Responsible AI Checklist (RAIC) for Health AI

Initial Planning Checkpoint

Coalition for Health AI

June 26, 2024

Checklist Document Versions

As this checklist is passed back and forth between different Reporters and Reviewers, Table 1 will help track versions of the document. Italicized information in the checklist serve as examples and should be replaced during use.

Versions						
Document Version	**Application & Model Version**	**Content Description**	**Reporter or Reviewer Name**	**Contact Information and Role**	**Organization**	**Date**
<1.0>	*<EHR-Based Pediatric Asthma Exacerbation Risk version 1.0 Model 2.0.>*	*<Documentation and evidence provided by implementer and development teams/specific departments from Mayo Clinic>*	*<Name>*	*<Reporter 1>* *E-mail:* *Phone:* *Title:*	*<Mayo Clinic>*	*<May 1, 2024>*
<2.0>	*<EHR-Based Pediatric Asthma Exacerbation Risk version 1.0 Model 2.0.>*	*<Documentation and evidence related to use and human-factors considerations provided by external consultant at ideas42>*	*<Name>*	*<Reporter 2>* *Email:* *Phone:* *Title:*	*<ideas42>*	*<May 5, 2024>*
<3.0>	*<EHR-Based Pediatric Asthma Exacerbation Risk version 1.0 Model 2.0.>*	*<Summary of findings and review of documentation and evidence provided by development and implementer teams at Mayo and consultants from ideas42>*	*<Name>*	*<Reviewer 1>* *Email:* *Phone:* *Title*	*<Mayo Clinic>*	*<May 7, 2024>*

Table 1. Checklist Document Versions

Table of Contents

Table 1: Stakeholder Roles, Professions, and Representative Organizations. Derived from CHAI Responsible AI Guide (Link)

1 Purpose and Use

1.1. Purpose

The Responsible AI Checklist (RAIC) is intended to guide the development and evaluation of a complete **AI solution** and **system** against CHAI content for trustworthy AI[1]. This tool is intended first for self-reporting and self-review, as well as a tool for self-reporting for independent review. The goal of the RAIC is to ensure that AI solutions and systems fulfill all five key, principle-based areas for trustworthy AI: 1. Usefulness, Usability, and Efficacy; 2. Fairness; 3. Safety; 4. Transparency and Intelligibility; 5. Privacy and Security. In alignment with these areas, the RAIC translates best practice considerations (detailed in the Responsible AI Guide) that meet core ethical and quality principles into detailed yes/no questions, or evaluation criteria, to determine whether best practices are met (see accompanying Responsible AI Guide). The relationship between evaluation criteria and their original considerations, as well as criteria that have been combined across multiple areas and considerations are mapped in a Traceability Matrix located in the Appendix of this document (Section 3.1). The RAIC encourages a holistic understanding of AI solutions in context, encompassing the interplay of human-factors, data, algorithms, infrastructure, and real-world workflows, facilitating conversations across developer and implementer teams, and As a self-review tool for developer and implementation teams, this iteration of the RAIC also serves as a starting point for facilitating conversation and alignment on best practices across the full AI lifecycle.

A secondary purpose of this version of the tool is to guide an understanding of the state of trustworthy AI in healthcare and the needs of representative stakeholders and healthcare organizations by stress-testing the checklist in the real-world. Specifically, utilization of this tool and feedback on existing end-to-end capabilities and practices will aid both in improving and iterating on the RAIC and its subsequent versions, as well as an understanding of the challenges that may influence the feasibility of best practices.

1.2. Intended Users

Intended users of the RAIC are developer and implementation teams within or outside of health systems with accountable Reporters from teams providing documentation and summaries for executive review. Multiple stakeholders (see section 3.3 in the Appendix and section 3.2 in the Responsible AI Guide) may be involved in the selection, procurement, development, and deployment process of an AI solution. This

[1] The RAIC was developed by forming expert workgroups for each principle area. Workgroups conducted a full landscape analysis and synthesized findings into a series of considerations and criteria for each lifecycle stage for their specific principle-based focus areas. These considerations and criteria were then compiled into a survey sent out to the broader CHAI community to gain multi-stakeholder feedback and ratings as part of a modified Delphi-process to gain consensus across multiple stakeholders. Results were then reviewed during the Fall convening and discussed further. Considerations that were rated as “Extremely Important” by at least 50% of the respondents, and/or were deemed extremely important following the second round of discussions, were included in this version of the Responsible AI Guide and Checklist. Additional considerations and criteria that were rated as either “Extremely Important” or “Very important” by at least 65% of survey respondents are included in the Traceability Matrix but not in this version of the Responsible AI Guide or Checklist.

iteration of the RAIC does not prescribe roles and responsibilities, however it outlines usage for those completing and reviewing the document (see Responsible AI Guide, pg. 2 for further details on this and plans for future versions). Developer and implementer teams may be entirely or in part internal or external to the healthcare organization looking to develop, procure, or implement an AI solution. As such, this tool may also be used as part of a collaborative process across developer and implementer teams to foster trust and alignment on best practices.

This checklist is most appropriate for products or devices that are themselves AI software (predictive or generative) or those that are AI assisted/AI enabled. At this point in time, AI tools often used in drug discovery and development (e.g. target selection or antibody design) in the pharmaceutical industry fall outside the targeted scope of the RAIC.
AI software examples: Payer/provider risk stratification or prediction, diagnostic algorithms, automated EHR coding, provider decision or administrative support, patient decision support, patient or provider facing chatbot used for education or assistance
AI assisted/AI enabled examples: AI enabled medical devices, AI assisted surgical robots, radiological technologies that are AI assisted or AI enabled for clinical (diagnostic/risk prediction) or nonclinical purposes (automated image quality enhancement.)

The **Reporter** is the individual tasked to gather responses and documentation from appropriate **"Providers of Evidence,"** or experts in the areas pertaining to RAIC items. The **Reviewer** can either be an internal executive responsible for checking the completeness and appropriateness of the explanations and documentation to guide the development, procurement, and/or implementation of an AI solution based on best practices, or an external independent Reviewer who will evaluate the overall AI system for alignment with best practices. Note that there may be multiple Reporters, Providers of Evidence, and Reviewers. For smaller organizations or health systems there may be fewer stakeholders available, or the need to consult with external experts to ensure best practices in specific areas. We do not expect that all best practices are feasible at this point and aim to further understand feasibility as they are stress-tested in the real world. Examples of user personas and scenarios are provided in the Appendix (section 3.4).

