



**DukeHealth**



**Duke**  
AI HEALTH

# Algorithm-Based Clinical Decision Support (ABCDS) Oversight

**A framework for the governance and evaluation of algorithms to be deployed at Duke Health**

**NCRAF**

April 25, 2023

---

Amanda Parrish, PhD

Nicoleta J Economou, PhD

# Promise of Artificial Intelligence/Machine Learning in Health Care



# AI/ML Fails in the Field

June 21, 2021

## The Epic Sepsis Model Falls Short—The Importance of External Validation

Anand R. Habib, MD, MPhil<sup>1,2</sup>; Anthony L. Lin, MD<sup>1</sup>; Richard W. Grant, MD, MPH<sup>3,4</sup>

» Author Affiliations

JAMA Intern Med. Published online June 21, 2021. doi:10.1001/jamainternmed.2021.3333

Related Articles

Sepsis accounts for nearly 1 million hospitalizations annually and is a major contributor to health care expenditures, and in-hospital mortality (ranging from 12.5%-15%).<sup>1</sup> Epic care teams to promptly implement goal-directed therapy to mitigate clinical deterioration. In JAMA Internal Medicine, Wong et al<sup>2</sup> report on their external validation of the Epic Sepsis Model (ESM) tool available within the Epic electronic health record that is designed to generate alerts for clinicians that patients may be developing sepsis. Based on their examination of 3,000 patients at the University of Michigan (Ann Arbor) between December 2018 and October 2019, the ESM had a sensitivity of 33%, specificity of 83%, positive predictive value of 12%, and negative predictive value of 95%, with an area under the curve of 0.63 (95% CI, 0.62-0.64). This falls short of the 0.76 to 0.83 that was jointly reported by Epic and University of Colorado Health. In 18% of all patients, the ESM did not detect sepsis in 67% of patients with sepsis.



Science

RESEARCH

### RESEARCH ARTICLE

ECONOMICS

## Dissecting racial bias in an algorithm used to manage the health of populations

Ziad Obermeyer<sup>1,2\*</sup>, Brian Powers<sup>3</sup>, Christine Vogel<sup>4</sup>, Sendhil Mullainathan<sup>5,6\*</sup>

Health systems rely on commercial prediction algorithms to identify and help patients with complex health needs. We show that a widely used algorithm, typical of this industry-wide approach and affecting millions of patients, exhibits significant racial bias: At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled illnesses.

Remediating this disparity would increase the percentage of Black patients receiving additional help from 17.7% to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness...

There is growing concern that algorithms may reproduce racial and gender disparities via the people building them or through the data used to train them (1-3). Empirical work is increasingly lending support to these concerns. For example, job search ads for highly paid positions are less likely to be presented to women (4), searches for distinctively Black-sounding names are more likely to trigger ads for arrest records (5), and image searches for professions such as CEO produce fewer images of women (6). Facial recognition systems increasingly used in law enforcement perform worse on people

researcher-created algorithms (10-13). Without an algorithm's training data, objective function, and prediction methodology, we can only guess as to the actual mechanisms for the important algorithmic disparities that arise. In this study, we exploit a rich dataset that provides insight into a live, scaled algorithm deployed nationwide today. It is one of the largest and most typical examples of a class of commercial risk-prediction tools that, by industry estimates, are applied to roughly 200 million people in the United States each year. Large health systems and payers rely on this algorithm to target patients for "high-risk

that rely on past data to build a predictor of future health care needs.

Our dataset describes one such typical algorithm. It contains both the algorithm's predictions as well as the data needed to understand its inner workings: that is, the underlying ingredients used to form the algorithm (data, objective function, etc.) and links to a rich set of outcome data. Because we have the inputs, outputs, and eventual outcomes, our data allow us a rare opportunity to quantify racial disparities in algorithms and isolate the mechanisms by which they arise. It should be emphasized that this algorithm is not unique. Rather, it is emblematic of a generalized approach in the health sector of for-profit and government-funded organizations beyond what our algorithm. First, this algorithm is used in other sectors: The predicted risk of some future outcome (in our case, health care needs) is widely used to target policy interventions under the assumption that the treatment effect is monotonic in that risk, and the methods used to build the algorithm are standard. Mechanisms of bias uncovered in this study likely operate elsewhere. Second, even beyond our particular finding, we hope that this exercise illustrates the importance, and the large opportunity, of studying algorithmic bias in health care, not just as a model system but also in its own right. By any standard—e.g., number of lives affected, life-and-death consequences of the decision—health is one of the most important and wide-

**"At a given risk score, Black patients are considerably sicker than White patients, as evidenced by signs of uncontrolled illnesses. Remediating this disparity would increase the percentage of Black patients receiving additional help from 17.7% to 46.5%. The bias arises because the algorithm predicts health care costs rather than illness..."**

Healthcare IT News

Global Edition Analytics

## Research suggests Epic Sepsis Model is lacking in predictive

Study in JAMA Internal Medicine finds that the model can identify two-thirds of sepsis patients and generate false alarms.

2021 | 12:44 PM





# “Wild West” of Algorithms





# Regulatory Landscape Changing Rapidly




**HHS.gov**  
**Civil Rights**

## Fact Sheet: Nondiscrimination in Health Programs and Activities Proposed Rule Section 1557 of the Affordable Care Act


The Department of Health and Human Services (HHS) has issued a proposed rule to advance health equity and reduce disparities in health care. The proposed rule, *Nondiscrimination in Health Programs and Activities*, revises the implementing regulation for [Section 1557 of the Affordable Care Act](#) (ACA), and proposes robust provisions that will be more effective in protecting people from discrimination.

Section 1557 of the ACA prohibits discrimination on the basis of race, color, national origin, sex, age, or disability in certain health programs or activities and is one of the government's most powerful tools to ensure nondiscriminatory access to health care. In addition to proposing revisions to the Section 1557 implementing regulation, this rulemaking also includes proposed revisions to nondiscrimination



# Local Government Taking Action

State of California Department of Justice



**ROB BONTA**  
*Attorney General*






Search

Translate Website | Traducir Sitio Web

HOMEABOUTMEDIACAREERSREGULATIONSRESOURCESPROGRAMSAPPOINTMENTSCONTACT

## Attorney General Bonta Launches Inquiry into Racial and Ethnic Bias in Healthcare Algorithms

Press Release / Attorney General Bonta Launches Inquiry into Racial and Ethn...



Wednesday, August 31, 2022

Contact: (916) 210-6000, [agpressooffice@doj.ca.gov](mailto:agpressooffice@doj.ca.gov)

*Sends letters to 30 hospital CEOs across the state requesting information regarding the use of commercial healthcare decision-making tools*

**OAKLAND** – California Attorney General Rob Bonta today sent letters to hospital CEOs across the state requesting information about how healthcare facilities and other providers are identifying and addressing racial and ethnic disparities in commercial decision-making tools. The request for information is the first step in a DOJ inquiry into whether commercial healthcare algorithms – types of software used by healthcare providers to make decisions that affect access to healthcare for California patients – have discriminatory impacts based on race and ethnicity.

All information provided to our office will be treated as confidential in accordance with California Government Code section 11180 et seq. Please provide the requested information and documents to Deputy Attorney General Anna Rich at [Anna.Rich@doj.ca.gov](mailto:Anna.Rich@doj.ca.gov), or 1515 Clay St., 20<sup>th</sup> Floor, Oakland, CA, 94612, by October 15, 2022. We thank you in advance for your cooperation.

Sincerely,



ROB BONTA  
Attorney General



# “Our House” circa 2018

- Total 6 models
  - Sepsis
  - Early Warning Score
  - First Admission
  - Readmission
  - Falls
  - Pressure Injury



# “Our House” Today


- Over 40 registered tools





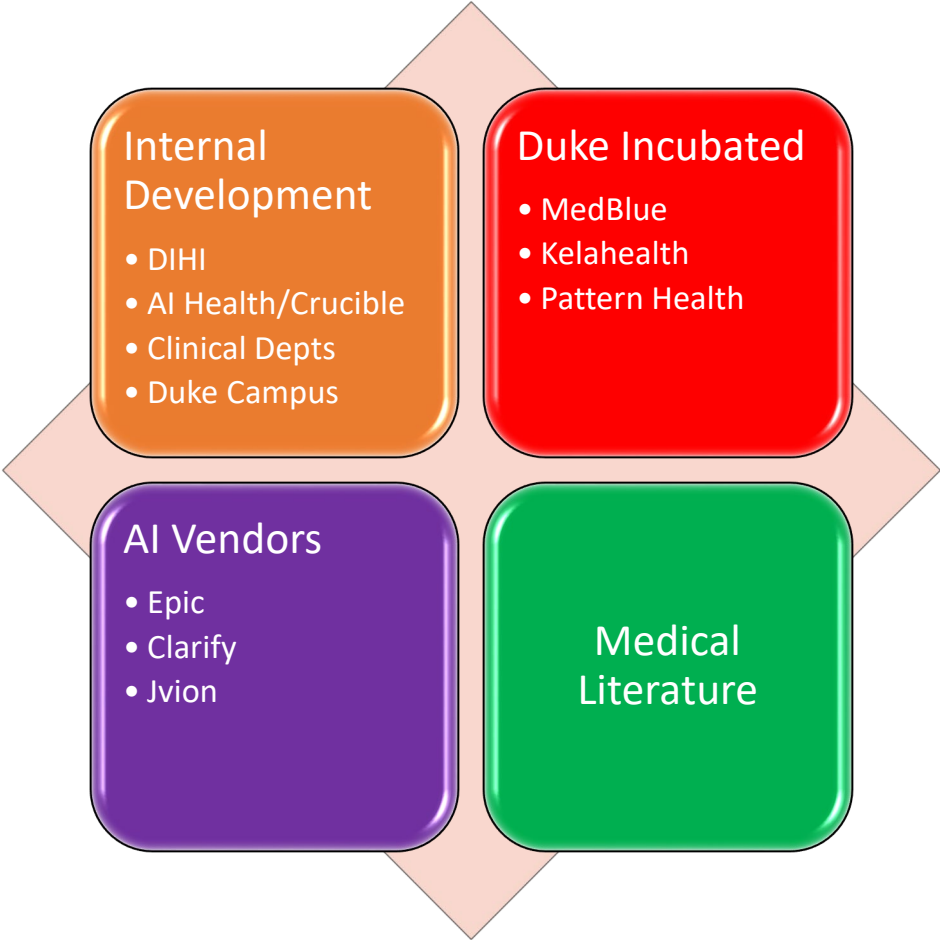
# Complex Environment

- Different skills
- Different knowledge bases
- Different resources available
- Different make up of project teams

