

To: All Users and Developers of Clinical Decision Support Software

From: Bradley Merrill Thompson, General Counsel, CDS Coalition

Date: April 27, 2017

Re: Request for Comment on Draft Industry Guidelines

The CDS Coalition has developed the attached voluntary guidelines for use by clinical decision support, or CDS, software developers to guide the design of such software. The purpose of these guidelines is to ensure that software, as much as possible, leaves the user firmly in control of the clinical decision-making. As explained more in the guidelines themselves, we have been also influenced by some recent legislation that articulates the dividing line between FDA regulation and the practice of medicine.

The coalition is comprised of large and small software developers, including those traditionally from the medical space as well as those new to the area. The coalition also includes hospitals and other users of such software, as well as those who make pharmaceutical products and medical devices that are guided by the software. But we recognize that there is a vast universe of organizations out there impacted by CDS, and so we would like to solicit comments from all those affected.

Our plan is to collect comments through July 1, 2017, and then make appropriate revisions. After that, we plan to present these guidelines to FDA for their feedback.

If you have any questions or concerns, please let me know. Please direct your comments to me, Brad Thompson, at bthompson@EBGLaw.com.

We would very much appreciate hearing your thoughts.

Voluntary Guidelines for the Design of Clinical Decision Support Software to Assure the Central Role of Healthcare Professionals in Clinical Decision- Making

Developed by the CDS Coalition

In December 2016, Congress enacted the 21st Century Cures Act, adding a new subparagraph 520(o)(1)(E) on Clinical Decision Support (CDS) software to the Federal Food, Drug and Cosmetic Act. That new subparagraph carves out certain software from the scope of FDA regulation. Ultimately, the United States Food and Drug Administration (FDA) will be responsible for administering and enforcing that new provision.

At the same time, the CDS Coalition believes that the CDS industry is best served by self-regulating through adopting basic software design guidelines. In particular, the coalition believes the CDS industry should develop guidelines that foster the design of software in such a manner that healthcare professional users will be able to independently review the basis for the recommendations the CDS software produces, such that the professionals will not need to rely primarily on the recommendations.

New subparagraph 520(o)(1)(E) on CDS addresses that same topic. As a statute, it establishes the minimum design elements necessary to avoid FDA regulation. These industry guidelines are not an attempt to simply apply that legislation, but rather establish voluntary industry guidelines for accomplishing those same design features and more.

In the development of these guidelines, the statute was relevant in the sense that the coalition wanted to establish guidelines that, if met, would assure that the software is not regulated by FDA. At the same time, the coalition wanted to go beyond merely repeating the statute. Instead, the coalition sought to identify guidelines that would more broadly assure some consistency in approach by software developers with regard to design features intended to assure the central role of healthcare professionals in clinical decision-making. The coalition believes that leaving the healthcare professional in control of the decision-making helps to ensure the appropriateness of the clinical decision-making. Software that does not meet these guidelines can be quite safe and effective, but may require additional validation because of the risk that such software will supplant the judgement of healthcare professionals.

To summarize, compliance with these guidelines during software development should assure that:

1. healthcare professional users remain fully in control of their own decision-making when using such software; and
2. such software is not regulated by FDA

These guidelines are intended not to dip below the requirements of the new law, but may well at times exceed those requirements. Software that does not meet these guidelines necessarily relies more on the proper functioning of the software over the judgment of the healthcare professional, and therefore may require a higher level of validation and potentially FDA oversight.

The coalition acknowledges that software developers do not need to follow these guidelines to avoid FDA regulation, and that there are other ways of meeting the requirements of the new subparagraph 520(o)(1)(E) on CDS. In addition, the coalition acknowledges that FDA, not industry, determines whether a given approach meets the requirements of the law.

Nonetheless, the coalition has done its best in these guidelines to identify one set of design elements that the coalition thinks meets the requirements of the statute. The coalition includes its analysis of the statutory requirements in Appendix D of these guidelines. The coalition's plan is to consult FDA to see if the agency agrees that these guidelines present one way a developer can avoid FDA regulation.

Scope of these Guidelines

Clinical decision support software is software that:

- Uses patient-specific information and organized clinical knowledge;
- Performs some analysis using that information and knowledge (rather than simply display or transmit the information);
- Produces a particular actionable result for the diagnosis, treatment or management of a disease or condition for a particular patient (as opposed to multiple options, which is not CDS); and
- Is standalone software, and not an accessory to a medical device.

That last point is important in distinguishing CDS from software that used in conjunction with a specific medical device to analyze the output of that device.

These guidelines further only address CDS intended for use by healthcare professionals.

