cycle of RWT plans and results publicly available via the Certified Health IT Product List (CHPL)
 - Measures demonstrating conformance to requirements, self-identified by developer
- Summary of intervention risk management practices made publicly available
- Detailed risk management practices made available to ONC upon request from ONC

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Technical & Performance

- Information about how the predictive DSI “works” made available to users, in plain language and via direct display, drill down, or link out:
 - Output and intended use, out of scope use(s), description of training data, external validation, update schedule, etc.
 - Like a “nutrition label”; leverage existing “source attributes” certification requirement
- Supportive of health equity by design:
 - Identification of REL, SOGI, SDOH, & Health Status data elements used
 - Information on validity and fairness of prediction in test and local data (if available)
- Additional enhancements that enable:
 - Authoring and revision capability for users
 - User feedback capabilities and feedback exports for quality improvement of DSIs



Transparency: Emerging guidelines and best practices

- Academia and industry are developing ways to demonstrate technical and performance dimensions of predictive algorithms in health care
 - Reporting guidelines, such as model cards & datasheets for datasets (aka algorithmic nutrition labels) that provide information on
 - Predictive model details, development processes, performance, and maintenance requirements (to identify “model drift”)
 - [Model Cards for Model Reporting](#)
 - [Datasheets for Datasets](#)
- Government, academia, and industry are coalescing on the need to manage risks at the organizational level
 - AI Governance Models
 - Duke [Algorithm-Based Clinical Decision Support](#) (ABCDS) Oversight
 - Risk management practices
 - [Office of the Comptroller of the Currency](#) handbook for the financial sector
 - [NIST Risk Management Framework](#) that sector agnostic



Sources of Source Attributes

We emphasized source attribute information that

1. Were most commonly included in existing, reviewed reporting guidelines
2. Would be most meaningful and interpretable in the context of health IT users and developers
3. Were focused on health equity, fairness, and identifying issues of bias
4. Were intended to show that the model would perform effectively outside of the specific context in which it was developed

Goals

- Identify minimum necessary attributes
- Based on existing model reporting guidelines
- Balance prescriptiveness and flexibility to accommodate varied applications, contexts, and use cases
- Align with existing reference material (e.g., NIST AI RMF, WH Blueprint, WH E.O.s)
- Support emerging industry-led efforts (e.g., CHAI and Health AI Partnership)

DSI - Health IT Modules are not required to enable or interface with predictive DSIs, but...

If a Health IT Module enables or interfaces with predictive DSIs, we are proposing that the module must make information about additional Source Attributes available to provide users transparency on how the predictive DSI was designed, developed, trained, evaluated, and should be employed.

- **Intervention Details**

1. Output of the intervention
2. Intended use of the intervention
3. Cautioned out-of-scope use of the intervention

- **Intervention Development**

1. Input features of the intervention including description of training and test data
2. Process used to ensure fairness in development of the intervention
3. External validation process, if available

- **Quantitative Measures of Intervention Performance**

1. Validity of prediction in test data
2. Fairness of prediction in test data
3. Validity of prediction in external data, if available
4. Fairness of prediction in external data, if available
5. References to evaluation of use of the model on outcomes, if available

- **Ongoing Maintenance and Use**

1. Update and continued validation/fairness assessment schedule
2. Validity of prediction in local data, if available
3. Fairness of prediction in local data, if available

Intellectual Property

- The proposals in § 170.315(b)(11)(vi)(C) would not require disclosing or sharing intellectual property (IP) existing in the developer's health IT (including other parties' IP)
- The proposed requirement would not provide information about or report any details of the specific code, pipeline, statistical processes, or algorithms used to generate model predictions, which might be considered the developer's intellectual property



Source Attributes Prescriptiveness

- We request comment on whether there are items contained within the proposed source attributes that we should explicitly require as elements of source attributes information.
 - Specific attention to three Source Attributes with multiple “should” components:
 - “Intended use of the intervention,”
 - “Input features of the intervention including description of training and test data”
 - “External validation process, if available”

Intervention Details

Output of the intervention is a description of the value that the model produces as an output, including whether the output is a prediction, classification, or other type of output