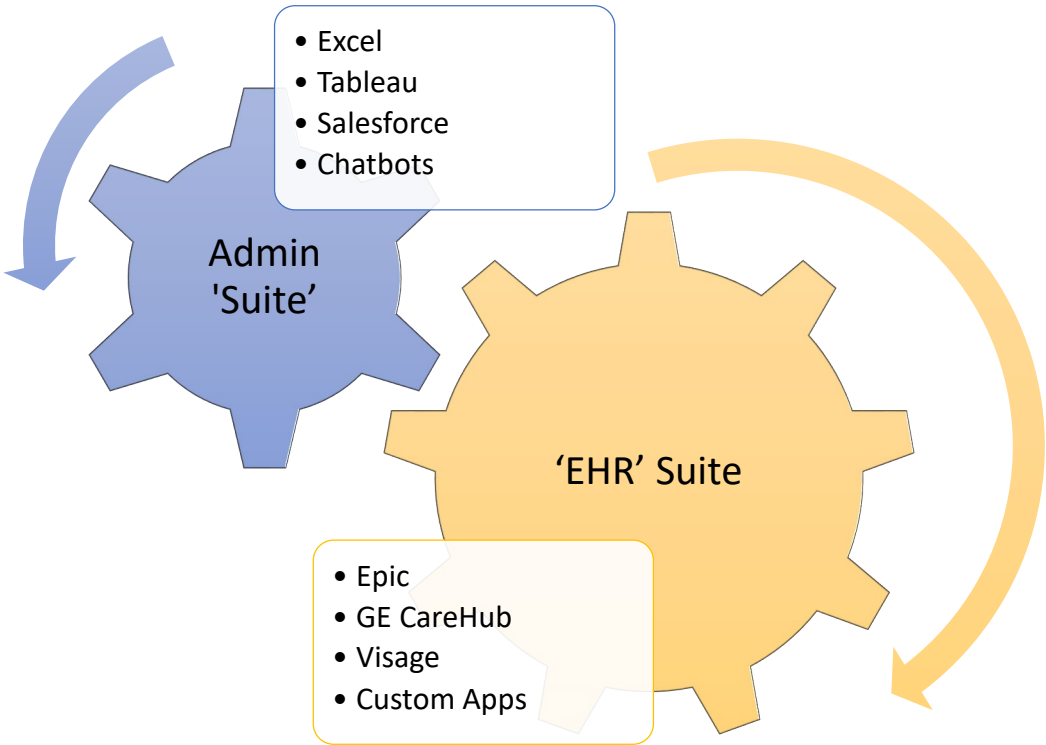


# Complex Environment

## Sources of Models



## Deployment Mechanisms



## Target Audience





# The Formation of the ABCDS Oversight Committee

In recognition of this changing landscape the Duke Health Chancellor and the Dean of the School of Medicine charged Duke Health leadership to form an oversight framework.

# Mission Statement

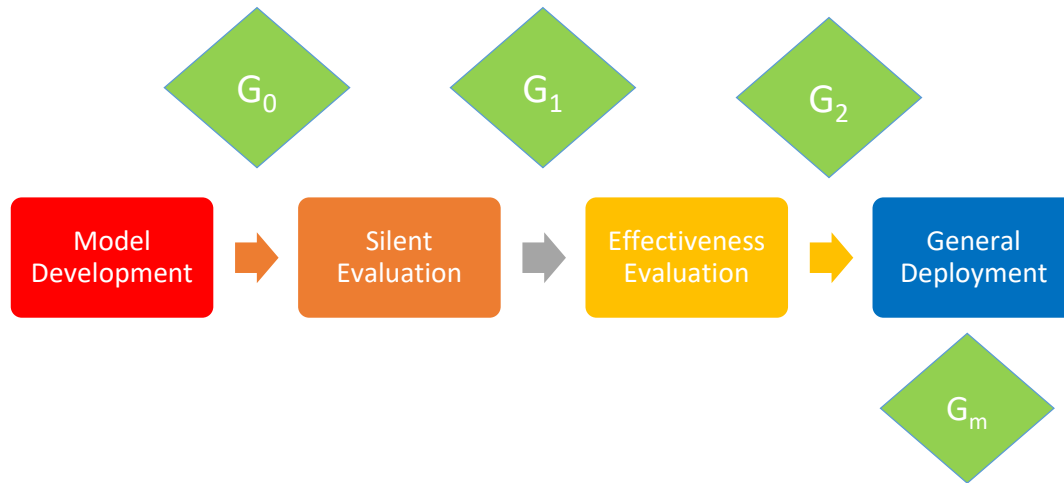
*Out of our primary focus on patient safety and high-quality care, our mission is to guide algorithm-based clinical decision support (ABCDS) tools through their lifecycle by providing governance, evaluation, and monitoring.*



Duke  
AI HEALTH



# ABCDS Lifecycle & Our Framework



*What are the clinical outcome and performance metrics?*

*How has the model been evaluated?*

*Who is the Clinical Owner?*

*Who will cover maintenance costs in production?*

*Will this ABCDS tool be used outside of Duke Health?*

*Is this a standard of care model?*

*How will the model be used in the clinic and how is it integrated with the workflow?*

...

‘Just-in-time’ Check-Points (**G**ates) Help Model Owners Get Ready for What’s Ahead

# Implementing Quality & Ethics with Our Framework

## Ethical Principles and Requirements

There are six general *ethical principles*<sup>3</sup> that any AI system must preserve and protect based on fundamental rights as enshrined in the Charter of Fundamental Rights of the European Union (EU Charter), and in relevant international human rights law:

1. **Respect for Human Agency:** human beings must be respected to make their own decisions and carry out their own actions. Respect for human agency encapsulates three more specific principles, which define fundamental human rights: **autonomy, dignity** and **freedom**.
2. **Privacy and Data governance:** people have the right to privacy and data protection and these should be respected at all times;
3. **Fairness:** people should be given equal rights and opportunities and should not be advantaged or disadvantaged undeservedly;
4. **Individual, Social and Environmental Well-being:** AI systems should contribute to, and not harm, individual, social and environmental wellbeing;
5. **Transparency:** the purpose, inputs and operations of AI programs should be knowable and understandable to its stakeholders;
6. **Accountability and Oversight:** humans should be able to understand, supervise and control the design and operation of AI based systems, and the actors involved in their development or operation should take responsibility for the way that these applications function and for the resulting consequences.




Transparency & Accountability

Impact & Safety

Fairness & Equity

Usability & Adoption

Regulatory Compliance



**Fig. 4.** Characteristics of trustworthy AI systems. Valid & Reliable is a necessary condition of trustworthiness and is shown as the base for other trustworthiness characteristics. Accountable & Transparent is shown as a vertical box because it relates to all other characteristics.

# People



Duke  
AI HEALTH




# People: ABCDS Oversight Committee

Co-Chairs:



M Pencina



E Poon

Director:



N Economou

ABCDS Oversight  
Committee


Additional Committee Members:




S Balu



M Cary



M Lipkin




K Lytle

ABCDS Regulatory  
Subcommittee

ABCDS Evaluation  
Subcommittee

ABCDS Implementation and  
Monitoring  
Subcommittee

Co-Chairs:



A Parrish



S Elengold




S Ellison

Co-Chairs:




B Goldstein



E Jelovsek

Co-Chairs:



A Bedoya



C O'Brien

Ops Team:



S Bessias



N Walden



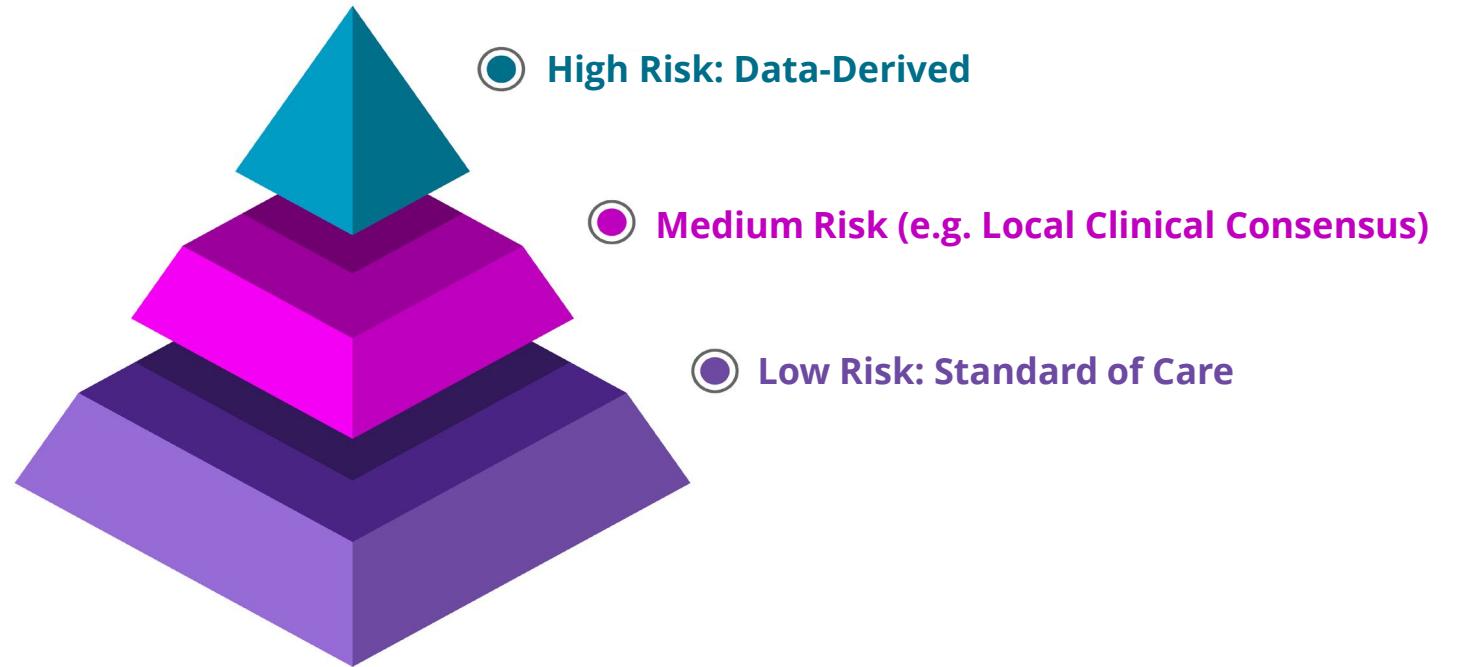
Duke  
AI HEALTH

# Process

# Scope of ABCDS Oversight Framework

**ABCDS Tool = Algorithm(s) + Interface Algorithms Are Presented In**

All electronic algorithms that could impact patient care at Duke Health fall within the scope of the ABCDS Oversight Committee and must undergo registration