Overview of Voluntary Guidelines for Enabling Independent Review of the Basis for CDS Recommendations

Healthcare professionals should be free to practice medicine to the best of their ability and in reliance on their own professional judgment, within the community standard of care. To empower medical professionals to use their best professional judgment, it is important that CDS software enable healthcare professional users to independently review the basis of the recommendations the software makes. By so doing, the software developers can ensure that the


healthcare professionals will not need to rely primarily on the software, but rather will be able to rely primarily on their own professional judgment.

Thus, to comply with these guidelines, developers must answer two fundamental questions affirmatively:

1. Are healthcare professional users able to independently review the basis of the recommendations the software makes?
2. Under the intended circumstances of use, will healthcare professional users not need to rely primarily on the software?

As already explained, software that does not meet these guidelines necessarily relies more on the proper functioning of the software in lieu of the judgment of the healthcare professional, and therefore may require a higher level of validation and potentially FDA oversight.

To answer the first of those two questions, CDS software developers can use three criteria to determine if the CDS is intended to enable the healthcare professional user to independently review the basis of the recommendations the software makes.



Each of these criteria would need to be met to enable the healthcare professional user to independently review the basis of the recommendations, as explained below.

1. **Transparency.** Does the software provide enough information for the user to understand and be able to evaluate the clinical basis for the software recommendation? This includes disclosure of the following:
 - a. What the software does and does not do. It is important in this regard that the labeling for the software communicate to the user the limits of the software's functionalities.
 - b. The information inputs used by the software. This includes (i) patient specific information, and (ii) the source of the clinical information or decision rules such as practice or professional guidelines that the software uses to analyze the patient information.

Furthermore, if a decision informed by the recommendation of the software could lead to serious injury, permanent impairment or death for the patient, the developer may also need to provide:

- c. An indication of the certainty or reliability of the output, including as appropriate confidence levels and/or ranking of alternatives

- d. The clinical rationale for the recommendations. This goes beyond merely identifying the source of the clinical rules, and includes a reasonable explanation of the clinical logic by which the software arrived at its specific recommendation based on patient specific information.

Most CDS is intended to aid a trained user in decision-making, but not to be a substitute for the user's expertise and judgment. On the other hand, if the CDS does not enable the intended user to sufficiently understand the recommendation made by the software and equally importantly, the basis for the recommendation, such CDS runs the risk of being used as a substitute for the user's expertise and judgment. In such cases, the software may need to be validated to a higher degree, and potentially subjected to FDA regulation.

This first criterion – transparency – requires more elaboration depending on the various types of data inputs and other design factors. As a result, the coalition has developed specific design guidelines to ensure the transparency of CDS, and includes those in Appendix A of these guidelines.

2. **Competent Human Intervention.** Is the intended user competent – through training, experience or otherwise - to make the clinical decision in question without the CDS? The education and experience required to be competent depends on the nature of the decision. A nurse, a primary care physician, and a specialist are each competent to make different types of decisions, as are pharmacists, home health aides and other health professionals and care-givers.

CDS intended to be used to extend a user's decision-making ability beyond his/her qualifications could mean that the healthcare professional is not able to independently review the basis for the recommendation. However, CDS that merely assists the user in applying her existing qualifications does enable independent review. CDS that, for example, collects, calculates, sorts, or otherwise gathers and presents information which the user is competent in interpreting (while easing the burden of data gathering or processing) should not, by itself, preclude the professional from independently reviewing the basis of the recommendations.

It is important to note that it may be prudent for decision-makers to also consider data outside of what the CDS has collected. Competent decision-makers will recognize that need and incorporate such information into their decision making process.

Like the first criterion, this second criterion –competent human intervention– requires more elaboration depending on the various types of data inputs and other design factors. As a result, the coalition has developed specific design guidelines to ensure competent human intervention, and includes those in Appendix B of these guidelines.

3. **Sufficient Time to Reflect.** Based on the intended use, is the user expected to have enough time to reflect on the software output before making a decision? The amount of time available to reflect will depend on the acuity of the condition, and how much time can lapse before the patient receives medical care without risk. The amount of time needed to reflect may also depend on the complexity of the decision.

If the intended user does not have enough time to independently consider the data inputs, in practice the user may not be able to independently review the basis of the recommendation.

Like the previous two, this third criterion –sufficient time to reflect– requires more elaboration depending on the various types of data inputs and other design factors. As a result, the coalition has developed specific design guidelines to ensure sufficient time to reflect, and includes those in Appendix C of these guidelines.

A Mitigating Factor: Other Relevant Clinical Information

To answer the second of the two questions posed by these guidelines, after determining that a user is able to independently review the basis for a recommendation, when assessing a user’s need to rely primarily on CDS software, the manufacturer can take into account the likelihood that the user has access to other useful sources of clinical information concerning the individual patient useful to making the diagnostic or treatment decision. Access to information about the patient outside of the CDS software mitigates the need of the user to rely primarily on the recommendations from the software.