- Users evaluating the model or deciding whether to use it should know what the model is predicting to ensure that the output is directly relevant to the way in which the users intend to use it

Intended use of the intervention is a description of the intent of the model developers in how the model is meant to be deployed and used, including its intended role in the identified use case. This information should clarify:

- Whether the model is intended for specific or general tasks and what those tasks are;
- Who the intended patient population is;
- Who the intended users of the model are, as well as the intended action of the user;
- The role of the model (e.g., whether it informs, augments, or replaces clinical management), which may be most clearly conveyed through use of a taxonomy such as those described by the International Medical Device Regulators Forum (IMDRF), American Medical Association, Consumer Technology Association, and others; and
- The logic underlying the model; for instance, the exact question the algorithm is supposed to answer, how it fits into specific clinical decision-making, and in what ways the inputs are appropriate to answer that question and, if appropriate, how that logic is associated with how the model should be used.

Intervention Development

Cautioned out-of-scope use of the intervention is a description of tasks, situations, or populations to which the model developer cautions a user against applying the predictive model. This description should include:

- Known risks, inappropriate settings, inappropriate uses, or known limitations of the model
- Description should inform users about tasks, situations or populations related to the intended use of the model in which the model may not perform as expected

Input features of the intervention including description of training and test data should include:

- Exclusion and inclusion criteria that influenced who was included in data sets;
- Statistical characteristics—including sample size—of the demographic and other key variables in these data to assess representativeness;
- The source and clinical setting from which the data was generated
- The extent of missing values in the training and testing data sets; and
- Other attributes related to data quality, such as the comprehensiveness of the data and the process of collecting the data should be included as the developer determines what is relevant while examining the data during pre-processing, creation, and testing of the model.

Intervention Development

Process used to ensure fairness in development of the intervention is a description of the approach the model developer has taken to ensure that the model output is fair. This should include:

- Approaches to manage, reduce, or eliminate bias in models and could be similar to a brief synopsis of risk mitigation practices and outcomes related to fairness for this DSI
- Many such approaches exist; however, there is no universal best process to ensure fairness
- For example, this attribute might state that in pre-processing the data before training the model, the developers employed a “disparate impact remover” transformation across race or ethnicity groups based on a well-known approach

External validation process, if available is a description of how and in what source, clinical setting, or environment a model’s validity and fairness has been assessed other than the source training and testing data. This should include:

- Who conducted the external testing (e.g., the model developer, developer of certified health IT, or an independent third party);
- The setting from which the external data was derived;
- The demographics of patients in external data; and
- A brief description of how external validation was carried out.





Quantitative Measures of Intervention Performance

- **Validity of prediction**

- In test data and, if available, external data and local data is the presentation of the measure or set of measures related to the model's validity (often referred to as performance) tested, respectively, in data derived from the same source as the initial training data, in data from an external source, and in data local relative to its current use.
 - This proposal would not prescribe the specific performance or validation measures to be used or included as part of the source attributes requirements but would require that some performance or validation measure(s) be used and included in the source attribute.

- **Fairness of prediction**

- In test data and, if available, external data and local data is the presentation of the measure or set of measures related to the model's fairness (evaluation of fairness in a model) in terms of the accuracy of its output across certain groups in data derived from the same source as the initial training data, in data from an external source, and in data local relative to its current use.
 - Numerous approaches and related measures exist to measure the fairness of model outputs. Examples of potential fairness measures include positive predictive parity, false positive error rate balance and false negative error rate balance, equivalent calibration within groups, and mean residual difference

- **References to evaluation of use of the model on outcomes, if available** are bibliographic citations or links to evaluations of how well the intervention, or model on which it is based accomplished specific objectives such as reduced morbidity, mortality, length of stay or other important outcomes



Ongoing Maintenance and Use

- **Update and continued validation or fairness assessment schedule** is a description of the process and frequency by which the model's performance is measured and monitored in the local environment and corrected when risks related to validity and fairness are identified
 - Information should also include how often performance is evaluated and how often the model is updated to provide users with insight into the likelihood that the model may have degraded (i.e., no longer provides valid or accurate predictions) since it was last updated
- **Validity and Fairness in Local Data**
 - As previously described



