COALITION FOR HEALTH AI

# Use Case Best Practice Guide

*Electronic Health Record (EHR) Information Retrieval*

# TABLE OF CONTENTS

# EHR Information Retrieval

*For Developers & Implementers*

## What Does This Guide Include?

Use case specific best practice guides provide high level industry and consensus defined insights and recommendations, for the application of responsible AI principles to a specific use case. This guide focuses on an AI-enabled EHR information retrieval solution. The guide is organized by role (developer/implementer) and responsible AI principle areas where applicable (see figure below).

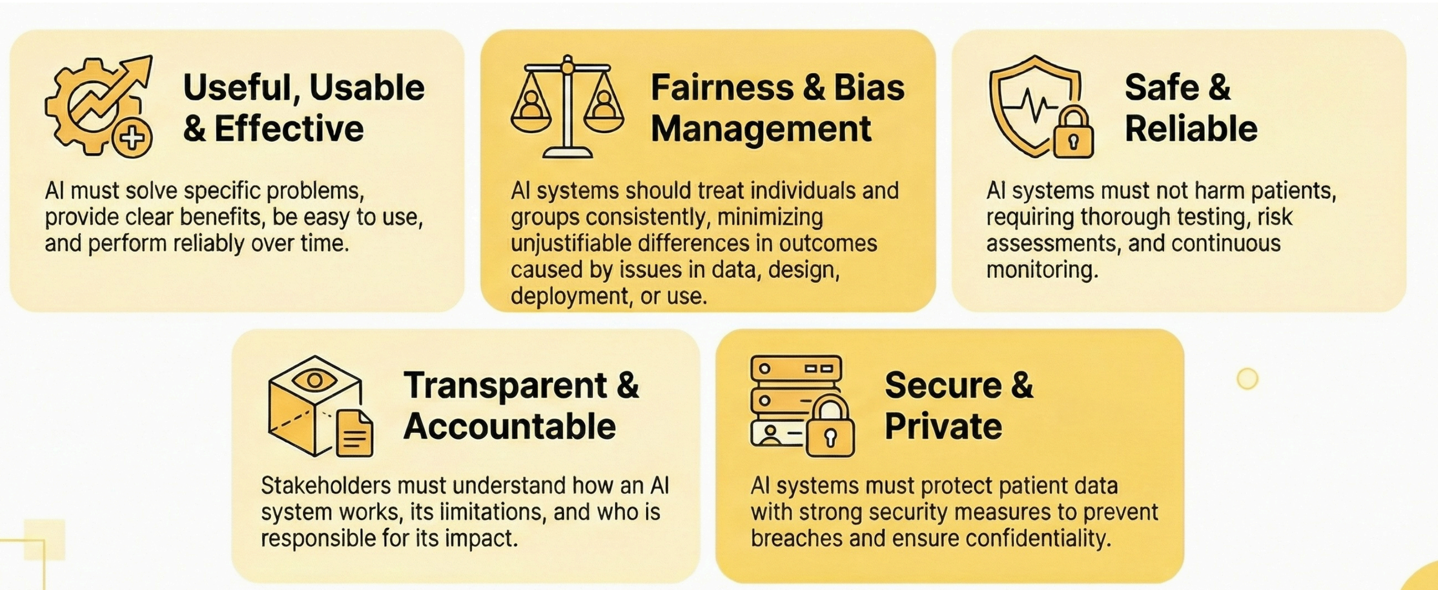


## Use Case Description

To arrive at the best practice statements, work group members grounded in a specific use case. The AI-enabled EHR Information Retrieval use case outlines a workflow for retrieving, standardizing, and processing diverse data formats from Electronic Health Records (EHRs) and related databases. The goal is to transform structured and unstructured healthcare data into a standardized format for further analysis and decision-making. The effectiveness of downstream tasks such as summarization and inference depends on properly cleansing and structuring raw data before it is processed. Ensuring data quality at this stage is critical for meaningful analysis.