In cases where the software inputs include all the relevant information concerning the individual patient, this simply means the same as the inputs portion of the transparency criterion above – the user has access to all of the information concerning the individual patient that has been inputted into the software available for direct review. In other words, this becomes redundant with the transparency criterion.

But in cases where the software considers a narrower set of inputs concerning the individual patient, this means the decision maker has access to additional information concerning the patient helpful to the decision. For example, it may mean the patient is available to the physician for direct physical examination, or that the physician is likely to have direct access to additional, useful diagnostic information such as radiological images or laboratory test results.

The bottom line is that the manufacturer can take into account, when assessing possible reliance by the user on the software, whether the user is likely to have available to her additional information outside of the CDS recommendation concerning the individual patient to independently arrive at the particular decision. To be clear, the developer must still assure the healthcare professional users are able to independently review the basis of the recommendations the software makes. This additional information merely helps the software developer assure that, in the end, the user does not need to rely primarily on the software for making the decision, and the second condition is therefore met.

Appendix A

Transparency Guidelines

If a given piece of software meets the Guidelines below, the transparency criterion under the test for determining whether software enables the user to independently review the basis for its recommendations would be met. To assure that healthcare professionals are able to independently review the recommendations of the software, as explained above, developers will also need to assure that they target appropriate healthcare professionals and use cases, and that users can be expected to have adequate time to reflect.

The Coalition recognizes that there may be other means to achieve transparency for the software in lieu of the steps recited below. These Guidelines are not intended as rigid requirements, but rather as a high level framework for manufacturers to follow. Manufacturers should analyze the purposes of these various guidelines and think about how the objectives can best be achieved for their particular software.

These Transparency Guidelines address two general questions:

1. What needs to be disclosed? We discuss four (4) types of disclosures:
 - i. The intended use
 - ii. The patient information and sources of clinical content used by the software
 - iii. The software output or recommendations
 - iv. The clinical rationale for the output or recommendation
2. How does it need to be disclosed?

Part of the complexity that these guidelines must face is that the information inputs can come in a wide variety of forms, and the recommendation outputs can likewise take many different forms. Without limiting the appropriate inputs and outputs, examples include:

Category of Information	Examples
Patient-specific information	<ol style="list-style-type: none">1. Data from electronic health record automatically transferred2. Data manually entered by patient3. Data manually entered by the user4. Data manually entered by another healthcare professional, not the user
Medical knowledge	<ol style="list-style-type: none">1. Peer-reviewed medical literature, either individual articles curated for this purpose or commercially-available databases of such literature2. Medical abstracts from meetings3. Consensus clinical guidelines from medical societies

	<ol style="list-style-type: none"> 4. Institution specific clinical guidelines 5. Expert medical opinion from individual professionals 6. Medical insights gleaned by software using machine learning, through digesting raw medical data from a given population of patients 7. Proprietary algorithms and calculators developed by drug or medical device companies 8. Clinical trial results not vetted through peer reviewing
Recommendations	<ol style="list-style-type: none"> 1. Quantitative calculated indices 2. Qualitative likely diagnosis 3. Qualitative likely best treatment option 4. Qualitative diagnostic or treatment options with quantitative rankings

These guidelines thus must contemplate all of those possible scenarios.

Proposed Disclosures by Number of Required Clicks

Most software, when it comes to presenting information to the end-user, is designed to have multiple access levels represented by the number of clicks required to access the relevant information. These access levels permit the end-user to go deeper and deeper into the details as the user clicks on links that take the user ultimately to source information.

When it comes to ensuring transparency, the goal is to balance the need to avoid clutter with the need to reveal important information. With regard to the various elements required for transparency, in the table below the guidelines set forth risk-based guides with regard to how many clicks should be necessary to get to the particular level of information, starting from either the page on which the recommendation is found, or an earlier page that the user will necessarily review before getting to the recommendation. For example, some software is designed so that the user is required to review an input page before getting to the recommendations. In that case, information regarding those inputs would be measured as against that input page.