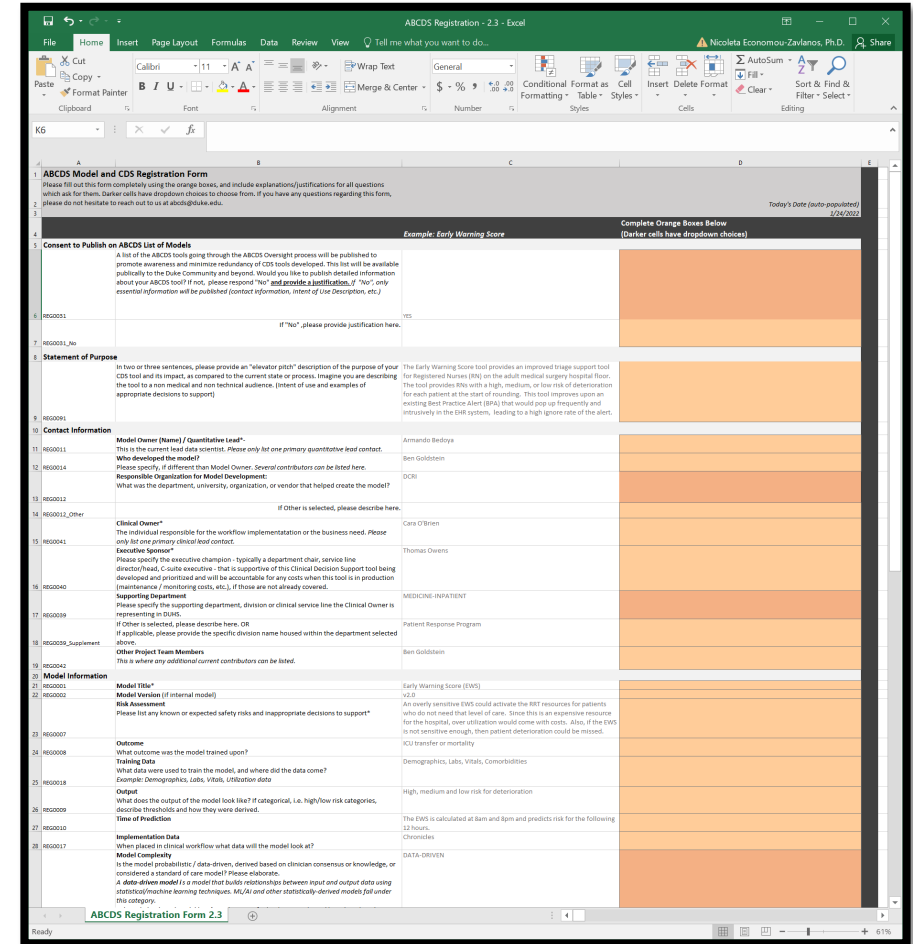




# The ABCDS Registration Form

## Information Requested

- ✓ Consent to publish
- ✓ Purpose
- ✓ Contact Information
- ✓ Model Information
- ✓ Use Case Information
- ✓ Regulatory Information
- ✓ (ONLY Standard of Care – Literature, society material)



ABCDS Registration - 2.3 - Excel

File Home Insert Page Layout Formulas Data Review View Tell me what you want to do...

ABCDS Model and CDS Registration Form

Please fill out this form completely using the orange boxes, and include explanations/justifications for all questions which ask for them. Darker cells have dropdown choices to choose from. If you have any questions regarding this form, please do not hesitate to reach out to us at [abcds@duke.edu](mailto:abcds@duke.edu).

Today's Date (auto-populated): 3/24/2023

Complete Orange Boxes Below (Darker cells have dropdown choices)

**Consent to Publish on ABCDS List of Models**

A list of the ABCDS tools going through the ABCDS Oversight process will be published to promote awareness and minimize redundancy of CDS tools developed. This list will be available publicly to the Duke Community and beyond. Would you like to publish detailed information about your ABCDS tool? If not, please respond "No" and provide a justification. If "No", only essential information will be published (contact information, intent of use description, etc.)

8. RECD001: If "No", please provide justification here. YES

9. RECD001\_Just: If "No", please provide justification here.

**Statement of Purpose**

In two or three sentences, please provide an "elevator pitch" description of the purpose of your CDS tool and its impact, as compared to the current state or process. Imagine you are describing the tool to a non-medical and non-technical audience. (Intent of use and examples of appropriate decisions to support)

The Early Warning Score tool provides an improved triage support tool for Registered Nurses (RN) on the adult medical surgery hospital floor. The tool provides RNs with a high, medium, or low risk of deterioration for each patient at the start of rounding. This tool improves upon an existing Best Practice Alert (BPA) that would pop up frequently and intrusively in the EHR system, leading to a high ignore rate of the alert.

10. RECD001: Model Owner (Name) / Quantitative Lead? Armando Bedoya

11. RECD001: Who developed the model? Ben Goldstein

12. RECD001: Responsible Organization for Model Development? DCRI

13. RECD001: What was the department, university, organization, or vendor that helped create the model? DCRI

14. RECD001\_Other: If Other is selected, please describe here.

15. RECD001: Clinical Owner? Care O'Brien

16. RECD001: Executive Sponsor? Thomas Owens

17. RECD001: Supporting Department? MEDICINE-INTENSIVE

18. RECD001: Please specify the supporting department, division or clinical service line the Clinical Owner is representing in DCRI. Patient Response Program

19. RECD001: If Other is selected, please describe here. OR

20. RECD001: If applicable, please provide the specific division name housed within the department selected above. Ben Goldstein

21. RECD001: Other Project Team Members

22. RECD001: This is where any additional current contributors can be listed.

**Model Information**

23. RECD001: Model Title? Early Warning Score (EWS)

24. RECD001: Model Version (if internal model)? V2.0

25. RECD001: Risk Assessment? An early sensitive EWS could activate the RRT resources for patients who do not need that level of care. Since this is an expensive resource for the hospital, over utilization would come with costs. Also, if the EWS is not sensitive enough, then patient deterioration could be missed.

26. RECD001: Outcome? ICU transfer or mortality

27. RECD001: What outcome was the model trained upon? Demographics, Labs, Vitals, Comorbidities

28. RECD001: Training Data? What data were used to train the model, and where did the data come? Example: Demographics, Labs, Vitals, Utilization data

29. RECD001: Output? High, medium and low risk for deterioration

30. RECD001: What does the output of the model look like? If categorical, i.e. high/low risk categories, describe thresholds and how they were derived. The EWS is calculated at 8am and 4pm and predicts risk for the following 12 hours.

31. RECD001: Time of Prediction? Chronicles

32. RECD001: Implementation Data? When placed in clinical workflow what data will the model look at?

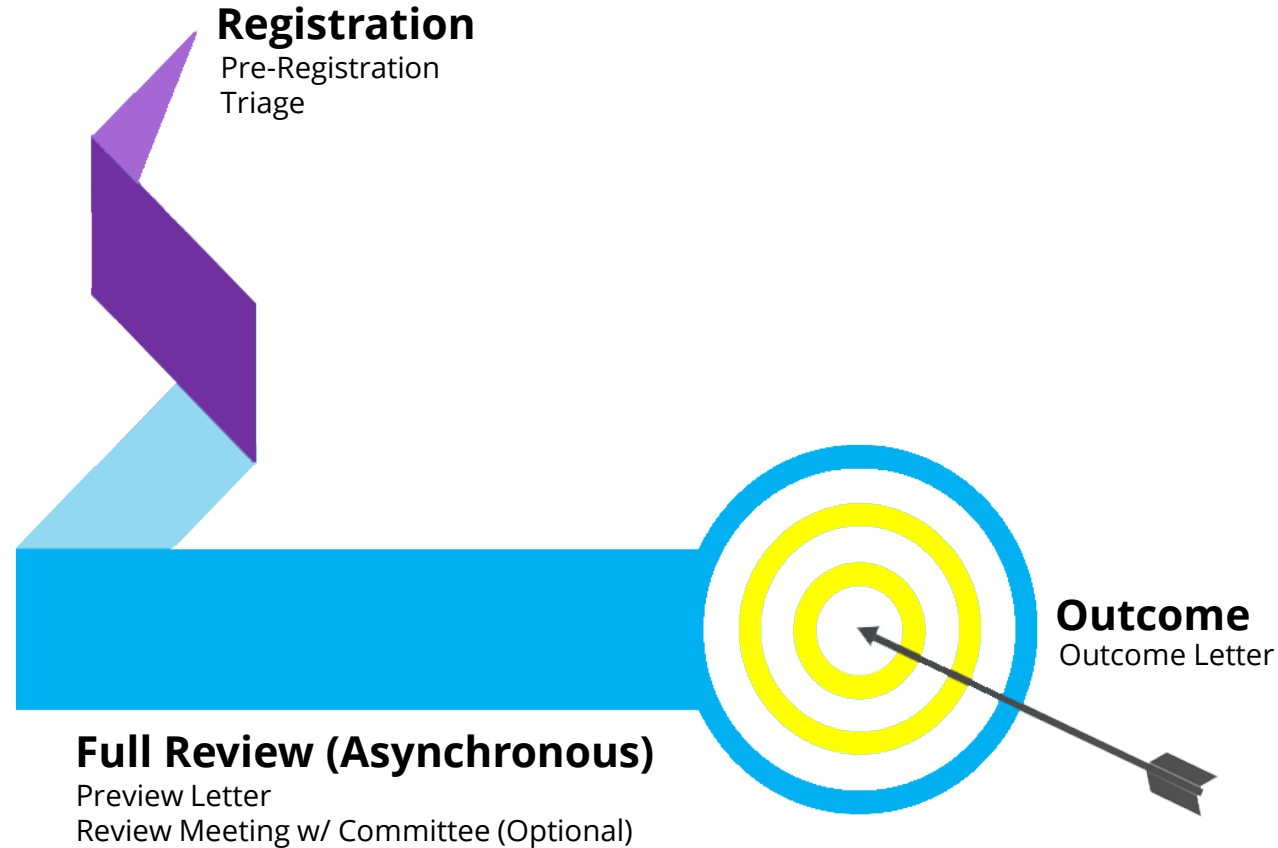
33. RECD001: Model Complexity? Is the model probabilistic / data-driven, derived based on clinician consensus or knowledge, or considered a standard of care model? Please elaborate. DATA-DRIVEN

34. RECD001: A data-driven model is a model that builds relationships between input and output data using statistical/machine learning techniques. ML/AI and other statistically-derived models fall under this category.