1.3. Usage

Usage of the RAIC is guided by the AI Lifecycle (Figure 1). The AI Lifecycle can be an iterative and non-linear/agile outline of the processes required for effective and trustworthy design, development, and use of a health AI system from end-to-end. To facilitate the agile process, we have identified a **planning checkpoint** and several **responsible AI checkpoints** that aim to help teams ensure that the necessary steps have been taken, prior to moving a tool into real-world use. The four checkpoints are summarized below. Examples of user personas and scenarios are provided in the Appendix (section 3.4).

1. The **planning checkpoint** follows Stage 1, where both developer and implementer teams (independently or together) are asked to define the specific problem and plan adequately for a potential AI solution. This checkpoint primarily helps teams:
    a. Appropriately consider the risks, benefits, costs, and needs for an AI solution both at the clinical and population levels
    b. Consider the risks, benefits, costs, and needs around purchasing or developing an AI solution in house
    c. Gain multi-stakeholder insights to help guide human-centered AI solution design, development (or purchasing) and downstream needs to maximize real-world effectiveness and trust
2. **Responsible AI checkpoint one** appears when progressing from iterations through design, development, and assessment processes, to the small-scale pilot phase. The goal of this checkpoint is to address readiness for piloting and to prepare for real-world risks and needs. Any updates to clinical and population risk summaries should be made based on new insights from the design, development, and silent-evaluation process. An important note is that this checkpoint is not only meant for developer organizations. There are items that assess for readiness for the implementer/purchasing organization, items to

guide conversations around responsibilities between developer and implementer organizations, items that speak to the larger AI system design and development (e.g. safety, privacy, security, and monitoring planning), and items that a purchasing/implementing organization may use to understand vendor best practices. An organization or health system acquiring or purchasing an AI solution may choose to use this checkpoint as part of their procurement process. For example, they may require developer organizations to provide relevant evidence in support of best practices during design, development, and evaluation to help make purchasing decisions to foster transparency. It is also recommended that purchasing/implementing organizations review the planning checkpoint items alongside the developer organization to ensure appropriate planning, risk determination, and usability for the broader AI system (beyond the AI solution alone).

3. **Responsible AI checkpoint two** appears when progressing from piloting to at-scale deployment of the AI system, which requires evaluation of readiness and preparation for the broader needs and wider scope of risk. Any updates to clinical and population risk summaries should be made based on new insights from initial real-world piloting.
4. **Responsible AI checkpoint three** appears following full scale deployment to evaluate for longer-term readiness for monitoring, managing, and updating the AI system. This checkpoint is repeated throughout regular monitoring of the AI solution, at appropriately timed intervals depending on the use case, and as dictated by the developer and/or implementer organization. As in previous checkpoints, updates should be made to clinical and population risk summaries based on insights gained from regular monitoring of AI solutions and systems.

Within each checkpoint checklist, relevant evaluation criteria are listed and given an identifier. The color coded Evaluation Criteria Identifier (EC Identifier) links each criterion to the original consideration as defined within principle area workgroups (see Traceability Matrix in the Appendix 3.1; See Section 1.5 for further details.)

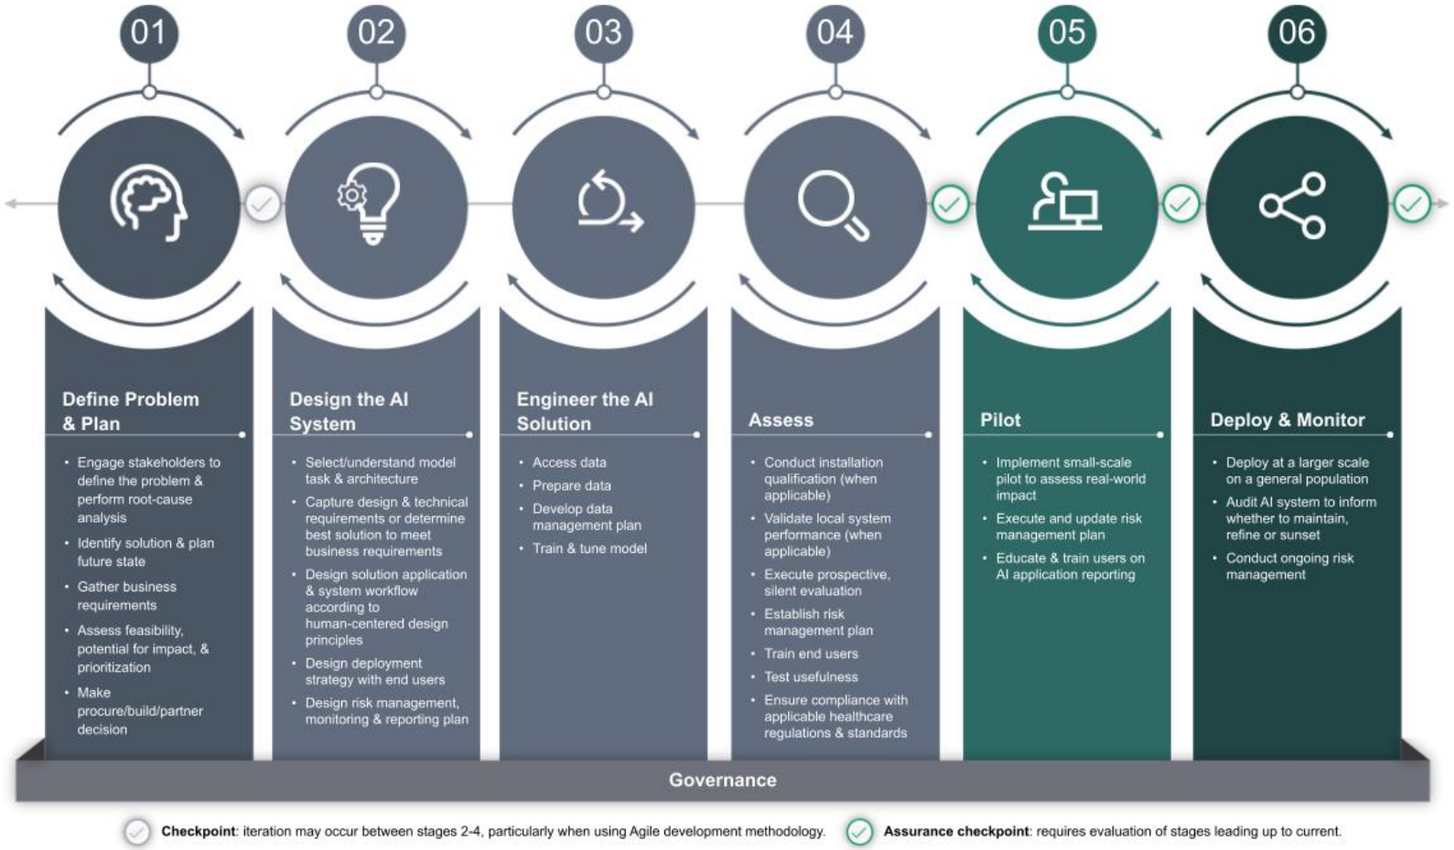

Figure 1: The CHAI AI Lifecycle Framework. Derived from CHAI Responsible AI Guide. The gray checkmark represents the Planning Checkpoint, while the green checkmarks correspond to Responsible AI Checkpoints 1-3.