Additional Considered Source Attributes Example

- Intervention Details
 - Information on explainability and interpretability
 - Whether a DSI meets the definition of a medical device under the FDA definition
- Intervention Development
 - Details on how model prediction and classification cut-points were selected
 - Security and privacy-preserving approaches included in model development
- Quantitative Measures of Intervention Performance
 - Model calibration or calibration curve
 - Confidence or prediction intervals or other measures of uncertainty
- Ongoing Maintenance of Intervention Implementation and Use
 - Whether the model is 'online' or 'unlocked'
 - Any additional organizational or technical controls in place to evaluate the impact of the online or unlocked updating and results of that evaluation.
 - The controls in place to update the descriptions of source data to reflect the changing composition of the data.



Availability of Source Attributes to the Public

- We solicit comment on whether we should require developers of certified health IT with Health IT Modules certified to proposed § 170.315(b)(11) to make all source attributes information publicly available or accessible, for example, on a website, similar to the existing API documentation requirement in § 170.315(g)(10)(viii)(B).
- We solicit comment on whether having this information publicly available would be beneficial for potential users that purchase models or associated technology or software, and would help inform them prior to procurement of certified health IT and procurement of predictive DSIs integrated with certified health IT.
- We also solicit comment on whether having this information publicly available would improve public confidence in predictive DSIs by enabling research on source attribute information.

Patient Access to Source Attributes

- Patients want to know if AI is being used in their care, and understand how and why it is being used in their care. We understand an emerging trend is for health care providers to inform patients about the use of these technologies, including predictive DSIs, in making decisions about their care.
- We solicit comment on whether existing Program requirements in the Communications condition and maintenance of certification requirements in § 170.403 are sufficient to ensure open and transparent discussion regarding the use of predictive DSIs in patient care – including discussion between users of certified health IT and patients. We are especially interested in whether we should require developers of certified health IT to provide the technical capability for users to support patients electronically accessing underlying source attribute information (e.g., through a patient portal) for predictive DSIs or otherwise indicate to a patient when a predictive DSI was used to make decisions about the patient in the course of the patient's care.





Consensus Metrics and Standards

- We also solicit comment on testing or assessment tools that might further support transparency and trustworthiness including
 - Consensus metrics and technical standards for evaluating fairness (assessing for bias) and validating performance (including testing performance in different populations and evaluating applicability or generalizability) of predictive models that are enabled by or interface with Health IT Module(s) prior to and during deployment
 - Development and engineering of algorithmic impact assessments (AIAs)
 - Development of documentation of datasets used, such as datasheets for datasets and data cards as well as tools that could be useful in these areas so that Health IT Modules certified to §170.315(b)(11) can demonstrate it meets a given requirement on an ongoing basis



Authoring and Revising Source Attributes

- We propose in § 170.315(b)(11)(vi)(E) that Health IT Modules enable users to author attributes and revise attributes beyond what is proposed in to support the ongoing evolution of what source attributes are important to users to make informed decisions regarding the DSI's recommendation(s).
 - Pertains to both evidence-based DSIs and predictive DSIs
 - Means that a Health IT Module would need to support the technical ability for a limited set of identified users to create new or revised attribute information alongside other source attribute information proposed
 - Example: a hospital that develops its own predictive DSI that is interfaced with a certified Health IT Module would be able to create new or revise existing source attributes information related to that predictive DSI that is made available through the certified Health IT Module without the developer of certified health IT's direct involvement.



DSI Feedback Loops

- In the 2015 Edition Proposed Rule, we proposed to adopt new functionality that would require a Health IT Module to be able to record at least one action taken, and by whom it was taken, when a CDS intervention is provided to a user
 - For example, whether the user viewed, accepted, declined, ignored, overrode, provided a rationale or explanation for the action taken, took some other type of action not listed here, or otherwise commented on the CDS intervention) (80 FR 16821).
 - We also proposed that a Health IT Module certified to § 170.315(a)(9) be able to generate either a human readable display or human readable report of the responses and actions taken and by whom when a CDS intervention is provided (80 FR 16821).
- In the 2015 Edition Final Rule, we noted that many commenters stated that current systems already provide a wide range of functionality to enable providers to document decisions concerning CDS interventions and that such functionality is unnecessary to support providers participating in the EHR Incentive Programs (80 FR 62622).