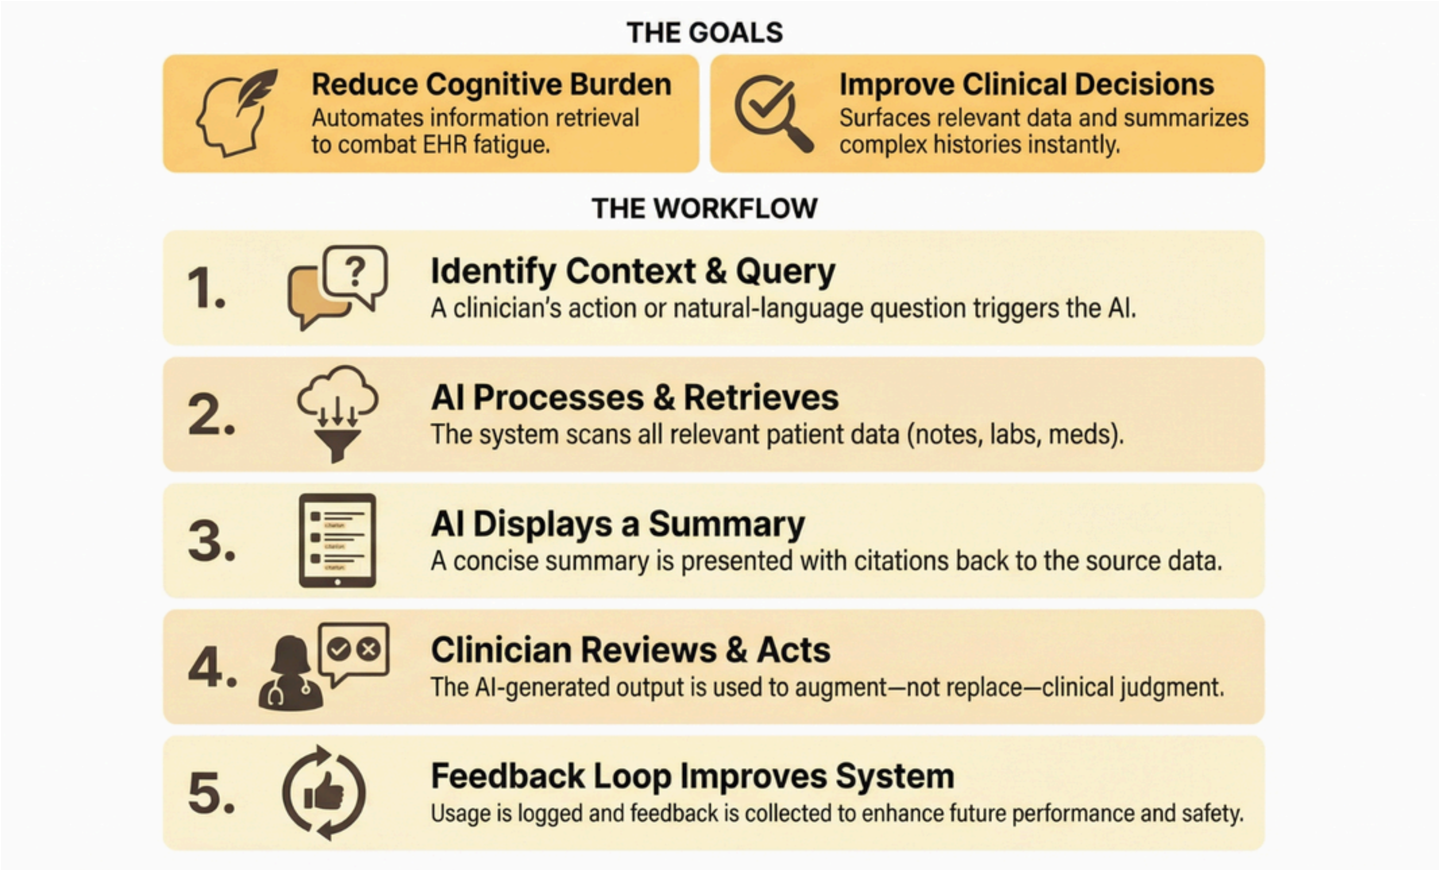


**Note**

As you make use of this guide, note that references to "clinical teams" or "clinical specialties" may be interchangeable with "user teams", or "user types" for EHR information retrieval solutions that apply to use-case applications outside of clinical contexts.

This guide focuses largely on clinician-initiated workflows within the electronic health record (EHR). While core clinical care may benefit from EHR-native solutions, there are other high-value workflows, such as utilization review and administrative tasks, that may be better served through purpose-built vendor platforms. These external systems may offer more specialized functionality, have improved purpose-specific ROI, and be designed for the needs of target end-users. When evaluating and deciding which solutions are a good fit for your team and the problem that is being solved, consider whether the use-case is better addressed by an external point solution or by EHR-native solutions. In addition to evaluating things like fit for use, efficacy, safety, etc., consider engaging downstream users/stakeholders early in this process to help with change management and to identify potential workflow needs and issues that might influence decisions. Non-clinician users may include care coordination, administrative, and operational staff, among others, depending on the use case.