ABCDS Registration Form 2.3

Ready

# What to Expect: ABCDS Checkpoint Review

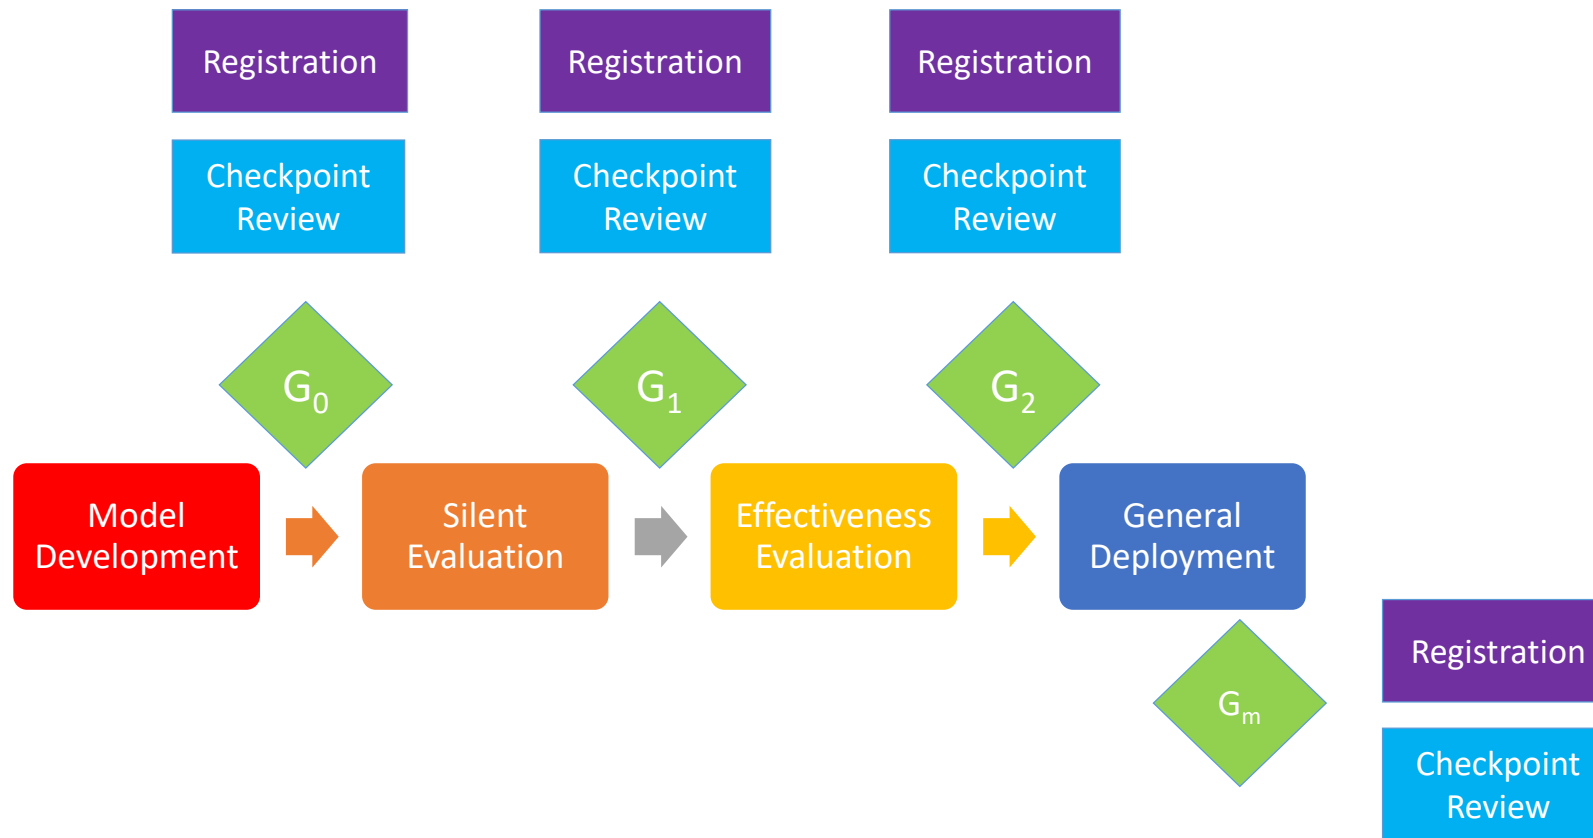


## Outcomes

- Approval
- Approval w/ Contingencies
- Re-review
- Denial

\* Clinical consensus-based models submit their review material during the registration process; these models will typically not require a full checkpoint review.