DSI Feedback Loop Proposal

- We propose that a Health IT Module certified to § 170.315(b)(11) must be able to export such feedback data, including but not limited to the intervention, action taken, user feedback provided (if applicable), user, date, and location, so that the exported data can be associated with other relevant data.
- We propose that such feedback data be available for export by users for analysis in a computable format, so that it can be associated with other relevant data, such as diagnosis, other inputs into the DSI, and the outputs of the DSI for a particular patient, to evaluate and improve DSI performance.
- In addition to quality improvement of the DSI, such an export would facilitate research, associating feedback data with other relevant data, and linking the DSI to patient health outcomes, including assisting in identifying and reducing health disparities and possible discriminatory outcomes.



Intervention Risk Management

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Technical & Performance

- Information about how the predictive DSI “works” made available to users, in plain language and via direct display, drill down, or link out:
 - Output and intended use, out of scope use(s), description of training data, external validation, update schedule, etc.
 - Like a “nutrition label”; leverage existing “source attributes” certification requirement
- Supportive of health equity by design:
 - Identification of REL, SOGI, SDOH, & Health Status data elements used
 - Information on validity and fairness of prediction in test and local data (if available)
- Additional enhancements that enable:
 - Authoring and revision capability for users
 - User feedback capabilities and feedback exports for quality improvement of DSIs

Governance

- Public disclosure regarding how certified health IT developer manages risks and govern predictive DSIs:
 - Risk analysis (8 risk types): validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy
 - Risk mitigation of those risks
 - Governance processes, including data management
- Summary documentation must be:
 - Publicly accessible through hyperlink without precondition
 - Reviewed annually for updates
- Detailed documentation must be:
 - Available to ONC upon request from ONC for each predictive DSI the certified health IT enables or interfaces with
 - Reviewed annually for updates

Oversight

- Conformance to proposed new requirements through Real World Testing (RWT) Program:
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2024 plans
 - Annual cycle of RWT plans and results publicly available via the Certified Health IT Product List (CHPL)
 - Measures demonstrating conformance to requirements, self-identified by developer
- Summary of intervention risk management practices made publicly available
- Detailed risk management practices made available to ONC upon request from ONC

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Governance & Risk Management

- Public disclosure regarding how certified health IT developer manages risks and govern predictive DSIs:
 - Risk analysis (8 risk types): validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy
 - Risk mitigation of those risks
 - Governance processes, including data management
- Summary documentation must be:
 - Publicly accessible through hyperlink without precondition
 - Reviewed annually for updates
- Detailed documentation must be:
 - Available to ONC upon request from ONC for each predictive DSI the certified health IT enables or interfaces with
 - Reviewed annually for updates

NIST Risk Management Framework

“ AI risk management can drive responsible uses and practices by prompting organizations and their internal teams who design, develop, and deploy AI to think more critically about context and potential or unexpected negative and positive impacts. Understanding and managing the risks of AI systems will help to enhance trustworthiness, and in turn, cultivate public trust.

- Govern 6 – Policies and procedures are in place to address AI risks and benefits arising from third-party software and data and other supply chain issues.
- Map 4 – Risks and benefits are mapped for all components of the AI system including third-party software and data.
- Measure 2 – AI systems are evaluated for trustworthy characteristics.
- Manage 1 – AI risks based on assessments and other analytical output from the MAP and MEASURE functions are prioritized, responded to, and managed.