Additionally, there are EHR information retrieval solutions that do not require workflows to be initiated by a user and instead depend on near real-time data feeds and use "always-on" virtual assistants that proactively surface insights at key parts of the workflow. These solutions can be especially valuable at key decision points, providing real-time insights that can help users make informed decisions about next steps quickly.

## Who is This Guide For?

**Developers:** individuals involved in the software development process, including requirements gathering, design, coding, testing, and maintenance of software applications (derived from IEEE, 12207:2017)

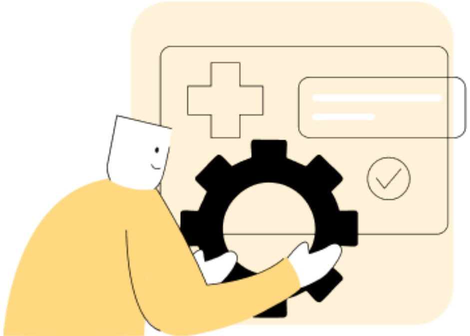

**Implementers:** individual(s) responsible for the procurement, deployment, and/or overall realization of a system or component in accordance with a specified design (derived from IEEE 829 and IEEE 730)

## Listening In: A Summary of Challenges & Insights

Below is a summary of some of the challenges and take aways that emerged from the work group conversations.

### Challenge # 1: There is heterogeneity across EHR systems.

- Clinical concepts are represented inconsistently across vendors and even across instances of the same system (e.g., different labels for the same note types).
- Local customizations, undocumented code sets, and inconsistent use of standards (e.g., LOINC, SNOMED) hinder reliable retrieval.

### Challenge # 2: There are data completeness and continuity issues.

- Patients with fragmented or short-term care histories (e.g., ED-only visits) often have sparse records, leading to gaps in retrieval and inaccurate AI summaries or risk scores.
- AI models may misinterpret or hallucinate information due to poor input data quality.

### Challenge # 3: There are transparency and provenance gaps.

- Clinicians lack tools to verify AI-generated outputs due to missing or opaque data lineage and provenance tracking.
- Retrieval systems often reconcile structured and unstructured data silently, which masks discrepancies.

## Challenge # 4: There are bias risks.

- Systematic under-representation of certain populations (e.g., pediatric, non-English speakers) or record types (e.g., narrative notes vs structured data) introduces bias into both retrieval and summarization.
- Inconsistent mapping across EHRs may skew outputs toward more common or dominant data formats.

## Challenge # 5: There are workflow alignment challenges and adoption barriers.

- If retrieval tools are not integrated into existing clinical, operational, or administrative workflows, users are less likely to adopt them.
- Output formats that are too long, unstructured, or specialty-inappropriate reduce usefulness and slow decision-making.

## Challenge # 6: There is a need for feedback loops and monitoring.

- Many systems lack lightweight mechanisms for clinicians to flag errors or omissions.
- Insufficient monitoring of retrieval outputs and system performance over time leads to blind spots in system quality and bias management.

### Key Takeaways

- **Start with high-quality, normalized use cases:** Pilot deployments in clinical domains with more standardized and complete data to reduce early errors and build trust.
- **Invest in semantic and structural interoperability:** Use a combination of structured search, NLP, and ontology mapping (e.g., SNOMED, RxNorm) to maximize retrieval accuracy across diverse EHRs. Track unrecognized or mismatched data for ongoing quality improvement and bias mitigation.
- **Make discrepancies transparent:** Design systems to surface, not hide, inconsistencies between data sources (e.g., conflicting meds in notes vs. med lists). Allow users to see provenance at the sentence level and verify key claims.
- **Center clinicians in evaluation:** Conduct usability testing across specialties and ensure outputs meet their time and cognitive load constraints. Use standard evaluation metrics (e.g., SUS, accuracy, recall time) and include clinicians in validation loops.
- **Ensure ongoing monitoring and governance:** Implement (semi-)automated quality checks, track data drift and performance over time, and document version control. Establish clear governance policies around error handling, feedback incorporation, and data quality monitoring.

# Best Practices for Developers

## Usefulness, Usability, & Efficacy

### Build a Reliable Retrieval Foundation

- Base retrieval on the most reliable field for the given use case—rather than defaulting to coded elements—and allow synonyms or alternate labels at query time to improve relevance.