- Click 0 means the information must be provided on the same page on which the inputs are recorded or the recommendation itself is provided
- Click 1 is a page that directly links to the input or recommendation page
- Click 2 is a page that directly links to a click 1 page

Transparency Element	Clicks
<p>1. Disclosure of the intended use. The software should explain clearly:</p> <ul style="list-style-type: none"> a. The role in the clinical practice that the software is to play. b. Limitations on the use of the software, and in particular warnings and contraindications about not using it too broadly. c. The intended user, as precisely as possible. d. The intended use setting, as precisely as possible. 	1
<p>2. Disclosure of information inputs</p> <ul style="list-style-type: none"> a. Patient specific data <ul style="list-style-type: none"> i. Data the user enters manually into the CDS program is presented back as a confirmation screen. ii. Data entered by access to an electronic record is: <ul style="list-style-type: none"> 1. Defined precisely by reference to the specific electronic record(s). 	1
<ul style="list-style-type: none"> 2. Presented to the user in one or more screens with regard to the data that the software deems relevant for its recommendation. 	2
<ul style="list-style-type: none"> b. Clinical intelligence content. Such content can come from many different sources, and those sources raise different issues. There are two different purposes served by the disclosures under this section: (1) allowing the user to understand and gauge the credibility, quality, completeness and appropriateness of the source of the content; and (2) allowing the user to review the original source material for the clinical content such that they can conduct their own review of the basis for the recommendation. Thus, in addition to providing the clinical rationale under part 4 below, the software must disclose the source or sources, as applicable, of its clinical intelligence as follows: <ul style="list-style-type: none"> i. Databases. Examples include commercially available databases such as MEDLINE®, and databases of clinical content curated specifically for the purpose of the software. <ul style="list-style-type: none"> 1. To be transparent, the software must provide enough information for the user to evaluate the database for completeness, currentness, quality, and relevance. 2. The database should be accessible to the user, but need not be within 2 clicks. ii. Specific, written, published materials. Examples include journal articles and medical society guidelines. <ul style="list-style-type: none"> 1. To be transparent, the software must provide enough information for the user to find the published materials. 2. Options for providing published materials. <ul style="list-style-type: none"> A. If the materials are provided through the software, they should be within 2 clicks of the clinical content page. B. If they are not provided, they should be 	1 unless 2 is specified

Transparency Element	Clicks
<p>reasonably accessible to the user. Further, the time required to access them must be considered under the time for reflection guidelines below.</p> <p>iii. Specific, written, unpublished materials. Examples include clinical guidelines unique and proprietary to the developer or user institution.</p> <ol style="list-style-type: none"> 1. The full text of such written materials needs to be provided within 2 clicks. 2. The text must be accompanied by a brief explanation of the source of the materials and how they were developed, including, for example, the authors and their qualifications. 3. The software also should disclose any financial interest in the recommendation held by the individual or institution, akin to the disclosures required in FDA’s Good Reprint Practices. <p>iv. Content that has not been summarized. Examples include software-embodied algorithms based on input received from a clinical expert or committee of experts, or from machine learning applied to raw data.</p> <ol style="list-style-type: none"> 1. If the source is people, such people need to be described by reference to their qualifications for the task. Again, the purpose of this transparency is to allow the user to assess the credibility and qualifications of the source. The summary should also include the particular topics on which these experts gave input. 2. If the source is machine learning or otherwise not human, the software must provide the user with appropriate metrics to give the user an understanding of the accuracy of the source. In language that is understandable to a clinician who does not have software development expertise, the program must contain a page that explains how the software has been validated, and provides meaningful quantitative assessments of the accuracy of the software. The software should provide a thorough explanation of the data sets used to feed and test the machine to provide important context and assurance to the clinician. The discussion should include any limitations or potential biases in the methods used to gather the data. Further, the software should explain if the machine learning model is designed to be adaptive, using both retrospective and prospective 	

Transparency Element	Clicks
<p>data inputs for incremental evolution of the decision algorithms for improving the accuracy. This page should be regularly updated as the software matures.</p>	
<p>Furthermore, depending on the role of the software and the risk associated with the intended use, the developer may also need to provide additional information. The following information is necessary if a decision informed by the recommendation of the software could lead to serious injury, permanent impairment or death for the patient.</p>	
<p>3. Disclosure of the output. The software must disclose enough to fully reflect the variability and uncertainties of medicine.</p> <p>a. Qualitative outputs from an expert system</p> <ul style="list-style-type: none"> i. The software must provide the full range of possible recommendations reasonably derived from the clinical inputs. ii. In addition, if the software developer would like, the software developer may identify a primary recommendation or suggest a ranking or confidence statements that need to be supported by the clinical rationale described below in section 4. <p>b. Quantitative outputs from an expert system</p> <ul style="list-style-type: none"> i. Using appropriate statistical tools, the software must communicate results, allowing the end user to determine the accuracy and reliability of the results presented. <p>c. Machine learning outputs</p> <ul style="list-style-type: none"> i. The software must provide the full range of possible recommendations reasonably derived from the clinical inputs. ii. In addition, if the software developer would like, the software developer may identify a primary recommendation or suggest a ranking or confidence statements that need to be supported by the clinical rationale described below in section 4. iii. If a recommendation is presented as a singular option, the recommendation must be accompanied by the associated confidence level, together with an assessment of both the likely benefits and risks. 	0
<p>4. Disclosure of the clinical rationale for the recommendations or rankings. In addition to describing the source of its clinical knowledge under part 2 above, the software must identify the particular clinical rationale that serves as a basis for its recommendations.</p> <p>a. Qualitative outputs from expert systems. The software must provide an explanation of the clinical reasoning such that the end user can understand the clinical rationale underlying the recommendation.</p>	1