1.4 How to complete this checklist

1.4.1 General

Who Should Complete This Checklist?

Each checkpoint checklist should first be completed by at least one **Reporter.** While there may be multiple stakeholders involved in sharing evidence necessary to respond to criteria, the Reporter is the individual responsible for requesting this information (if available), making sure available evidence is clearly documented for relevant evaluation criteria in the checklist, and indexing it in a centralized place for ease of Reviewer access. They will also provide a summary at the end of each checkpoint that provides reviewers with a broad overview of the potential or observed benefits, costs, risks, and/or adverse events associated with that checkpoint. Example roles, professions, and representative organizations are shown in Table 1, Section 3.3 in the Appendix and described in more detail in the CHAI Responsible AI Guide.

Reporters will then pass the checklist off to at least one **Reviewer** who is internal to either developer and/or implementer organizations (such as an area specific executive). Ideally, organizations will also pursue independent and external third-party review. The Reviewer will go over the responses to evaluation criteria and evidence, and indicate whether best practices for each criteria have been met. They will also provide a summary of findings based on the available evidence and any observed gaps. This feedback can be used to improve processes, help guide teams on next steps, or help build/design solutions to fill gaps in best practices.

Reporter Responsibilities for Completion (Planning Checkpoint)

1. All Reporter required sections of the checklist or summaries are denoted with dark blue coloring.
2. During the planning phase (Stage 1) of the AI solution, Reporters should gather information from relevant stakeholders in order to complete an initial summaries of **Clinical Risk** (see Section 1.4.2 for further instructions) and **Population Impact** (see Section 1.4.3 for further instructions).
3. After using the available tools, the Clinical Risk classification and summary and Population Impact summary should be included in **Table 3** located at the start of each checkpoint. This summary will provide context to reviewers as they evaluate responses to additional criteria. It is important to update clinical risk and population impact tables as new related insights are gained ***at every checkpoint***.
4. The Reporter will then complete the Planning Checkpoint Checklist (Stage 1) providing a brief explanation and document code in the "Evidence and Explanation & Metadata/Documentation Code" column of the checklist, with supporting evidence indexed within the "Evidence & Explanation Metadata Table" (see Section 2.5 for further instructions and Table).
5. The Reporter will complete the "Executive Summary of Anticipated and Observed Benefits, Risks, and Limitations" section for the planning checkpoint.
6. Reporter responsibilities for a checkpoint will end by updating the document version table (Page 2) and up-versioning the document header, prior to sending the checklist and associated evidence to the appropriate Reviewer.

Reviewer Responsibilities for Completion (Planning Checkpoint)

1. All Reviewer required sections of the checklist or summaries are denoted with light blue coloring.
2. The Reviewer will go through information provided in the checklist by the Reporter along with accompanying documentation listed in Evidence and Explanation Metadata table.

3. Reviewers will then complete the Summary of Findings table (Section 2.4), summarizing findings provided in the checklist by the Reporter in the context of anticipated and observed benefits, risks, and limitations of the AI solution.
4. Reviewers will then update the document version table on Page 2 and up-version the document header.

Example Reporter Role Responses

Checklist: Stage 2-4 | Design, Engineer, and Assess the AI Solution								
EC Identifier	**Evaluation Criteria**	**Evidence and Explanation & Metadata/Document Code**	**N/A or Cannot Complete (CC): Describe in comment**	**Reporter Initials & Date**	**Evidence & Explanations Provided? (Yes/No/ Partial/NA)**	**Benefits, Limitations, or Adverse Outcomes**	**Criteria Met? (Yes/No/ Partial/NA)**	**Reviewer Initials & Date**
	Responsible AI Checkpoint 1: Readiness for Real World							
LS2.F.C1.EC2	Will the real-world/clinical outcome measure be available for evaluation within an adequate time frame and in a manner that accurately represents the target population?	*Evidence and explanation: Real-world retrospective data was used for evaluation of model performance and comparable to target population.* *Metadata/Document Location: <insert link to bias assessment document and relevant data showing summary of real-world retrospective data population descriptives and demographics and comparison to target population descriptives and demographics.)*		*M.G. 05/06/2024*				
LS2.F.C1.EC3	Will real-world/clinical outcomes be systematically compared for impartiality across all relevant socio-demographic subgroups, ensuring fairness and addressing potential bias?	*Evidence and explanation: Overall ER admission rates are lower following use of the AI solution. Clinical outcomes are similar for all subgroups except for Black Patients, who show higher ER admissions following discharge at the same population level risk threshold compared to the sample majority group and compared to the population mean.* *Metadata/Document Location: <insert link to bias assessment document and relevant*		*M.G. 05/06/2024*				

		data showing likelihood of ER admissions following discharge (as measure of clinical outcomes that AI solution aimed to impact)						

Example Reviewer Role Responses

Checklist: Stage 2-4 | Design, Engineer, and Assess the AI Solution									
EC Identifier	**Evaluation Criteria**	**Evidence and Explanation Metadata/Document Code**	**Not applicable: describe in comment**	**Cannot be completed: Describe in comment**	**Reporter Initials & Date**	**Evidence & Explanations (Yes/No/Partial/NA)**	**Limitations or Adverse Outcomes**	**Criteria Met (Yes/No/Partial/NA)**	**Reviewer Initials & Date**
	Responsible AI Checkpoint 1: Readiness for Real World								
LS2.F.C1.EC2	Will the real-world/clinical outcome measure be available for evaluation within an adequate time frame and in a manner that accurately represents the target population?	*Evidence and explanation: Real-world retrospective data for ER admission rates are available and will be used for evaluation of model's impact on clinical outcomes. Data is comparable to target population.* *Metadata/Document Location: <insert link to bias assessment document and relevant data showing summary of real-world retrospective data population descriptives for measure and demographics and comparison to target population descriptives for measure and demographics of sample.)*			*M.G. 05/06/2024*	*Yes*	*No, None stated*	*Partial, Provide justification for why this clinical outcome was selected.*	*N.E. 05/10/2024*
LS2.F.C1.EC3	Will real-world/clinical outcomes be systematically compared for impartiality across all relevant socio-demographic subgroups, ensuring fairness and addressing potential bias?	*Evidence and explanation: Overall ER admission rates are lower following use of the AI solution. Clinical outcomes are similar for all subgroups except for Black Patients, who show higher ER admissions following discharge at the same population level risk threshold compared to the sample majority group and compared to the population mean.* *Metadata/Document Location: <insert link to bias assessment document and relevant data showing likelihood of ER admissions following discharge (as measure of clinical*			*M.G. 05/06/2024*	*Partial, provide information on what threshold was selected and why.*	*Yes, Black patients have poorer outcomes at the chosen threshold*	*Partial*	*N.E. 05/10/2024*

		outcomes that AI solution aimed to impact)							