**Example**

*When clinicians search for "cardiology consult," the retrieval system should not rely only on document type codes (e.g., LOINC or local IDs), which may be incomplete or inconsistently applied. Instead, it should weigh the document title (e.g., "Cardiology Consultation"), free-text headers, and coded metadata, selecting whichever source is most trustworthy for that scenario.*

- Use ontology mapping (e.g., SNOMED, LOINC, RxNorm) and semantic search so that clinicians can ask questions in their own terms without needing to understand codes.

### Optimize for Specialty-Specific Needs

- Design retrieval and summarization pipelines that efficiently process large patient records by segmenting or prioritizing relevant sections based on the intended use case. Rather than prescribing one method (e.g., section-by-section summarization), focus on achieving concise, context-appropriate summaries that preserve accuracy, relevance, and clinical value.
- Fine-tune retrieval and summarization models on domain-relevant datasets and ontologies (e.g., oncology staging terms, trauma care guidelines) to ensure accurate, high-value information retrieval.

### Adapt and Refine Through Clinician/User Feedback

- Design AI-generated summaries that are configurable by clinical team type (e.g., palliative care, cardiac, dietary), or user-type based on iterative user feedback.

## Fairness & Bias Management

**Harmonize Data and Monitor for Bias**

- Ensure retrieval systems can map equivalent entities across heterogeneous EHR implementations, which refers to the diverse ways different EHR vendors or instances represent the same clinical concepts.

**Example**

*One EHR may label a note as "History & Physical," while another calls the same type of documentation a "Progress Note." A schema-aware system identifies these as equivalent so clinicians see consistent retrieval results. Where alignment is not possible, implement error logging to capture instances where data cannot be reconciled, and integrate bias detection routines to check whether certain populations (e.g., pediatric vs. adult records) or record types (e.g., imaging, labs, narrative notes) are systematically under-represented or misclassified.*

- Incorporate frequency and continuity of care as features in model development to flag and mitigate representation gaps; continuity of care can be a problem when a health system is working with EHR data, especially if patients tap into the health system temporarily and then go elsewhere.

**Note**

Not all bias is due to population or subgroup-level differences and not all bias is harmful. Bias can be:

- A statistical reality in the data that might require supplementation or methodological considerations
  - Example: certain health conditions are found at different rates in the general population compared to specific subgroups of the population due to a range of factors.
- Intentional
  - Example: A solution specifically designed for care management for older adults may not generalize when used on younger adults.
- Due to contextual factors that impact the quality of experience or method of use
  - Example: differences in retrieval by type of EHR or automation bias resulting from overly confident sounding outputs.

It is also not feasible to conduct extensive evaluations of all subgroups or contexts for solutions, however identifying skews in data, variability in performance, and variability in use, for a particular use case, and being transparent about what is known and unknown, is important.

## Safety & Reliability

**Prepare and Validate Data Befrore Training and Inference**

- Implement model input validation and pre-processing pipelines to remove low-quality records and define thresholds for acceptable data completeness; this means developers must pre-screen training and inference data to detect gaps (e.g., missing notes, labs, or demographics) that could compromise model reliability. Removal or flagging of such records supports safe, meaningful model outputs and can reduce unintended behaviors from poor inputs.
- Train retrieval models using datasets that are enriched with semantic and syntactic representation (e.g., ontologies, common data models). Properly structured training data reduces misclassification risk and improves reliability of downstream inference tasks.

**Design Robust Retrieval Logic for Comprehensive Clinical Search**

- Use a mix of search methods: structured search for coded data and AI/NLP for free-text notes. This helps give complete answers when clinicians ask for everything of a certain type (e.g., all allergies, all past echocardiograms). Add simple filters for common items like medications, labs, or allergies so nothing is left out, even if recorded in different ways.
- Use checklist-based validation (e.g., critical event detection) and entity-matching to ensure that no essential data points are missed. Critical event detection refers to systems that monitor patient

data, such as, vital signs, lab results. or clinical notes to automatically identify signs of potential deterioration or serious health incidents, enabling timely clinical interventions;
    - Entity-matching is the data process of identifying and linking records that refer to the same real-world entity (e.g., a patient, medication, or clinical concept) across disparate sources, despite variations in representation such as misspellings, formats, or identifiers.

**Continuously Evaluate Model Quality Using Clinical and Benchmark Testing**

- Run (semi-)automated benchmarking combined with clinical evaluations (e.g., Likert scoring of accuracy and readability) after every major model update.
- Establish evaluation protocols for hallucination (i.e., the generation of information not supported by source data) and factual accuracy (i.e., the degree to which outputs correctly reflect verifiable information). Testing may include benchmarking against curated test sets with known ground truth or conducting an error analysis to identify common failure modes (e.g., confusing medical terms, inventing lab values).
- Build retrieval systems so they automatically log how searches are used and what results are returned, creating a record that can be reviewed and annotated later. Require minimum inter-annotator agreement (e.g., Fleiss' Kappa $\geqslant 0.7$) before deployment.