Transparency Element	Clicks
<p>i. If based on specific clinical guidelines:</p> <ol style="list-style-type: none"> 1. An explanation of the clinical decision pathway through the guidelines. This may include decision trees, flow diagrams or other visual models. 2. A summary of any additional considerations factored into the recommendation beyond what was stated in the guidelines. 3. An explanation of any deviations from the published guidelines. 	
<p>ii. If based on specific published literature:</p> <ol style="list-style-type: none"> 1. Identification of the specific literature that underlies the recommendation. If multiple articles address the recommendation, the software should present up to three that support the recommendation, as well as up to three that present a contrary view (if any). 2. A summary of the clinical reasoning based on the factors identified in those articles. 	1
<ol style="list-style-type: none"> 3. Copies of the identified articles. If the developer elects not to provide copies directly, the time associated with the user being able to obtain copies must be factored into the time for reflection. 	2
<p>iii. If based on any other source, provide:</p> <ol style="list-style-type: none"> 1. An explanation of the clinical thought process, including any assumptions that underpin the recommendation to a degree typical of a second opinion. 2. A description of the clinical evidence, if any, that supports the recommendation, to the degree typical of a medical abstract, or an appropriate reference to published evidence. 	1
<p>b. Quantitative output from expert systems</p> <ol style="list-style-type: none"> i. The full and precise mathematical formulae used in the calculation. ii. A narrative explanation of the source or clinical basis for each step in the calculation, or a reference to the source of the formula. iii. For any such source, follow any relevant disclosure above in part 4.a. for qualitative information. 	1
<p>c. Machine learning outputs</p> <ol style="list-style-type: none"> i. Software built on machine learning needs to explain that it has spotted an association through machine learning, and state as precisely as it can the association. For example, with regard to a particular patient, the software might note that there is an association between patients with X, Y, and 	1

Transparency Element	Clicks
Z symptoms and improvements in those symptoms when taking Drug A.	

DRAFT

Appendix B

Competent Human Intervention Guidelines

If a given piece of software meets the Guidelines below (as well as transparency and time-to-reflect Guidelines defined in separate appendices), under the CDS Coalition's guidelines, the healthcare professional would be able to independently review the basis for the recommendation. These Guidelines are not intended as requirements, but rather, as one approach for determining whether the healthcare professional is able, by virtue of education and experience, to independently review the recommendations of the software. The Coalition recognizes that there may be other means to demonstrate competent human intervention.

These Guidelines recognize the need of a healthcare professional to exercise his/her clinical expertise and judgment within his/her applicable practice area and utilizes traditional licensure classifications. To enable the end user to independently review the basis of the recommendations, the software should be intended to be used to support a decision the user is qualified to make without the use of the software. Consequently, software intended to be used to extend a user's decision making ability beyond his/her qualifications, license, certification or practice area would not ensure that the user is able to independently review the basis for the recommendations.

I. Competent Human Intervention Definition

Competent human intervention means the intended user must possess any necessary knowledge, skill, judgment, training, experience and license to understand the information presented, appreciate the consequences of acting or not acting on that information and make a decision or take action based upon the information presented.

In determining the level of knowledge, skill, judgment, training, experience and license necessary to qualify as competent human intervention, we look at (1) the abilities and licensure of the intended user and (2) the nature and complexity of the clinical decision-making process.

II. Competent Human Categories

We outline three categories of intended healthcare professional users below.

Healthcare Practitioner is a non-physician professional who provides clinical or other healthcare service to individuals and has been trained and certified or licensed in her profession. Healthcare professionals represent a variety of degrees and specialties, including but not limited to registered nurses, nurse practitioners, physician assistants, radiology technicians, pharmacists, midwives, dietitians, therapists, psychologists, chiropractors, phlebotomists, physical therapists, respiratory therapists, occupational therapists, audiologists, speech pathologists, optometrists, emergency medical technicians, paramedics, medical laboratory scientists, and medical prosthetic technicians.

General Practitioner is a person who is legally qualified and validly licensed to practice medicine as a primary care physician. Primary care means providing continuing and comprehensive medical care to individuals and families, not limited by the cause, organ system or diagnosis and includes the specialties of internal medicine, family medicine, general practice and pediatrics.

Specialist is a physician who is board certified in a particular specialty or sub-specialty by the American Board of Medical Specialties, other than internal medicine, family medicine, general practice or pediatrics.