1.4.2 Clinical Risk Evaluation

*Risk should be assessed from both the **clinical** and **population** perspective. For **clinical risk**, we adopt the International Medical Device Forum's (IMDRF's) categorization system for assessment of clinical risk (See Table 2). This should be done by a licensed clinician based on the FDA IMDRF guidance.*

Table 2. Assessment criteria for clinical risk level. Levels are described in detail in "Software as a Medical Device": Possible Framework for Risk Categorization and Corresponding Considerations" by IMDRF Software as a Medical Device (SaMD) Working Group (2014).

Clinical Risk Classification			
State of Healthcare situation or condition	**Significance of information provided to healthcare decision**		
	Treat or diagnosis	Drive clinical management	Inform clinical management
Non-Serious	II	I	I
Serious	III	II	I
Critical	IV	III	II

Clinical risk classification and summaries should be provided in **Section 2.1, Table 3. Clinical Risk and Population Impact Evaluation Summaries**

1.4.3 Population Impact Evaluation Tool:

Population risk refers to how systemic, individual, and group-level tendencies when combined with decision-making demands across the AI lifecycle, can impact health and well-being for entire subgroups and over longer periods.

While it is common to refer to systemic, individual, and group-level tendencies as "biases"—it is important to note that they are often the result of things like:

- Historical Norms/policies
- Current Societal Norms/policies
- Scope of Skills/Responsibilities
- Natural limitations/variability in cognitive resources/awareness
- The burden of increasing clinical/administrative demands
- Role specialization (and therefore less insight into other roles or expertise)

It is normal for us to:

- Not have all knowledge about a topic
- To want to use data that is readily available or easily accessible
- To be focused on our role-specific responsibilities and not aware of the roles/responsibilities of others
- To focus on resolving a specific problem (e.g. sepsis prediction), without considering how it might unintentionally harm a subgroup of individuals due to bias in data/measurement
- To want to follow shortcuts

The following questions will help stakeholders involved in purchasing or developing an AI solution, together with other relevant stakeholders (see Section 3.3 in Appendix) to evaluate population risk and impact in a way that will improve current practices and minimize population-level harm across several domains. This will allow teams to leverage the power of health AI to positively impact patients and providers and reduce healthcare gaps and inequities, rather than perpetuate or prolong them. These questions are best explored with patient advocacy/population health and medical area experts present or consulted. Given that bias in AI is unavoidable, this tool will also help organizations evaluate and prioritize bias mitigation efforts towards algorithms with greater risk and/or those that may be impacted by ethical/legal guidelines. Using this tool aims to improve current practices and minimize population-level harm. (Tool adapted to health-specific context in part from ethicstoolkit.ai)

Identify who will be impacted by the AI system:

Primary Impacted: Who or what may be or is directly impacted based on the objectives of the AI system? (e.g. patients, family caretakers, physicians, nursing, organization, business operations, etc.)

Secondary: Who or what may be or is impacted downstream based on those primarily impacted? (e.g. if physicians and their clinical workflows are primarily impacted, downstream effects may be experienced by nursing staff, or radiology technicians)

Unexpected/Unintended: Who or what may be impacted unexpectedly/unintentionally at the population or location level? Examples may include:

- Patients who do not speak English or their children
- Physicians working in community hospitals vs. academic medical centers
- Patients without insurance
- Acquired hospitals that use a different (non-integrated) electronic medical record system
- Members of a specific socio-demographic subgroup
- Individuals with visible or invisible disabilities

Se

ced care management services)

Select Degree: Minor Impact | Moderate Impact | Major Impact

Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations

Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Access to Direct Health Services/Healthcare:** Algorithms that impact who or how someone does/does not have access to necessary direct health care services (transportation coordination, medicine or health service approval, preventative care appointments, specialty care services, diagnostic testing, mental health screening, etc.)
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Emotional Health/Well Being:** These algorithms impact the emotional health or well-being of an individual or group. (Time waiting for health services/benefits, effort required to arrange for services)
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Life/Safety:** These algorithms directly impact individual or group safety or life (e.g. diagnostic, treatment, recommended treatments)
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Financial:** These algorithms impact the costs associated with healthcare for individuals, groups, or in specific areas. (e.g. health plan premiums, cost of care)
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Privacy:** These algorithms impact the privacy of personal health information for an individual or group.
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Trust:** These algorithms impact the trust that an individual or group may have in the healthcare system, clinician(s), or other healthcare professional.
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

- **Freedom/Agency/Rights:** These algorithms impact an individual's freedom/agency/rights as it pertains to their healthcare or health information.
  Select Degree: Minor Impact | Moderate Impact | Major Impact
  Select Scale: Small Proportion | Substantial Proportion OR Primarily one or more Vulnerable Subpopulations | Nearly Every Person OR Majority of one or more Vulnerable Subpopulations
  Select Direction: Positive Impact | Mostly Positive Impact | Mostly Negative Impact | Negative Impact

Is it possible that the degree or scale of impact could vary by context (population subgroup or location implemented).

- **No likelihood of systematic variation in scope of impact by context**
- **Small likelihood of systematic variation in scope of impact by context, but variability is due to known and validated clinical or social needs**
- **Small likelihood of systematic variation in scope of impact by context**
- **Medium likelihood of systematic variation in scope of impact by context, but variability is due to known and validated clinical/social needs**
- **Medium likelihood of systematic variation in scope of impact by context**
- **High likelihood of variation in scope of impact by context, but variability is due to known and validated clinical/social needs**
- **High likelihood of variation in scope of impact by context**

1.5 How to Interpret This Checklist

The checklist is designed not as a binary pass-fail assessment, but rather as a comprehensive tool to evaluate the risk-benefit profile of the AI solution and its associated system and to guide best practices across developer and implementer teams. Given the inherent complexity of each use case and implementation, a nuanced approach is essential. The checklist aims to facilitate transparency and furnish reviewers with substantial evidence, empowering relevant parties to make informed go/no-go decisions. Furthermore, it underscores the importance of additional measures that may be undertaken by the implementation or developer organization. These measures are crucial for preventing and mitigating adverse outcomes, as well as ensuring that the AI solution is employed judiciously in contexts where its limitations are acknowledged and respected.