# Full Checkpoint Review





# Implementing Quality & Ethics with Our Framework

Transparency & Accountability

Impact & Safety

Fairness & Equity

Usability & Adoption

Regulatory Compliance



# Implementing Quality & Ethics with Our Framework

Transparency & Accountability

Impact & Safety

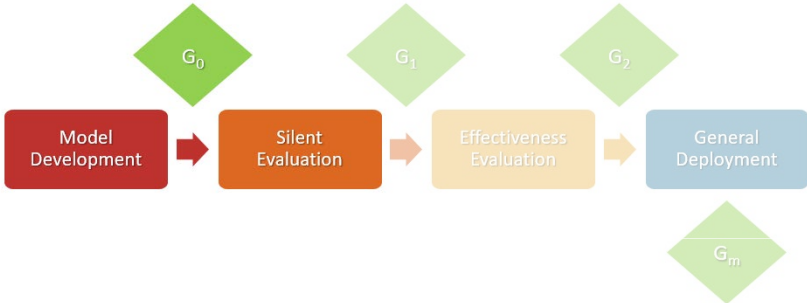
Fairness & Equity

Usability & Adoption

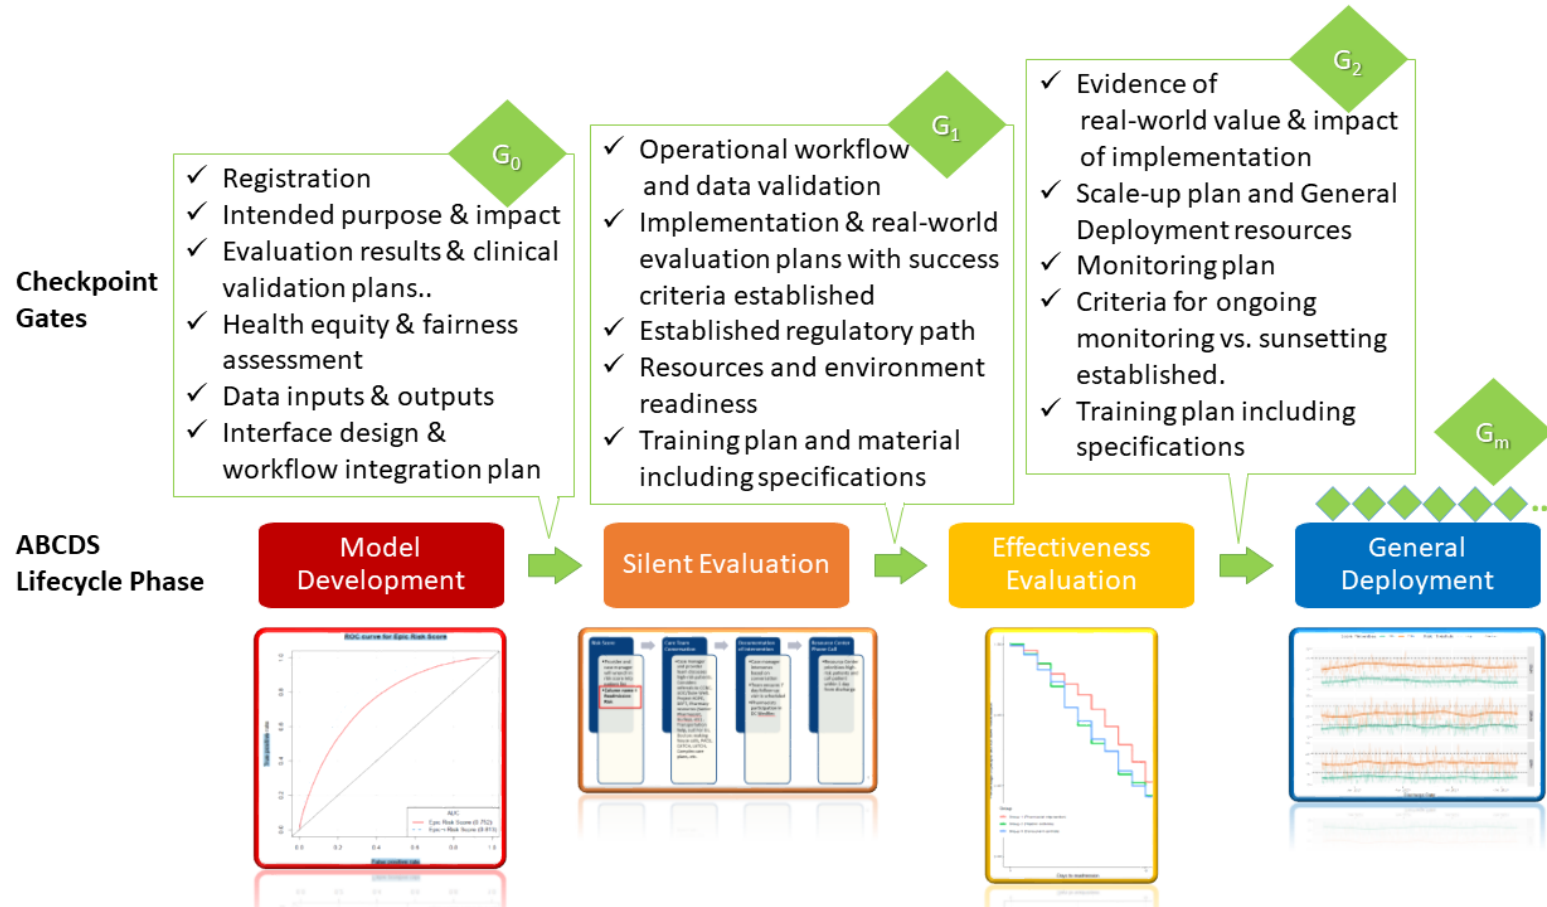
Regulatory Compliance

Principle	Criteria	Submission Materials
Clinical Impact & Safety	The ABCDS software, in comparison to current state, stands to improve clinical care.	<ul style="list-style-type: none"><li>✓ Evidence that the tool has potential to impact clinical outcomes or processes</li><li>✓ List of key impact metrics (clinical outcomes and/or process improvement) with definitions, following TRIPOD guidelines<sup>5</sup></li><li>✓ List of core performance metrics (e.g. sensitivity, PPV, etc.) and results from development</li><li>✓ Calibration curves, threshold selections and justification if applicable</li></ul>
	Plans for Silent Evaluation will inform the decision to proceed with pilot implementation in clinic.	<p>Silent Evaluation Plan, including:</p> <ul style="list-style-type: none"><li>✓ Summary of benefits you expect to demonstrate and criteria to proceed into Effectiveness Evaluation</li><li>✓ Study design, including in/exclusion criteria, timeframe and sample size considerations</li><li>✓ Core performance metrics with shell tables</li><li>✓ Data analysis plan</li><li>✓ Data quality evaluation plan</li></ul>

*Sample evaluation criteria supporting the principle of clinical impact & safety at the G<sub>0</sub> Checkpoint evaluation between pilot implementation and general deployment*



# ABCDs Oversight Full Review




Bedoya, A. D., et al. (2022). "A framework for the oversight and local deployment of safe and high-quality prediction models." Journal of the American Medical Informatics Association.



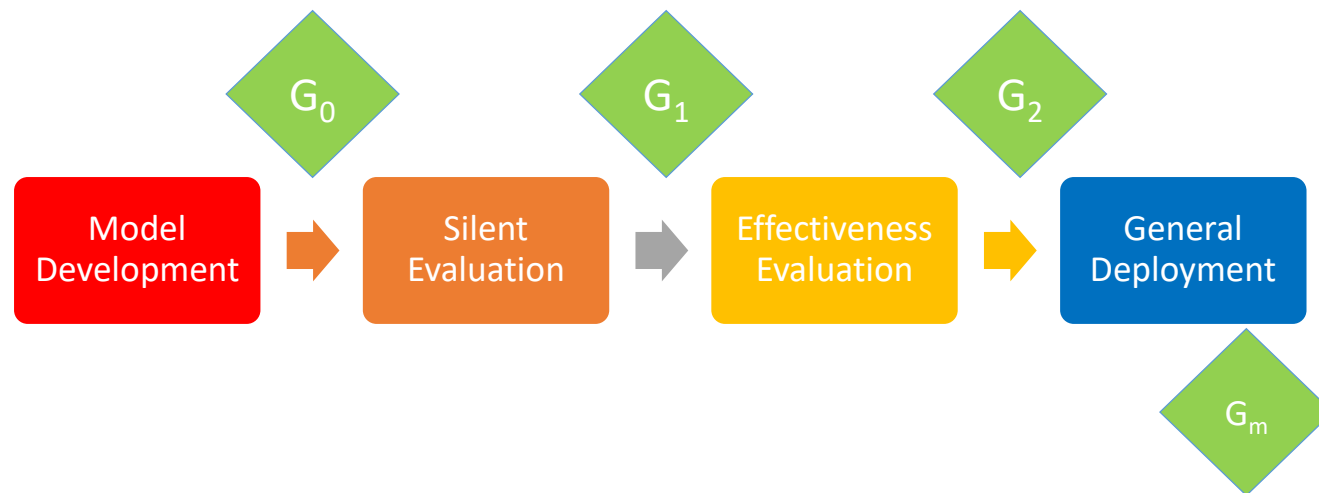
# Scope of ABCDS Oversight Framework





# Fast-Track Checkpoint Review

## Full Review

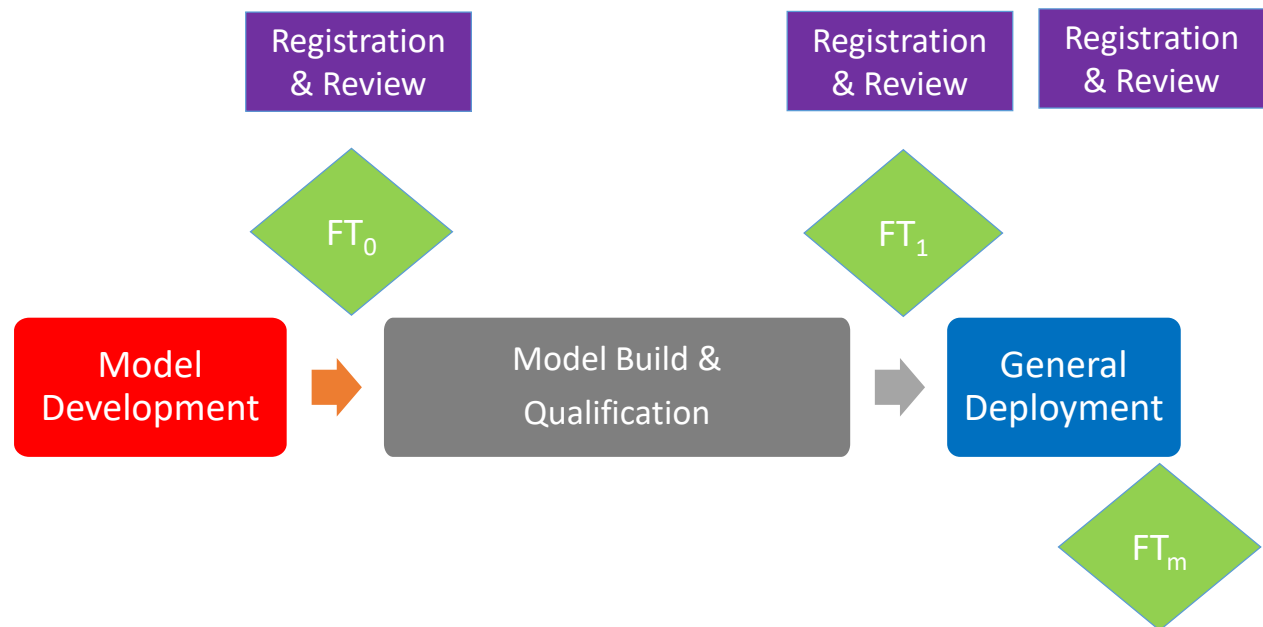


## Fast-Track Review

Upon registration:

\*Clinical Validation

\*Technical Specification Document



# Fast-Track Evaluation – Alignment without Guidelines

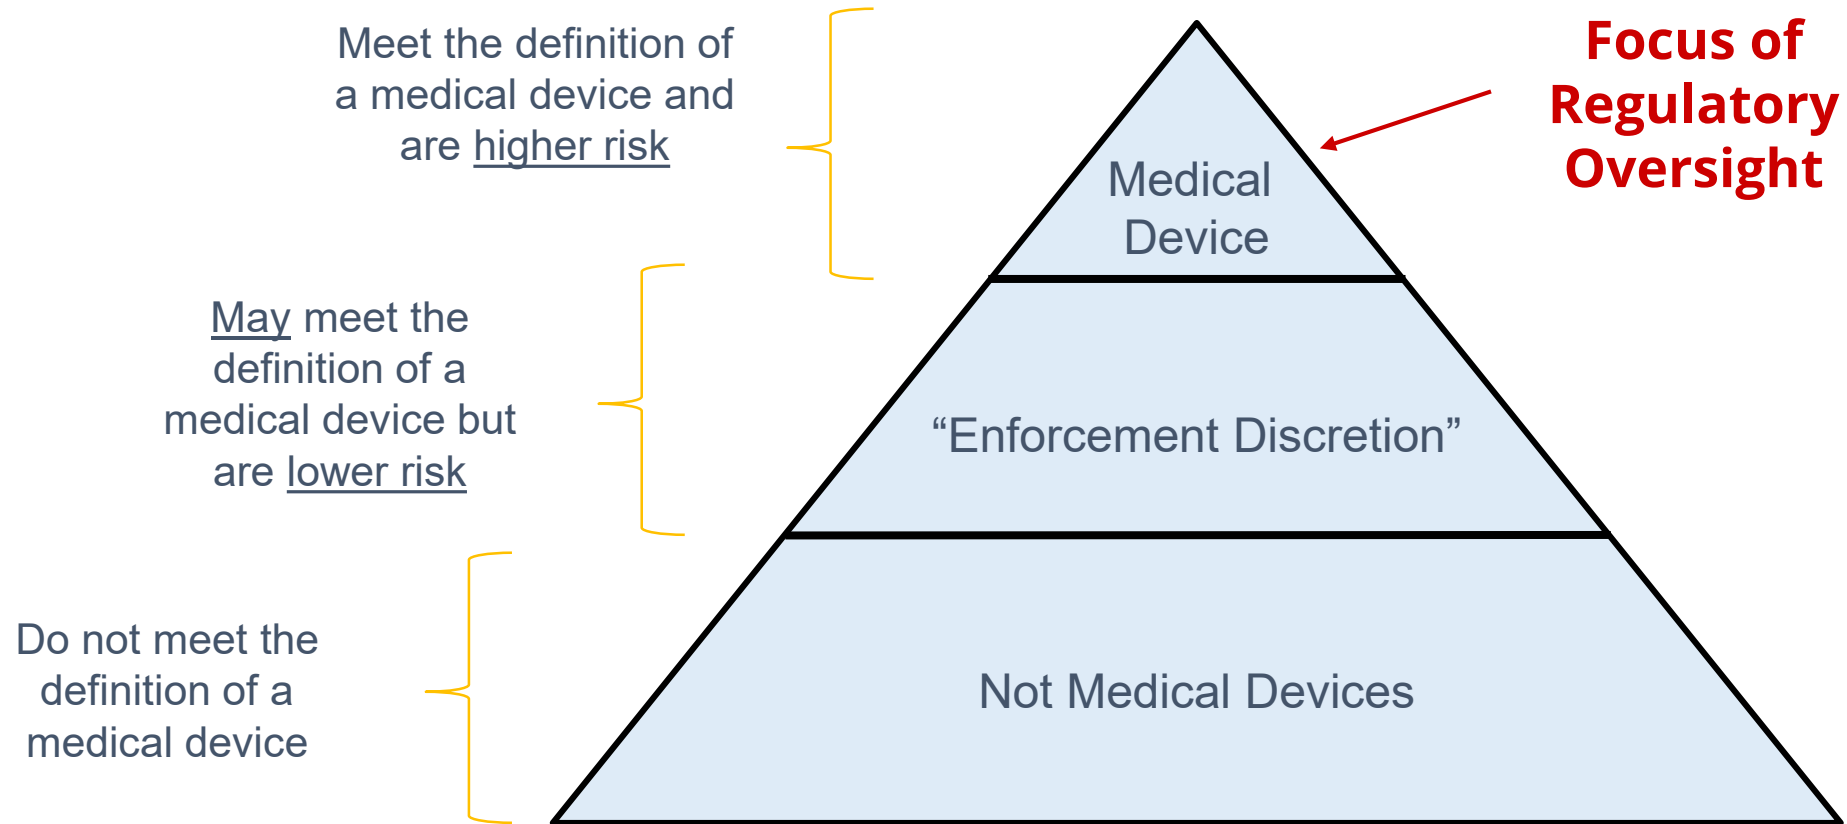
Clinical Evaluation		
Valid Clinical Association	Analytical Validation	Clinical Validation
Is there a valid clinical association between your SaMD output and your SaMD's targeted clinical condition?	Does your SaMD correctly process input data to generate accurate, reliable, and precise output data?	Does use of your SaMD's accurate, reliable, and precise output data achieve your intended purpose in your target population in the context of clinical care?