III. Clinical Decision Support Categories

We outline the three corresponding categories of decision support based upon the nature and complexity of the clinical decision-making process.

Care Management CDS is defined as tools that support routine clinical decisions related to the execution or management of a care plan that requires basic clinical knowledge, simple triage decisions or other determinations of when the intervention of a general practitioner or specialist is recommended. Care Management CDS does not involve diagnostic or treatment decision-making.

General Practitioner CDS is defined as tools that assist general practitioners in making a diagnosis, selecting a treatment or therapy, developing a medical treatment plan, or help general practitioners determine whether the intervention of a specialist is recommended.

Specialist CDS is defined as tools that (1) address those diagnostic and treatment decisions involving diseases or conditions, types of patients or methods of treatment that required specialized knowledge gained through advanced medical training or (2) identifies when a different specialist or sub-specialist is recommended.

IV. Competencies and Decision Type Combinations That Allow For Independent Review

Combining those factors, the Xs indicate where independent review of the basis of the recommendation by the software is possible because there is competent human intervention.

	Healthcare Practitioner	Generalist	Specialist
Care Management	X	X	X
General Practitioner		X	X
Specialist			X

V. Intended Use Requirements in Labeling and Promotional Activities

The CDS software manufacturer must clearly identify the intended user, including the necessary qualifications of the user, in its description in marketing materials based upon the nature of the clinical decision supported by the software. In addition, manufacturer's promotional practices must be consistent with the notion that the CDS software is intended to be used by applicable intended users. This means that the sales and marketing activities are targeted at only the applicable intended user category. By way of example, the tradeshow attended, magazine and media outlets in which the product is advertised and invitees for meetings or presentations should be consistent with and reflect the applicable intended user for such product.

DRAFT

Appendix C

Time to Reflect Guidelines

If a given piece of software meets the Guidelines below (as well as transparency and competent human intervention guidelines defined in separate appendices), under the CDS Coalition's guidelines, the healthcare professional would be able to independently review the basis for the recommendation. These Guidelines are not intended as requirements, but rather, as one approach for determining whether the healthcare professional is able, by virtue of having enough time to reflect, to independently review the recommendations of the software. The Coalition recognizes that there may be other means to demonstrate adequate time to reflect.

To enable the end user to independently review the basis of the recommendations, the software must be intended to be used in a setting that allows the user sufficient time to evaluate and consider the CDS output before making a decision or taking an action. Consequently, software intended to be used in a clinical setting where the acuity of the condition combined with the complexity of the decision making does not allow the user the amount of time necessary to reflect on the output would not ensure that the user is able to independently review the basis for the recommendations.

I. "Time to Reflect" Definition

Time to reflect means the intended user has the amount of time necessary to evaluate the basis of the recommendation of the software before taking action. In determining whether there is sufficient time to reflect, we examine two factors: (1) the acuity of the condition and (2) the complexity of the clinical decision making process.

II. Acuity of the Condition

The acuity level is the time that can lapse without material risk to the patient before that patient receives medical care. Acuity levels can be classified as follows:

Immediate. There is no time for reflection. This acuity level refers to conditions that could result in death or serious irreversible injury or where there is imminent risk of deterioration to such condition if there is any delay in treatment or intervention.

Urgent. Treatment is time-sensitive. This acuity level refers to conditions that could result in death or serious irreversible injury or where there is risk of deterioration to such condition or requiring emergency care if treatment or intervention is not commenced promptly and without undue delay. Urgent conditions generally require treatment within a fifteen minute to two hour timeframe.

Non-Urgent. There is sufficient time for reflection. This acuity level refers to conditions for which investigation or interventions can be delayed or referred to another healthcare provider or health care system without risk of death, serious irreversible injury or deterioration requiring emergency care.

III. Complexity of the Clinical Decision Making Process

The complexity of the clinical decision making refers to the amount of time the user needs to understand the clinical rationale underlying the decision and evaluate it for accuracy.

Simple. The clinical decision making process or clinical rationale for such decision is so well established, or the calculation is so easily understood, that the software output can be evaluated nearly instantaneously without review of detailed supporting materials.

Moderate. The clinical decision making process or clinical rationale for such decisions is such that the software output can be understood with minimal consideration or quick reference to detailed supporting materials.

Complex. The clinical decision making process or clinical rationale for such decision is so complicated, multifaceted or multifactorial in nature that the software output requires an extended time to be evaluated, or needs to be interpreted in connection with the detailed supporting materials.

In connection with this factor, the software developer must also consider how quickly, with the software, the user can consult the underlying guidelines that serve as a basis for the recommendation by the software.