Throughout the checklist, each evaluation criteria has received one or more coding tags in the left-hand column (example: LS1.U.C1.EC1). These identifiers are designed for traceability to the considerations in the Responsible AI Guide, and they are color-coded by principle area. Some evaluation criteria are based on considerations that space multiple principle areas or span multiple considerations within a principle area.

- Usefulness, Usability, Efficacy: (Principle Area Denoted with U)
- Fairness: (Principle Area Denoted with F)
- Safety: (Principle Area Denoted with S)
- Transparency, Intelligibility, and Accountability: (Principle Area Denoted with T)
- Privacy and Security: (Principle Area Denoted with PS)

(example: LS1.U.C1.EC1 would denote Lifecycle Stage 1, Usefulness, Usability, and Efficacy Principle Area, Consideration 1, Evaluation Criteria 1.)
Note: once the review of the checklist is complete, we'll be creating more streamlined, sequential tags. For now, the color coding will give you what's most important, as many evaluation criteria reflect overlaps in different principle-based considerations through the lifecycle.

2 Reporting Checklist

Columns and sections to be completed by the Reporter are denoted in dark blue and by Reviewer in light blue.

2.1 Clinical Risk & Population Impact Evaluation Summary

Clinical Risk and Population Impact Evaluation tools are provided in sections 1.4.2 and 1.4.3 respectively. ***Reporters*** *should provide a summary of clinical risk (including classification level) in Table 3 below, and a summary of population impact initially in the Planning Phase (Stage 1). If not completed during the Planning Phase* ***and*** *as insights are gained during subsequent Checkpoints, tools in sections 1.4.2 and 1.4.3 should be revisited and information in Table 3 should be updated.* ***Reviewers*** *should go over this information to gain context for the information that follows in the checklist (Section 2.3).*

Table 3. Clinical Risk and Population Impact Summaries

Clinical Risk Classification & Population Impact Summaries		
Domain		**Reporter Initials and Date**
Clinical Risk Classification & Summary		
Population Impact Summary		

2.2 Checklist Stage 1

Checklist: Stage 1 Define the Problem & Plan									
Criterion Number	**EC Identifier**	**Evaluation Criteria**	**Evidence, Explanation & Metadata/Document Code**	**N/A or Cannot Complete (CC): Describe in comment**	**Reporter Initials & Date**	**Evidence & Explanations Provided? (Yes/No/ Partial/NA)**	**Benefits, Limitations or Adverse Outcomes**	**Criteria Met? (Yes/No/ Partial/NA)**	**Reviewer Initials & Date**
	Planning Checkpoint								
PC.CR1	LS1.U.C1.EC1 LS1.PS.C2.EC2, LS1.PS.C2.EC3	Is there a clearly defined problem posed for an AI solution that is consistent with organizational goals, end user needs, and risk tolerances, thereby ensuring its role is appropriate, clearly defined, and understood?							
PC.CR2	LS1.U.C1.EC2	In its intended use, will the AI solution directly target the stated problem?							
PC.CR3	LS1.F.C1.EC1	Does the problem definition account for socio-demographic differences and avoid inherent disadvantages or discrimination against specific groups?							
PC.CR4	LS1.F.C1.EC2	Is the problem definition and its associated solution sufficiently inclusive to address a wide range of scenarios across comprehensive socio-demographic subgroups and not just a subset of the population?							
PC.CR5	LS1.PS.C2.EC1	Are the evaluation processes and documentation regarding the purpose of the proposed AI solution sufficiently robust and defined in relation to specific mission/business objectives, including identification of specific tasks and funding sources?							

PC.CR6	LS1.T.C1.EC2, LS1.T.C1.EC3, LS1.U.C1.EC3	Given the problem statement and organizational objectives, does the AI solution provide a clear improvement over existing standard or alternative practices, justifying need and appropriateness and has this reasoning been documented?								
PC.CR7	LS1.T.C1.EC4	Is the medical context (including whether diagnostic or prognostic) and rationale for developing or validating the AI model appropriate?								
PC.CR8	LS1.T.C2.EC1	Has the intended use and purpose of the AI solution within the clinical pathway, including its purpose and intended users (e.g., healthcare professionals, patients, public), been documented?								
PC.CR9	LS1.T.C2.EC2	Considering the model's purpose, has the outcome of the model been defined and documented, including details on how and when the outcome is assessed?								
PC.CR10	LS1.U.C3.EC5, LS1.U.C3.EC6	If applicable, has a cost-benefit analysis been performed and documented, accounting for the magnitude and frequency of the benefits, risks, and costs associated with internal development and implementation compared to external development and internal implementation?								
PC.CR11	LS1.F.C9.EC4	Is there sufficient time allocated to conduct both model performance evaluation and fairness assessment?								
PC.CR12	LS1.T.C4.EC3, LS1.T.C4.EC4	Have details on the data (e.g. evaluation, training data) and model performance been documented?								
PC.CR13	LS1.PS.C1.EC3	Are reasonably detailed data maps available that outline the data processing activities associated with the AI systems?								
PC.CR14	LS1.T.C4.EC2, LS1.F.C13.EC1, LS1.F.C13.EC2, LS1.F.C13.EC3	With specific attention to the risk of bias, are there clearly documented guidelines specifying how and when end users should utilize the AI solution, as well as criteria for when it is permissible or advisable to incorporate additional information into a decision, or to override the solution's output?								