Background on IRM

- Given a lack of healthcare sector-specific guidance and the nascency of several emerging efforts for risk management of predictive software, our proposals would not require a specific framework, guideline, or approach that such developers of certified health IT must use – only that they employ or engage in IRM practices in accordance with proposed requirements in § 170.315(b)(11)(vii)(A) through (D)
- We view our proposals for risk management of predictive DSIs in § 170.315(b)(11)(vii) as complementary to our proposals for predictive DSI source attributes in § 170.315(b)(11)(vi)(C)
- The proposed source attributes information requirement is meant to provide users and implementers with sufficient information to understand how the model was designed, developed, and tested, including the model's purpose, known limitations, and intended use(s)
- Correspondingly, the proposals for intervention risk management would provide users, implementers, and the wider public, including patients, with information to understand how developers of certified health IT with Health IT Modules that enable or interface with predictive DSIs analyze, mitigate, and govern risks throughout the technology's life cycle



Pillars of Intervention Risk Management Proposal

Risk Analysis

- Should estimate the likelihood and magnitude of the negative impact (harm), or consequences, of each risk characteristic; to whom each risk applies (including, for example, individual, group, and societal harm); and the source of each risk

Risk Mitigation

- Should describe:
 - Practices used to prioritize or establish different levels of risk;
 - Practices to mitigate or minimize identified risks;
 - Change control plans or ongoing validation/updating processes
 - Processes to supersede, disengage, or deactivate deviations from intended use
 - Approaches to include SMEs in measuring / validating performance

Governance

- Should set an effective framework for risk management, with defined roles and responsibilities for clear communication of predictive DSI limitations and assumptions
- Should include setting and enforcing priorities for managing and using data as a strategic asset

Risk Analysis Categories - Validity

- NIST's AI RMF describes seven characteristics of trustworthy AI, and we propose to adapt these concepts and require that developers of health IT with certified Health IT Modules that enable or interface with predictive DSIs employ or engage in risk management practices related to the following characteristics:
 - Validity - Assessment of risk related to validity should include and consider the following areas:
 - Validation of the accuracy and completeness of data used in development and testing of the predictive DSI
 - Evaluation plans and results for validation in testing environments and ongoing evaluation in deployment;
 - Both technical validity and clinical validity, which is closely related to measurement of effectiveness such as those discussed in the proposed source attribute "References to evaluation of use of the model on outcomes" in § 170.315(b)(11)(vi)(C)(3)(v).



Risks to Reliability and Robustness

- “Reliability” indicates whether a model used in a predictive DSI consistently performs as required, without failure, for a given time interval, under given conditions. Assessment of reliability should include
 - Defining what range of behaviors is considered reliable for a model
 - The error rate considered acceptable
 - The results of evaluations that demonstrate reliability in both testing and deployed environments
- “Robustness” or generalizability is the ability of a model used in a predictive DSI to maintain its level of performance under a variety of circumstances. Assessment of robustness should
 - Evaluate limitations of the model based on the source of the training and testing data used and how features of that data and its source might relate to performance outside of the training and testing environment, which are likely to relate to information discussed in the proposed source attribute “input features of the intervention including description of training and test data”



Risks to Fairness and Intelligibility

- “Fairness,” as noted above in this section, is defined by a lack of bias against certain groups, and fairness enhancing (or bias managing) processes seek to ensure that models are fair. NIST has identified three major categories of AI bias that should be addressed and managed to enhance fairness of models:
 - Systemic
 - Computational and statistical
 - Human-cognitive
- “Intelligibility” refers to the extent to which the predictive DSI can be understood, often through a representation of the mechanisms underlying an algorithm’s operation and through the meaning of AI systems’ output in the context of its designed functional purpose. In assessing intelligibility, developers of certified health IT should
 - Delineate the expected and acceptable context of use, including the intended users and operational setting.
 - Assess whether the predictive DSI provides intelligible information as an output that will allow for its intended users to make effective interpretation of relevant predictive DSI behavior when applied or used in the expected operational setting.