IV. Acuity and Complexity Combinations That Allow For Independent Review

Combining those factors, the Xs indicate where independent review of the basis of the recommendation by the software is possible because there is adequate time for reflection.

	Immediate	Urgent	Non-Urgent
Simple	X	X	X
Moderate		X	X
Complex			X

V. Intended Use Requirements in Labeling and Promotional Activities

The CDS software manufacturer should consider the complexity of the clinical decision making process and clearly identify the acuity level of the condition for which the software is intended, or not intended, to be used. In addition, manufacturer's promotional practices should be consistent with the notion that the CDS software is intended to be used within a certain time for reflection given the complexity of the decision. This means that the sales and marketing

activities describe the typical medical care setting where the CDS will likely be used, and the complexity of the decision making.

DRAFT

Appendix D

Background: Understanding the Statute

As explained above, these guidelines are not meant to be an interpretation of the new statute. But at the same time, the goal of the coalition is to ensure that compliance with these guidelines also means that the software is not FDA-regulated. Thus, an understanding of what the new statute requires is foundational to the development of these guidelines. To provide that understanding, in the table below in the left-hand side we include the new language from the statute, and on the right-hand side provide a running commentary explaining the significance of the language.

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
(1) The term device, as defined in section 201(h), shall not include a software function that is intended—	Paragraph (1) of subsection (o) lists a variety of software functions that are excluded from the FDA definition of a medical device. Because this document focuses on clinical decision support software, we focus our analysis on subparagraph (E) that includes CDS.
(E) unless the function is intended to acquire, process, or analyze a medical image or a signal from an in vitro diagnostic device or a pattern or signal from a signal acquisition system,	<p>The introduction to subparagraph (E) is a bit jarring to read, because it starts with an exclusion. Think of it as a double negative. These are software functions that are <u>not</u> excluded from the FDA definition of a medical device.</p> <p>We have used formatting to make this introductory clause easier to comprehend.</p> <p>For decades, FDA has regulated software that analyzes medical images and certain signals from electronic medical devices.</p> <p>The term “medical image” is fairly clear. There are a wide variety of file formats routinely used for medical images, but the key is that the file represents an image used for a medical purpose. So conceivably, this could include a simple JPEG, if the photograph is, for example, of a mole and the software is being used to analyze that photograph of a mole for assessing the possibility of melanoma.</p> <p>The term signal probably refers to an electronic signal, and such signals typically are defined as</p>

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	<p>an electric current that represents information. There are two basic types of electrical signals: analog and digital. In analog signals, some continuously variable aspect of the electrical current represents the information.</p> <p>In vitro diagnostic device is a well-defined term in FDA regulation.</p> <p>A signal acquisition system is a commonly used phrase in electrical engineering.</p> <p>The essence of this exclusion seems to be electronic information that is not comprised of either words or numbers understandable by humans, but rather visual or graphic information requiring interpretation.</p>
for the purpose of—	<p>This phrase then marks the introduction of the types of software that are in fact excluded from the FDA medical device definition. Excluded software must meet all three of the following criteria (notice the word “and” after the second criterion.)</p>
(i) displaying, analyzing, or printing medical information about a patient or other medical information (such as peer-reviewed clinical studies and clinical practice guidelines);	<p>This first criterion is very broad. Indeed, it can be boiled down for our purposes to software that “analyzes” “medical information.”</p> <p>Given our focus on clinical decision support software, we will not analyze the “display” or “printing” functionality portions of this section.</p> <p>Clinical decision support typically involves taking generalized knowledge and applying it to a specific patient. Congress drafted this language to include the analysis both of the generalized medical information, as well as its application to the particular patient.</p> <p>It is noteworthy that the parenthetical is merely an example, and not words of limitation. The information does not need to be from either peer-reviewed clinical studies or medical</p>

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	<p>practice guidelines. The key statutory phrase is the wide open description of “other medical information.” Other than the information being medical, there is no limitation on the scope of that language.</p> <p>Thus the generalized medical information can come from a wide variety of sources, ranging from broadly accepted medical information to brand-new medical insights that have not been affirmed in any way by the medical establishment.</p> <p>Notice that the example “clinical practice guidelines” is not limited to consensus guidelines or the guidelines of established medical societies. Thus, even that example suggests that any guidelines, even if they are only specific to a particular medical provider, qualify.</p> <p>Beyond that, the open-ended reference to “other medical information” would include such sources as individual medical opinions of healthcare professionals, to new medical insights derived by analyzing large troves of patient data. There is simply no limit in the statute on what qualifies as medical information.</p>
(ii) supporting or providing recommendations to a health care professional about prevention, diagnosis, or treatment of a disease or condition; and	<p>This subparagraph squarely focuses on the function of clinical decision support software, in that the core purpose of such software is to provide or support recommendations for clinical decisions.</p> <p>This subparagraph also makes it clear that the threshold for the applicability of this entire provision is that the software must be directed toward recommendations for healthcare professionals, as opposed to consumers and patients.</p> <p>The phrase “about prevention, diagnosis, or</p>