PC.CR15	LS1.S.C2.EC1, LS1.S.C2.EC2, LS1.T.C5.EC1, LS1.T.C5.EC2	For the health AI solution being selected or developed, are there well-defined or standardized protocols in place and documented for determining patient population inclusion and exclusion criteria for training and application of the model, particularly in cases where certain populations may not fall under strict exclusion rules but where the validity of the AI solution may be compromised?								
PC.CR16	LS1.F.C3.EC1	Has a bias monitoring and mitigation strategy been defined, taking into account the AI solution's feasibility or effectiveness for various problem-relevant subgroups or end users based on workflow (e.g., language limitations, access constraints, insurance coverage, provider availability, patient load, etc.)?								
PC.CR17	LS1.F.C4.EC3	Could failure to consider socio-demographic subgroups potentially lead to harmful outcomes (at the patient or population level), or diminish the overall generalizability of the AI solution?								
PC.CR18	LS1.F.C3.EC2	Are security or mitigation measures implemented to safeguard relevant subgroups against intentional data contamination or model-based attacks?								
PC.CR19	LS1.F.C4.EC2	Is there a documented set of criteria outlining how AI fairness will be ensured across all socio-demographic subgroups?								
PC.CR20	LS1.T.C9.EC2	Does the healthcare organization have an established quality management system with which model development must comply?								
PC.CR21	LS1.T.C9.EC3	Have independent quality reviewers and auditors been identified, and has a method for reporting been established?								
PC.CR22	LS1.S.C4.EC1	Has an initial assessment been conducted to ensure compliance with federal rules and regulations, e.g. determining whether the health AI solution falls under the FDA's oversight (as guided by the FDA's Digital Health Policy Navigator), and establishing clear plans for adherence to								

		applicable local regulations?								
PC.CR23	LS1.S.C5.EC1, LS1.S.C5.EC2, LS1.T.C8.EC1, LS1.PS.C1.EC2	Have regulatory, ethical, and legal experts been contacted and consulted with to ensure compliance with requirements (safety, privacy, security, bias, transparency, etc.), to review existing/past related legal cases, and to facilitate planning for ethical and legal risks (regardless whether this is non-FDA approved vs. FDA approved AI model)?								
PC.CR24	LS1.T.C4.EC5, LS1.T.C11.EC3, LS1.T.C11.EC4	Have the terms and conditions for compliance with regulatory and ethical requirements, including exceptions and other related considerations been established and documented?								
PC.CR25	LS1.T.C8.EC3	Will IRB and FDA submissions be required for future applications of the model?								
PC.CR26	LS1.U.C5.EC1 LS1.F.C6.EC1 LS1.S.C1.EC1 LS1.S.C1.EC7 LS1.T.C12.EC1 LS1.T.C12.EC2	Are relevant stakeholders and end users involved in the AI solution's problem definition, articulation of the business need, AI solution selection process, and risk management planning process, have they commented on safety, security, and fairness related risks, and has their engagement and input been documented?								
PC.CR27	LS1.S.C1.EC2	Are safety and other risks actively identified by end users and relevant stakeholders for the current state and potential use of the AI solution to inform risk management practices?								
PC.CR28	LS1.S.C1.EC6, LS1.T.C8.EC2	Were human-centered design and human factors approaches employed throughout the current state analysis, selection process, and requirements gathering for the health AI solution?								
PC.CR29	LS1.PS.C1.EC2	Are designated personnel responsible for documenting and maintaining the inventory details of AI systems within the organization?								

PC.CR30	LS1.T.C9.EC1	Will the model be reported within a registry, inventory, or centralized data platform?								
PC.CR31	LS1.T.C4.EC1	Is there a documented overview of the model (i.e., who is developing the model, model date, model version, model type, citation details, license, etc.)?								
PC.CR32	LS1.T.C3.EC1	Has a format been identified (e.g., Model Card) in which project stakeholders, developers, end users, and patients can access documentation about the AI model and project-related information?								
PC.CR33	LS1.PS.C1.EC1	Is there complete documentation of AI solutions, including which models or systems are to be inventoried and attributes to be documented (such as documentation, links to source code, incident response plans, data dictionaries, and contact information for AI actors)?								
PC.CR34	LS1.F.C7.EC1, LS1.F.C7.EC2	Are there clearly pre-defined considerations, assumptions, or methods informing AI bias risk assessment and management for the relevant subgroups and have these been documented?								
PC.CR35	LS1.U.C3.EC1. LS1.U.C3.EC3, LS1.U.C3.EC4, LS1.T.C7.EC1, LS1.PS.C2.EC4	Is there a framework for evaluating the relative magnitude and frequency of benefits and risks to the AI solution, including identifying key performance indicators for impacts on patient care and end users, and has it been used to identify and document potential benefits and risks?								
PC.CR36	LS1.T.C1.EC5	Have key performance indicators been identified to justify the use of AI and measure the solution's impact on society and its ROI for the healthcare organization?								
PC.CR37	LS1.T.C10.EC1, LS1.T.C10.EC2 LS1.T.C11.EC1	Are goals and associated key performance indicators and success measures defined, quantifiable, and tracked, in a way that aligns with the AI solution's intended use and workflow (compared to non-AI workflow)?								
PC.CR38	LS1.T.C11.EC2	Are health and data standards (data provenance and diversity) defined?								

PC.CR39	LS1.F.C2.EC4, LS1.F.C2.EC1	Has the team explicitly defined fairness within the context of the problem being addressed by the AI solution and does the definition prioritize both minimizing harm and maximizing clinical access and benefits?								
PC.CR40	LS1.F.C2.EC2	Are there justifying criteria established for assessing AI fairness across sub-groups, such as equitable treatment alongside proportional representation, and parity in false positive and false negative rates?								
PC.CR41	LS1.F.C2.EC3	Has the team established a predefined approach for evaluating fairness based on the defined concept of fairness and the specific problem?								
PC.CR42	LS1.T.C6.EC1, LS1.T.C6.EC2	Has the team identified how risks will be evaluated, documented, and, if need be, made accessible as information for end users and/or patients?								
PC.CR43	LS1.S.C1.EC10, LS1.T.C9.EC4	Does the implementer organization have accessible standard operating procedures for risk management and safety reporting, ensuring that there is consistent monitoring and decision-making when the AI solution is deployed?								
PC.CR44	LS1.S.C1.EC3, LS1.S.C1.EC4, LS1.S.C1.EC5, LS1.S.C1.EC9	Is there a well-defined risk management process and plan in place that includes risk identification, assessment, and mitigation strategies, as well as Corrective and Preventative Actions (CAPAs), to address potential safety and other risks of each health AI solution for patients and end users?								
CR45	LS1.S.C1.EC8	Is there ongoing evaluation and consideration of defined human factors in the Return on Health / Return on Investment analyses of health AI solutions, guiding risk management decisions throughout the lifecycle?								
CR46	LS1.PS.C1.EC3	Are risk management processes in place for AI systems, as defined by privacy and cybersecurity policies?								