Risks to Safety and Security

- “Safety” as a concept is highly correlated with risk and generally denotes that the product is free from any unacceptable risks and the probable benefits outweigh any probable risk. Developers should assess
 - Who could be injured,
 - when injury could arise and how injury could arise, engaging external parties in this assessment when such risks are not obvious
 - Implement procedures for regularly evaluating safety
- “Security” (and relatedly resilience) is a predictive DSI's and model's ability to withstand adversarial attacks, or more generally, unexpected changes in its environment or use. In assessing security, developers should consider
 - Common IT security concerns related to the exfiltration of models, training data, or other intellectual property through the technology's endpoints
 - Potential weaknesses in the controls for the access, transmission, and storage of sensitive information

Risks to Privacy

- “Privacy” refers generally to the norms and practices that help to safeguard human autonomy, identity, and dignity, as well as data autonomy and intrusions on information about an individual. Analysis of privacy should
 - Consider the NIST Privacy Framework and application of NIST Privacy Risk Assessment Tool
 - Like safety and security, specific technical features of AI or ML-enabled technologies may promote or reduce privacy, and assessors can identify how the processing of data could create privacy-related problems



Risk Mitigation

- We propose in § 170.315(b)(11)(vii)(A)(2) “Risk Mitigation” to require implementation of practices to mitigate risks associated with predictive DSIs. Risk mitigation practices implemented by developers of certified health IT should cover the following:
 - Practices to prioritize (establish different levels of) risks based on their impact and likelihood
 - Practices to mitigate or minimize identified potential risks
 - Change control plans, including schedule of validation and updating processes
 - Processes to supersede, disengage, or deactivate an existing predictive decision support intervention that demonstrate performance or outcomes that are inconsistent with their intended use
 - Approaches to including subject matter experts in measuring and validating whether the system is performing consistently with their intended use and as expected in the specific deployment setting