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	treatment of a disease or condition” focuses the section on software that might otherwise qualify as a medical device, as that language is taken out of the medical device definition.
<p>(iii) enabling such health care professional to independently review the basis for such recommendations that such software presents so that it is not the intent that such health care professional rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient.</p>	<p>This is the meat of the subparagraph, and the toughest part of the test.</p> <p>The word “enabling” is one of the most important words in this particular subparagraph. The software has to be designed and marketed so as to “enable” healthcare professionals to independently review the basis for the recommendations. So we have to ask ourselves what it takes for healthcare professionals to be so enabled.</p> <p>For healthcare professionals to be able to review the basis for the recommendations, implicitly the software must meet at least three conditions:</p> <ol style="list-style-type: none"> 1. The healthcare professional has to be qualified to review the basis for the recommendation. In other words, it is not enough that someone is simply generically a healthcare professional. More specifically, the particular healthcare professional has to be qualified to conduct an independent review of the particular recommendation. So a psychiatrist, even though she is a healthcare professional, may not be qualified to review a recommendation by the software with regard to the treatment of cancer. The particular recommendation at issue has to be within the scope of competence of the healthcare professional targeted by the software. 2. The healthcare professional must have the time necessary to conduct the independent review. If the software focuses on a time critical function such

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	<p>as determining whether to shock a patient with an external defibrillator in the face of an urgent arrhythmia, to meet this criterion the basis for the recommendation would need to be very simple and susceptible to being reviewed in the amount of time available.</p> <p>3. The healthcare professional must be given access to the medical information – both patient specific and generalized – on which the recommendation is based so that the professional can review that basis. We refer to this concept as transparency, meaning the ability to see through the software to the underlying information.</p> <p>A substantial portion of these industry guidelines focuses on amplifying the meaning of those three factors.</p> <p>But there is a second, in a sense parallel, requirement in the section. That second requirement is that it is not the intent of the software developer for the user to “rely primarily on any of such recommendations.” In some cases, this seems like a redundant statement of the test, in that if the software gives healthcare professionals complete latitude to independently review the basis, then the user is not intended to “rely primarily on such recommendations.” However, statutory language is never dismissed as simply redundant if there is an additional meaning which it can be reasonably given. And in this case, there is another factor that should be considered.</p> <p>That factor is whether there is information beyond what the software provides on which the healthcare professional can base her therapeutic or diagnostic decision. The</p>

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	<p>existence of such additional information makes the healthcare professional less reliant on the recommendation from the software.</p> <p>When making a diagnostic or therapeutic decision, healthcare professionals typically consider a wide range of information including direct physical examination of the patient, a medical history that the healthcare professional might take during an interview of the patient, laboratory tests that the healthcare professional may review directly and other such information. The existence of such information, so long as it is available to the healthcare professional, makes the healthcare professional less reliant on the software's recommendation.</p> <p>The bottom line is that there are two ways to make it so that a healthcare professional is not primarily reliant on the software's recommendation:</p> <ol style="list-style-type: none"> 1. Access to all of the information on which the software bases its recommendation, 2. Access to other information outside of that considered by the software. <p>Both of those, considered together, will make a healthcare professional who is qualified to make a decision, and who has adequate time to make that decision, less reliant on the software's recommendation.</p> <p>Thus, these need to be considered together, and to some extent complement each other. If software is absolutely transparent in that the healthcare professional has very quick and easy access to all of the information on which the software bases its recommendation, that by itself will demonstrate that the healthcare professional is not intended to primarily rely on the recommendation. But if that transparency is</p>

Language of the new subparagraph 520(o)(1)(E) of the Federal Food, Drug And Cosmetic Act	Commentary
	<p>not so complete or so easily accomplished, then the developer may consider whether other complementary information that the healthcare professional can reasonably be expected to have outside of the software, considered together, means that the healthcare professional is not intended to rely primarily on the software's recommendation.</p> <p>At the same time, it must be recognized that the sentence construction makes it clear that the statute principally looks to the ability of the healthcare professional to independently review the basis of the recommendation in determining that the healthcare professional is not primarily reliant on the software. So while information the software does not consider – information that is directly available to the healthcare professional – can help accomplish the ultimate objective of avoiding primary reliance, it cannot compensate for software that does not allow for the review of the basis for its recommendation. The statute requires that the healthcare professional be able to review the basis of the recommendation. Outside information simply makes it easier to attain the ultimate goal of avoiding reliance.</p>