CR47	LS1.PS.C1.EC4	Have cybersecurity and privacy risk assessments been conducted on the AI systems?								
CR48	LS1.PS.C1.EC1	Does the organization demonstrate understanding of the privacy and cybersecurity risks of its AI system within the healthcare industry, aligning with its mission priorities and risk tolerances?								
CR49	LS1.PS.C3.EC1, LS1.PS.C3.EC2	Is there comprehensive documentation and clear rationale provided for the prioritization of and decisions made in response to identified privacy and security risks in the context of the proposed AI solution?								
CR50	LS1.U.C3.EC2	If the AI solution is intended for use in clinical decision-making, will its implementation lead to better outcomes than the current standard of care, considering the magnitude and frequency of benefits, risks, and costs?								
CR51	LS1.U.C5.EC2	Has the clinical validation success rate been measured against medical criteria?								
CR52	LS1.F.C8.EC1	Do the data use/sharing agreements of an externally acquired AI solution align with privacy and data security policies, such as Personally Identifiable Information (PII) and HIPAA?								
CR53	LS1.F.C9.EC1	Can the vendor provide documentation of bias evaluation steps taken, metrics, and outcomes, as aligned with the purchasing organization's AI/ML bias policies and relevant definition of fairness?								
CR54	LS1.F.C11.EC1	Does the vendor provide transparent, stepwise information on how the AI/ML system was developed and who developed it?								
CR55	LS1.F.C4.EC1	Is there a risk that the AI solution might amplify existing social inequalities or be retooled to cause inadvertent harm?								
CR56	LS1.F.C5.EC1, LS1.F.C5.EC5, LS1.F.C5.EC2	If applicable, are there systematic differences between the AI/ML training environment and the deployment context that could lead to bias or								

		disparities (e.g., workflow operations, treatment protocols, provider types, patient load, population representativeness, accessibility, data sources, and IT service integration)?								
CR57	LS1.F.C5.EC3	Has a plan been outlined for evaluating different potential sources of bias across various deployment sites?								
CR58	LS1.F.C5.EC4	If applicable, was the AI/ML model trained or tested in a setting similar to the context of its deployment?								
CR59	LS1.F.C9.EC2	Is the vendor willing and able to share model performance and parity information across relevant socio-demographic subgroups?								
CR60	LS1.F.C9.EC3	Is the vendor open to having a separate, third-party organization conduct a bias evaluation and provide results in accordance with the purchasing organization's needs, policies, and guidelines?								
CR61	LS1.F.C11.EC2, LS1.F.C11.EC3	Will the vendor disclose comprehensive information on the sources of data used to build the model, including demographic representation, representativeness with respect to the deployment context, and availability of relevant socio-demographic subgroup information to assess potential biases?								
CR62	LS1.F.C12.EC1, LS1.F.C12.EC2	Has the organization assessed how internal and/or vendor security practices may potentially expose models or data to external attacks and have procedures been defined for minimizing the scope and degree of impact? (e.g. attacks such as: data theft, data reconstruction, model altering, data altering, biased data distributions, or alterations to model attributes/functions, especially in a manner that might expose a specific subpopulation to greater risk of harm)								
CR63	LS1.F.C10.EC3	Can the vendor assure that patient data will not be utilized or shared to predict sensitive health information or identity data unrelated to the intended purpose of the tool?								

CR64	LS1.S.C3.EC2	Is the customer aware of any limitations in the underlying technology at the implementing site(s), and are they aware of any required alternatives or modifications to ensure patient safety, particularly when utilizing off-the-shelf (OTS) databases or similar components within the health AI solution?								
CR65	LS1.S.C3.EC1	Has an agreement been established early in the selection process, clearly delineating the responsibilities of the developer and implementer organizations and stakeholder involved throughout the entire lifecycle of the health AI solution, including aspects related to safety, effectiveness, and performance?								
CR66	LS1.T.C11.EC5	Has a joint plan been implemented between vendor and buyer to bring expectations into alignment with site-based goals, standards, terms and conditions?								
CR67	LS1.T.C1.EC1, LS1.U.C2.EC1	Has there been a thorough assessment of how the AI solution will integrate into existing workflows, and is this assessment and evidence of feasibility documented?								
CR68	LS1.U.C2.EC2, LS1.U.C2.EC3	With respect to integration of the AI system into the workflow, have things such as limiting impacts on flow of people or tasks in physical and digital environments, and limiting impacts on patient-clinician interactions been considered?								
CR69	LS1.T.C8.EC4	Will the integration of the model necessitate disclosures on the interface, or consent by end users and/or patients, as required by regulatory and legal standards?								
CR70	LS1.U.C4.EC3	Is there a documented assessment of the potential for confidence in the AI solution, weighing the risks and benefits, thus contributing to end users' trust in the model and its output?								
CR71	LS1.U.C4.EC4	Is there a defined pathway in place to address user concerns or lack of trust in subsequent iterations of the model?								

CR72	LS1.U.C4.EC1, LS1.T.C3.EC4	Is there transparent, accessible, and easily understandable information about the AI model that is made available to patients and end users to facilitate trust?								
CR73	LS1.T.C3.EC2, LS1.T.C3.EC3	Will documentation and transparency vary in scope, language, and specificity, based on who will be receiving the information (as appropriate), including end users and patients?								
CR74	LS1.T.C3.EC5	Are there pathways established for various stakeholders to equally access the project-related and model-related information?								
CR75	LS1.F.C10.EC1, LS1.F.C10.EC2, LS1.T.C7.EC3	Are patients informed about how their data will be used or shared and are there processes in place to uphold the privacy and security of patient data, ensuring compliance with use agreements and relevant privacy policies?								
CR76	LS1.T.C7.EC2	Can the patient opt out of the AI solution's use in their care?								
CR77	LS1.U.C4.EC2, LS1.T.C4.EC6	Is there clear, non-technical information available and documented on the limitations of the model for its intended use cases, both for internal use and for end-users to promote trust and transparency?								
CR78	LS1.S.C5.EC3	Is there a transparent process in place to inform patients about the use of AI in their care and request informed consent when applicable, ensuring coverage in case of adverse events or legal challenges?								
CR79	LS1.F.C7.EC3, LS1.PS.C3.EC3	Has a clear process been established to regularly update the framework or evidence used for risk and bias assessment across the lifecycle stages in light of new information or developments?								

2.3 Executive Summary of Anticipated Benefits, Risks, Adverse Outcomes, and Limitations

The **Reporter** should complete this section and provide an overall summary for reviewers based on responses to criteria above.

Executive Summary of Anticipated Benefits, Risks Adverse Outcomes, and Limitations

2.4 Summary of Findings

The **Reviewer** should complete this section and provide an overall summary of findings based on responses, summary, and evidence provided by the Reporter.