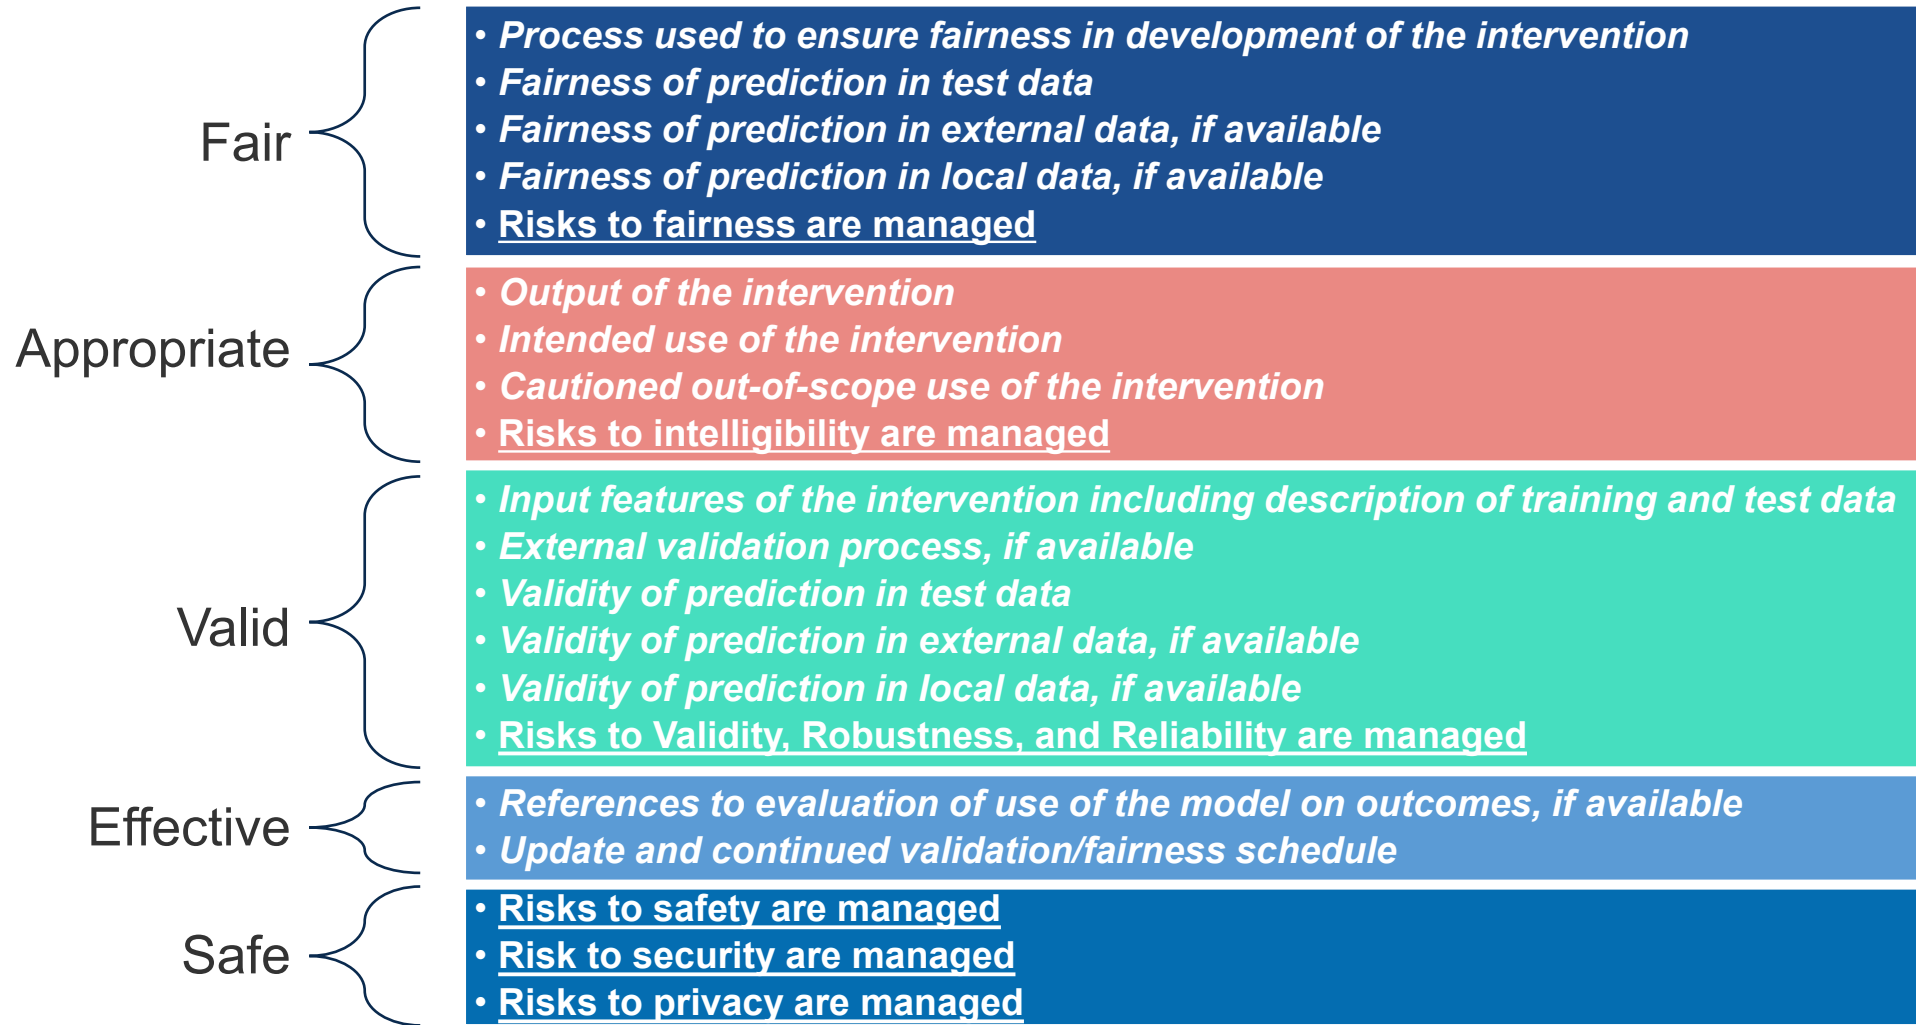


Governance

- We propose to require health IT developers to establish policies and implement controls for predictive decision support intervention governance, including how data are acquired, managed, and used in a predictive decision support intervention
 - Governance should encompass models, software and data developed or provided by other parties as well as internally developed interventions
 - We expect developers of health IT to consider how the policies and controls they implement for data governance ensure the responsible acquisition, management, and use of data, including how the developer of certified health IT factors in and addresses ethical, legal, and social implications (ELSI) underlying data collection (acquisition) and use
- Our use of the term “policies” means statements of management intent regarding the objectives and required components of intervention risk management.
- Our use of the term “controls” means a system of internal controls that the developer has in place to implement the associated risk management policies, including those at the organizational and technology level
 - For example, processes for controlling the quality of the data inputs; internal and external audits; process to escalate conflicting views between the model development and validation groups