**Figure 4- Clinical Evaluation Process**

Software as a Medical Device (SaMD): Clinical Evaluation, FDA Guidance, 2017

# Regulatory Considerations

# Risk-Based Approach for Regulation of Software Functions





# Medical Device

The Food, Drug, and Cosmetic Act (FD&C) defines a medical device as:

- An instrument, apparatus, implement, machine, contrivance, implant, in vitro reagent or other similar or related article or component part or accessory which:
  - Is intended for use in the diagnosis of disease or other conditions, or in the cure, mitigation, treatment, or prevention of disease or
  - Is intended to affect the structure or any function of the body; and
  - Does not achieve any of its primary intended purposes through chemical action within or on the body of man and is not dependent upon being metabolized for the achievement of any of its primary intended purposes

# Software functions that are NOT medical devices

**Software functions that could be used in a healthcare environment, in clinical care or patient management, but do not meet the definition of a medical device.**



FD&C Act does not apply;

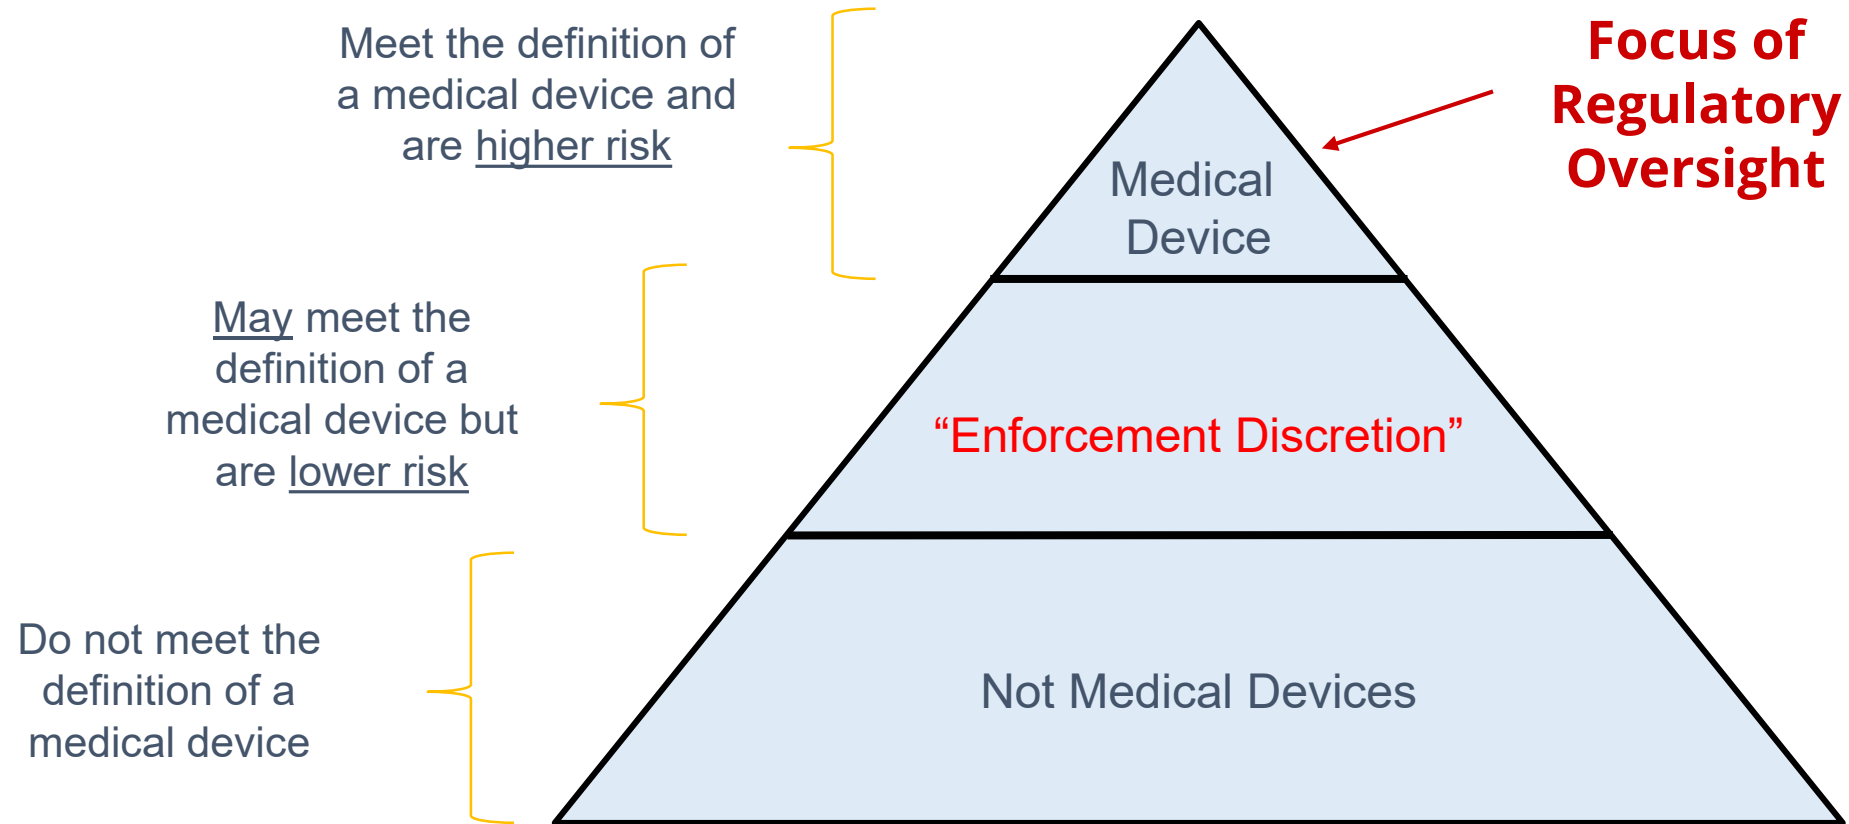
Not regulated by the FDA!

Examples: general purpose products, data transfer/storage only, **some** clinical decision support, **some** general wellness products


# General Wellness Software

- Not a Medical Device: Software intended to maintain or encourage general state of health
  - Weight management, stress, fitness, mental acuity, sleep, self esteem, etc.
- Enforcement Discretion: Intended use relates the role of healthy lifestyle with helping to reduce the risk or impact of certain chronic diseases or conditions
  - *Help living well with or reduce the risk of* heart disease, high blood pressure, type 2 diabetes, anxiety, etc.