Reviewer Summary of Findings

2.5 Evidence & Explanation Metadata:

*This section should be completed by **Reporters** to list all attached evidence documents and track the source of evidence and explanations listed in the checklist. **Providers of Evidence** include any stakeholders who provided documentation and evidence to the Reporter (See Appendix Section 3.3 for a non-exhaustive list of potential stakeholders that may be involved in providing evidence for various criteria.) The first line is an illustrative example of use.*

Evidence & Explanation Metadata				
Evidence Document Code	**Reporter Name and Role**	**Provider of Evidence Name(s), Title, Role, & Contact Information**	**Description**	**Evidence Archive Location**
E.g. <DataPlan.v1.2>	*<Enter Reporter Name, VP of Quality>*	*<Enter Name, Data Engineer, email@email.com>*	*Data Management Plan*	*<Link to Document Attachment or Location>*

3 Appendix

3.1 Link to Traceability Matrix

https://docs.google.com/spreadsheets/d/15cJEerA861o3cSV-rzL8n0H_X-65orTBk4uuybdTByg/edit?usp=sharing

3.2 Terms Defined

AI model: A conceptual or mathematical representation of phenomena captured as a system of events, features, or processes. In computationally-based models used in AI, phenomena are often abstracted for mathematical representation, which means that characteristics that cannot be represented mathematically may not be captured in the model. Often used synonymously with “algorithm,” though it may be conceptually distinct, prior to the transformation of inputs to outputs.

AI solution: A shorthand for the AI model or algorithm and required technical infrastructure (hardware, software, data warehousing, etc.).

AI system: A fully operational AI use case, including the model, technical infrastructure, and personnel in the workflow.

3.3 Representative roles in health AI industry

The roles of the developer vs. implementer organizations are unique to each AI solution and may vary throughout the lifecycle.

Stakeholder Roles	Example Stakeholder Professions	Representative Organizations
Data Science Developer	Data Scientists, Data Engineers, Data Analysts & Storytellers, Machine Learning Engineers, Product Managers	Academic Medical Centers Community Health systems Vendors Expert Consultants
Informatics and Information Technology	Biomedical Researchers and Informaticists, Software Developers, Front-End Engineers, Support Engineers, Data engineers, Quality Assurance Analysts, Security & Compliance Experts	
Design and Implementation Experts	Implementation Scientists, Human Factors Experts, User Experience Designers, Patient Safety Experts, Clinicians	
End Users	Health Care Providers (e.g. Clinicians and Nurses), Insurers and Payers, Healthcare Operations Workers, Patients and Caregivers	Health Systems such as: Academic Medical Centers Community Health Systems Integrated Healthcare Systems Primary Care Networks Urgent Care Networks
Health System Administration	Health Systems Leadership, Contract Administrators, Vendor Management Specialists	

Clinical Administration	Lab Managers, Nursing Managers, Other Clinical Decision-Makers	Independent Imaging Centers Providers in Private Practice
Impacted Groups	Patients and Caregivers, Patient Advocates	Patient Advocacy Organizations Patient Advisory Boards
Ethics and Regulation & Standards Organizations	Bioethicists, IRB Analysts, IRB Members and Leaders, Lawyers and Legal Advisors, Civil Servants, NGO Decisionmakers, Policy Analysts, Regulatory Experts and Consultants	Federal Government Local Government NGOs Law Firms Standards Organizations Medical and Nursing Societies Medical Licensing Bodies Medical Device Collaboratives, etc.

Table 1: Stakeholder Roles, Professions, and Representative Organizations. Derived from CHAI Responsible AI Guide (Link)

3.4 Example User Personas and Scenarios for Development, Procurement, and Implementation

Example 1:
Scenario: A health system or healthcare organization (e.g. payer, EHR company) that has internal developer and implementer teams and are looking to develop a model to predict risk of post-op complications.
Example Reporter(s): Chief quality officer is assigned the role of Reporter and project lead and contacts relevant stakeholders who will serve as Providers of Evidence (as appropriate) from the organization (e.g. data, informatics & security, policy/legal, human factors or social & behavioral sciences, clinical area expert, patient advocate). Ideally these individuals work together to complete the planning phase tasks and set a roadmap for the responsible AI checklist tasks and processes. When the model is ready to be piloted, teams and stakeholders will provide evidence to the Reporter for Responsible AI Checkpoint 1.

Example Reviewer(s): The Vice President of Quality reviews the evidence and makes a go-no-go decision about moving the project forward to piloting. If no-go decision is made, it may be because modifications and further evidence are required, at which point the AI solution undergoes further iteration. If a go decision is made, the project moves forward to piloting, with relevant stakeholders involved in gathering evidence for the next Responsible AI Checkpoint.

The Reporter and Reviewer for subsequent checkpoints may differ as appropriate for the success of the project and as determined based on expertise required.

Example 2:
Scenario: Health system or healthcare organization purchasing/acquiring an AI solution from an external developer team to assist with imaging diagnostics (mammography), with an internal implementation team.
Example Reporter(s): The Chief Medical Officer assigned the role of Reporter from the implementing/purchasing organization to work alongside relevant stakeholders (radiologists, radiology technicians, IT and security, patient privacy) to gather evidence on internal needs, processes, and capabilities to help guide the purchasing decision and design the broader AI system (e.g. end user engagement, operations, security and privacy capabilities, integration capabilities). They also work alongside the developer organization who assigns the Informatics Lead and Product Lead for the AI solution as Reporters from their respective organization, to address some of the Planning Checkpoint items and to gather evidence for best practice criteria in Responsible AI Checkpoint 1.
Example Reviewer(s): The procurement team may assign an internal reviewer (or consult with an external individual if further expertise is required), to review the evidence provided by the developer organization to help make a go-no go decision about purchasing. They may gather information from several potential vendors and use this checkpoint as a way of comparing vendor offerings, model performance, integration capabilities, transparency, privacy/security, etc. to guide the decision around which vendor to purchase from. The reviewer may instead choose to use this checkpoint as a way to select two vendors from which to pilot an AI solution internally, prior to making final purchase decisions. Once the decision to purchase or pilot is made, the implementing/purchasing organization may assign another reporter from the implementer team to help guide the initial pilot (which may lead to another go-no-go decision), or guide a small scale implementation process. Internal implementer and external developer teams will likely continue to collaborate to help troubleshoot problems that may arise during Responsible AI Checkpoint 2 and/or Responsible AI Checkpoint 3.

Additional Notes:

Developer organizations may choose to use the planning and other checkpoint checklists to help guide their development and piloting process, to help prepare for regulatory evaluation, and/or have external expert organizations review or validate the evidence they have provided. They may also choose to summarize the best practice evidence for respective checkpoints to share with potential clients, fostering transparency and trust.

In some cases, such as small community clinics or private practice settings, access to the full list of individuals required for an internal implementation or development team may not be available. In these cases these organizations may look for vendors who are already using best practices or who are willing to be transparent about their development process as outlined in the respective checklists. They may also choose to consult with external experts to help guide them through the purchasing and review processes in a way that is aligned with best practices and criteria defined here.