Source Attributes and IRM Information Help Users Determine the FAVES of a Predictive DSI





Oversight & Implementation

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Technical & Performance

- Information about how the predictive DSI “works” made available to users, in plain language and via direct display, drill down, or link out:
 - Output and intended use, out of scope use(s), description of training data, external validation, update schedule, etc.
 - Like a “nutrition label”; leverage existing “source attributes” certification requirement
- Supportive of health equity by design:
 - Identification of REL, SOGI, SDOH, & Health Status data elements used
 - Information on validity and fairness of prediction in test and local data (if available)
- Additional enhancements that enable:
 - Authoring and revision capability for users
 - User feedback capabilities and feedback exports for quality improvement of DSIs

Governance

- Public disclosure regarding how certified health IT developer manages risks and govern predictive DSIs:
 - Risk analysis (8 risk types): validity, reliability, robustness, fairness, intelligibility, safety, security, and privacy
 - Risk mitigation of those risks
 - Governance processes, including data management
- Summary documentation must be:
 - Publicly accessible through hyperlink without precondition
 - Reviewed annually for updates
- Detailed documentation must be:
 - Available to ONC upon request from ONC for each predictive DSI the certified health IT enables or interfaces with
 - Reviewed annually for updates

Oversight

- Conformance to proposed new requirements through Real World Testing (RWT) Program:
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2024 plans
 - Annual cycle of RWT plans and results publicly available via the Certified Health IT Product List (CHPL)
 - Measures demonstrating conformance to requirements, self-identified by developer
- Summary of intervention risk management practices made publicly available
- Detailed risk management practices made available to ONC upon request from ONC

Snapshot of Proposals to Promote Transparent & Trustworthy DSIs through the ONC Health IT Certification Program

Oversight & Implementation

- Conformance to proposed new requirements through Real World Testing (RWT) Program:
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2024 plans
 - Annual cycle of RWT plans and results publicly available via the Certified Health IT Product List (CHPL)
 - Measures demonstrating conformance to requirements, self-identified by developer
- Summary of intervention risk management practices made publicly available
- Detailed risk management practices made available to ONC upon request from ONC



Oversight through Transparency & Real World Testing

- Summary information for intervention risk management practices should be publicly available via a publicly accessible hyperlink that allows any person to directly access the information without any preconditions or additional steps.
 - Clinicians, patients, health systems, and the public could use this information to bolster their trust in the developers of certified health IT and those certified Health IT Modules that enable or interface with predictive DSIs.
- Developers of certified health IT with Health IT Module(s) certified to § 170.315(b)(11) would be required to submit real world testing plans and corresponding real world testing results, consistent with other “(b)-criteria” in § 170.405(a)
 - RWT for all DSI types (predictive, evidence-based, and linked referential) beginning for 2024 plans
 - Measures demonstrating conformance to requirements, self-identified by developer
 - Annual cycle of RWT plans and results publicly available via CHPL
- Propose to add (a)(9) to the list of applicable criteria for Real World Testing, effective as of a final rule until it expires



Office of the National Coordinator
for Health Information Technology

Contact ONC



Phone: 202-690-7151



Health IT Feedback Form:

<https://www.healthit.gov/form/healthit-feedback-form>



Twitter: [@onc_healthIT](https://twitter.com/onc_healthIT)



LinkedIn: [Office of the National Coordinator for Health Information Technology](https://www.linkedin.com/company/office-of-the-national-coordinator-for-health-information-technology)



Youtube:

<https://www.youtube.com/user/HHSONC>

HealthIT.gov

Subscribe to our weekly eblast
at [healthit.gov](https://www.healthit.gov) for the latest updates!