# Risk-Based Approach for Regulation of Software Functions



# Enforcement Discretion



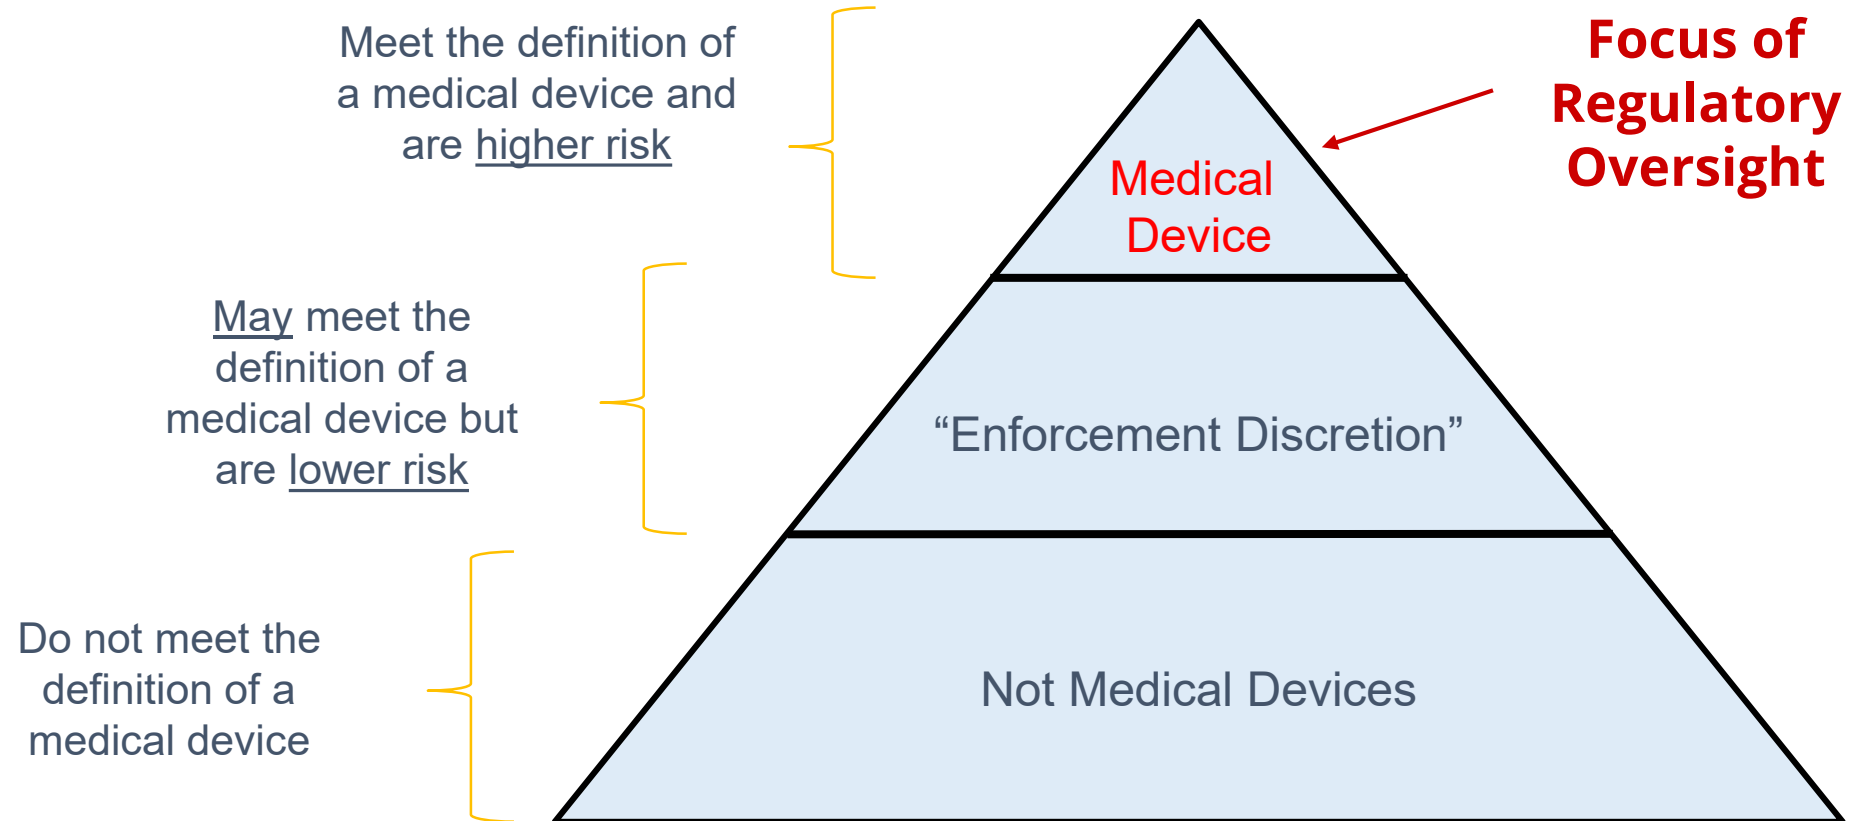
**Software functions that may meet the definition of a medical device for which the FDA intends to exercise ‘enforcement discretion’**

Under FDA jurisdiction, but “FDA intends not to pursue enforcement action for violations of the FD&C Act”.

Examples: **some** general wellness products, apps that coach/prompt, medical calculators




# Risk-Based Approach for Regulation of Software Functions



# Software Functions that are the Focus of FDA Regulatory Oversight

Software that meets the definition of a medical device and either is intended:

- to be used as an accessory to a regulated medical device; or
- to transform a mobile platform into a regulated medical device.



**FD&C Act will be enforced;  
FDA will regulate this software  
function.**

# FDA Final Guidance 2022

*Contains Nonbinding Recommendations*

## **Clinical Decision Support Software**

### **Guidance for Industry and Food and Drug Administration Staff**

**Document issued on September 28, 2022.**

**The draft of this document was issued on September 27, 2019.**

For a software function to be Non-Device CDS, it must meet all the following four criteria to be excluded from the device definition under section 520(o) of the FD&C Act.

1	Not 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
2	Intended for the purpose of displaying, analyzing, or printing medical information about a patient or other medical information
3	Intended for the purpose of supporting or providing recommendations to an HCP about prevention, diagnosis, or treatment of a disease or condition
4	Intended for the purpose of enabling an HCP to independently review the basis for the recommendations that such software presents so that it is not the intent that the HCP rely primarily on any of such recommendations to make a clinical diagnosis or treatment decision regarding an individual patient



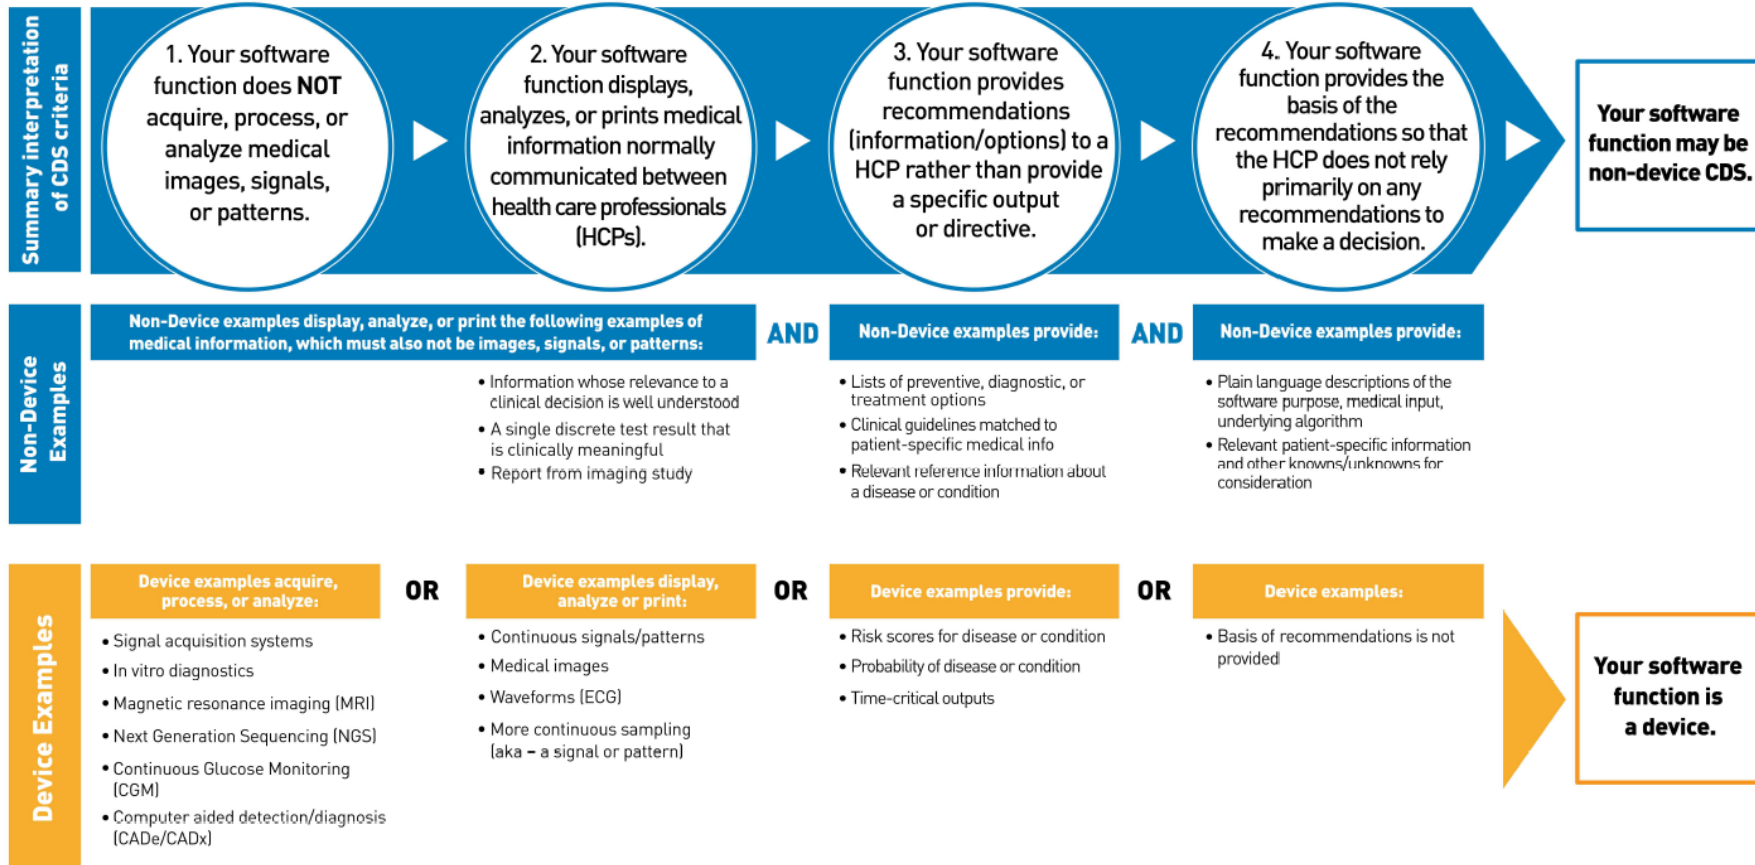
**Duke**  
AI HEALTH

# Your Clinical Decision Support Software: Is It a Device?



The FDA issued a guidance, Clinical Decision Support Software, to describe the FDA's regulatory approach to Clinical Decision Support (CDS) software functions. This graphic gives a general and summary overview of the guidance and is for illustrative purposes only. Consult the guidance for the complete discussion and examples. Other software functions that are not listed may also be device software functions. \*

## Your software function must meet all four criteria to be Non-Device CDS.



**\*Disclaimer:** This graphic gives a general overview of Section IV of the guidance ("Interpretation of Criteria in Section 520(o)(1)(E) of the FD&C Act"). Consult the guidance for the complete discussion. The device examples identified in this graphic are illustrative only and are not an exhaustive list. Other software functions that are not listed may also be device software functions.

### Criterion 3: Automation Bias and Time Criticality

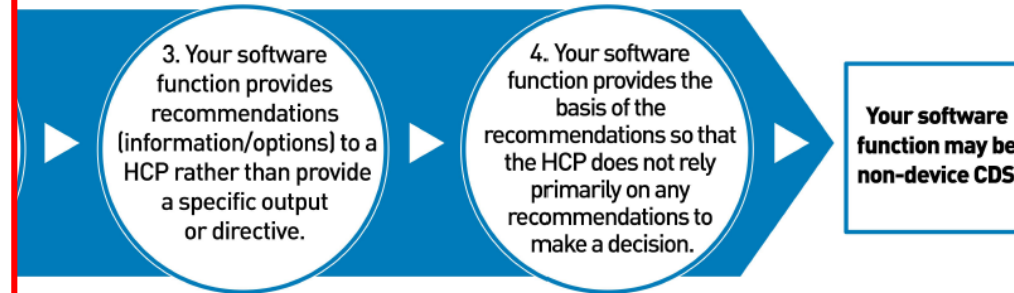
- Propensity of humans to over-rely on a suggestion from an automated system.
- Can result in errors of commission/omission.
- May be more likely to occur if software provides a user with a single, specific output/solution compared to a list of options or complete information to consider.
- Automation bias increases in time critical situations as the user may not have adequate time to consider other information.

## Clinical Decision Support Software: Is It a Device?

FDA

FDA's regulatory approach to Clinical Decision Support (CDS) software functions. This graphic gives a general overview of the guidance. Consult the guidance for the complete discussion and examples. Other software functions that are not listed

### all four criteria to be Non-Device CDS.



AND

Non-Device examples provide:

- Lists of preventive, diagnostic, or treatment options
- Clinical guidelines matched to patient-specific medical info
- Relevant reference information about a disease or condition

AND

Non-Device examples provide:

- Plain language descriptions of the software purpose, medical input, underlying algorithm
- Relevant patient-specific information and other knowns/unknowns for consideration

Non-Device Examples

- A single discrete test result that is clinically meaningful
- Report from imaging study

Device Examples

Device examples acquire, process, or analyze:

- Signal acquisition systems
- In vitro diagnostics
- Magnetic resonance imaging [MRI]
- Next Generation Sequencing [NGS]
- Continuous Glucose Monitoring [CGM]
- Computer aided detection/diagnosis [CADe/CADx]

OR

Device examples display, analyze or print:

- Continuous signals/patterns
- Medical images
- Waveforms [ECG]
- More continuous sampling [aka – a signal or pattern]

OR

Device examples provide:

- Risk scores for disease or condition
- Probability of disease or condition
- Time-critical outputs

OR

Device examples:

- Basis of recommendations is not provided

**Your software function is a device.**

**\*Disclaimer:** This graphic gives a general overview of Section IV of the guidance ("Interpretation of Criteria in Section 520(o)(1)(E) of the FD&C Act"). Consult the guidance for the complete discussion. The device examples identified in this graphic are illustrative only and are not an exhaustive list. Other software functions that are not listed may also be device software functions.



## Criterion 4: Independent Review

- Provides background information in plain language on the inputs, algorithm logic/methods, datasets, validation, and patient information (detailed list in guidance)
- Expected regardless of software complexity and proprietary nature
- Recommends considering usability testing

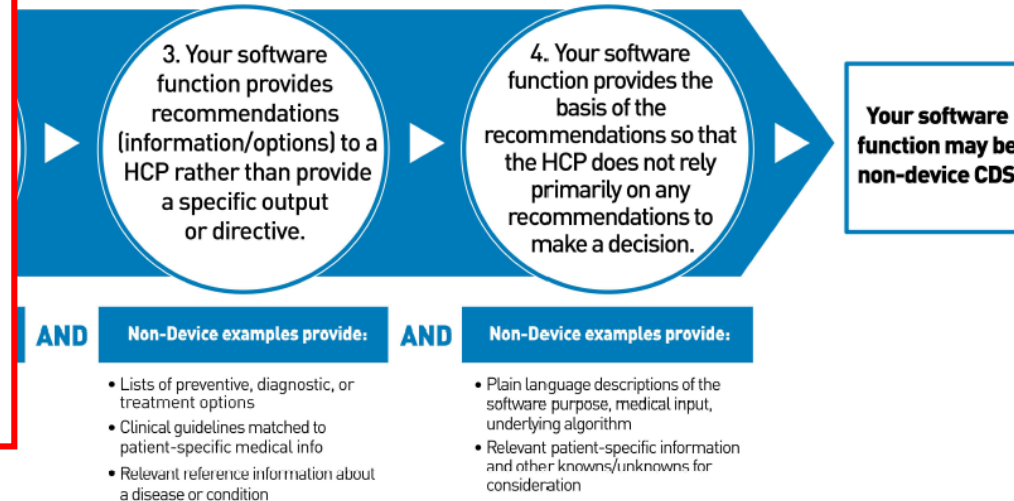
Note: References time criticality again highlighting that FDA does not consider software functions supporting a critical time sensitive task/decision to meet this criterion as HCP is unlikely to have sufficient time to do independent review.

## Clinical Decision Support Software: Is It a Device?

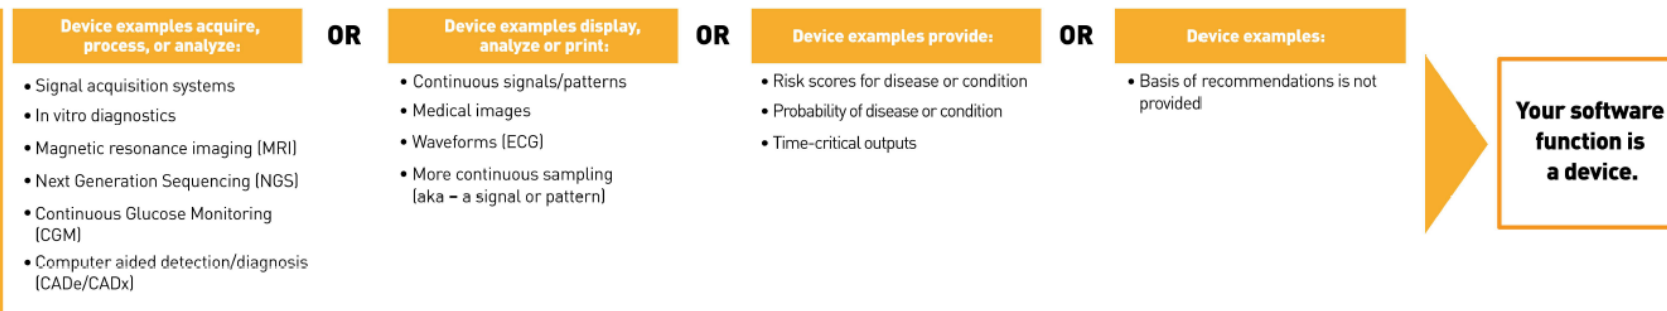
FDA

FDA's regulatory approach to Clinical Decision Support (CDS) software functions. This graphic gives a general overview of the guidance. Consult the guidance for the complete discussion and examples. Other software functions that are not listed

### All four criteria to be Non-Device CDS.



### Device Examples



**\*Disclaimer:** This graphic gives a general overview of Section IV of the guidance ("Interpretation of Criteria in Section 520(o)(1)(E) of the FD&C Act"). Consult the guidance for the complete discussion. The device examples identified in this graphic are illustrative only and are not an exhaustive list. Other software functions that are not listed may also be device software functions.

# Tools for Engaging FDA

- FDA Digital Health Inbox
- FDA Digital Health Policy Navigator  
<https://www.fda.gov/medical-devices/digital-health-center-excellence/digital-health-policy-navigator>
- Q-Submission



Software functions intended for administrative support of a health care facility are not devices under section 520(o) of the FD&C Act.

# Impacting How We Deliver Patient Care



# Lessons Learned

- Successful AI Governance is a Team Sport
  - Lots of skillsets, perspectives and languages to bring together
- Culture Shift is Hard
  - Governance's role as Coaches and Facilitators (*not* Punisher)
  - Show Teams how to succeed by addressing gaps in their knowledge, skillsets, and/or bandwidth
  - There is no such thing as over-communication in a complex system
- Benefits of Centralized Governance
  - Transparency of Process & Expectations
  - Institutional Visibility into all the 'skeletons in the closet'
- Conscious Decision (thus far) Not to Regulate Who Gets to Build AI Models





# Future Directions

- Translating FDA guidance to practice
- Imaging
- Centralized Model Monitoring



# Learn More...

<https://aihealth.duke.edu/algorithm-based-clinical-decision-support-abcds/>



## What is ABCDS?

Algorithm-Based Clinical Decision Support (ABCDS) Oversight is a “people-process-technology” framework for the governance and evaluation of clinical algorithms created for use at Duke Health. This framework fosters innovative, safe, equitable, and high-quality patient care by introducing checkpoints throughout the development lifecycle as well as after deployment to ensure that transparency, quality, and ownership are maintained for ABCDS algorithms and tools. The ABCDS Oversight is a collaborative effort between the Duke University School of Medicine and the Duke University Health System.

Bedoya, A. D., et al. (2022). "A framework for the oversight and local deployment of safe and high-quality prediction models." [Journal of the American Medical Informatics Association.](#)

## Questions & Feedback

Contact us at [abcds@duke.edu](mailto:abcds@duke.edu)



[Learn More](#)[Insights](#)[News](#)[Join Us](#)[Join a Working Group](#)

Providing guidelines for the  
responsible use of AI in  
healthcare

[Learn More →](#)

<https://www.coalitionforhealthai.org/>

# BLUEPRINT FOR TRUSTWORTHY AI IMPLEMENTATION GUIDANCE AND ASSURANCE FOR HEALTHCARE

COALITION FOR HEALTH AI

APRIL 04, 2023



Duke  
AI HEALTH

Thank you