**Flag uncertainty and escalate low-confidence results for review**

- Create confidence scoring and flag low-confidence retrieval results for human review. This ensures that errors at the data processing stage do not propagate downstream into summarization or inference tasks, where missing information could undermine decision-making.

**Promote Interoperability by Following and Improving Standards**

- Stay up to date on interoperability standards such as HL7 FHIR and USCDI, which define how health information should be structured and shared. Where these standards are incomplete or unclear, such as how to handle embedded images or mixed media in clinical notes, developers should propose clarifications or extensions back to the standards community.

### In Practice

This could look like:

1) Schema enforcement: build integration layers that validate whether incoming data matches expected formats, and reject or flag data that doesn't

or

(2) Advocacy: share findings with EHR vendors and standards bodies to help strengthen guidelines for consistent implementation.

## Transparency

**Start with Transparent Data Ingestion and Preprocessing**

- Design adaptable data ingestion and parsing pipelines (e.g., CDA → FHIR → JSON) that improve data consistency while preserving transparency about any transformations applied before model use. While large language models can process unstructured text effectively, maintaining clear provenance and documentation of preprocessing steps remains essential for trust.
- Demonstrate data lineage and semantic traceability, such as code mappings or provenance tags. This ensures safe deployment by allowing users to track how raw EHR data was transformed before being used in downstream workflows.

**Examples**

Code Mapping: "The local code H123 (internal homelessness indicator) was mapped to ICD-10 Z59.0 (Homelessness) using the developer's terminology service. This mapping is documented and version-controlled."

Provenance Tag: Metadata like "Derived from note authored by social worker on 03/15/25, mapped via NLP to SNOMED concept X."

#### Design Retrieval Systems that Expose Inconsistencies

- Design retrieval systems to first identify and surface data inconsistencies, such as mismatches between a medication list and a clinician's note, so users understand the source of the conflict and developers can assess the error type (i.e., a "hallucination" or other type of error).
  - Make it clear when issues arise from the underlying EHR data rather than from the AI model itself, helping clinicians interpret results with appropriate caution.
- Design retrieval models to surface discrepancies explicitly rather than silently reconciling them. For example, if structured and unstructured data disagree, flag both to the user with clear indicators (instead of making hidden inferences).

#### Ensure Citation and Trust Through Provenance in Summarization

- Design AI systems that extract from both structured and unstructured data (progress notes, documents) and display data provenance; this supports the value of multi-source extraction and explicit citation (i.e., provenance) to ensure clinicians trust and verify the AI summary content.
- When possible, design summarization pipelines with sentence-level provenance linking each statement to its original source note for transparent verification.

# Best Practices for Implementers

## Usefulness, Usability, & Efficacy

**Build Organizational Capacity and Expertise**

- Build capacity for public health and community health organizations to recruit, retain, and support clinical informaticists and data stewards to participate in AI governance, development, and implementation

**Note**
*"Build capacity" could mean investing in training, hiring, or partnering to ensure community health organizations have access to clinical informatics expertise — e.g., shared informatics resources, funded fellowships, academic-public health partnerships, or integrating informatics roles into grant-supported infrastructure*

- Collaborate with institutional education and training teams to strengthen frontline staff knowledge in data standards, terminologies, and representation

**Select and Customize Tools for Clinical Relevance**

- When procuring solutions, request vendors demonstrate how their system handles inconsistent coding practices and unstructured note types. Include usability testing with clinicians to validate that discrepancies are transparent and manageable.
- If adopting specialty-specific AI solutions, provide configurable summary templates and filters (e.g., "oncology-focused view" or "trauma events view") within the EHR retrieval interface to allow users to tailor the retrieved data and summaries to their specialty workflows. Ensure configureability is balanced with simplicity in the UI for the end user.
- Embed context-aware retrieval logic that can adjust relevance rankings based on patient-specific or role-specific context (e.g., primary care vs. cardiology). Where possible, allow implementers to configure relevance criteria for their setting.

### Integrate Seamlessly into Clinical Workflows

- Integrate AI summaries directly into existing clinical workflows  to reduce friction and encourage adoption. This integration minimizes the need for clinicians to switch between systems or perform additional logins or navigation. By making the AI summaries natively accessible within the EHR, implementers ensure that the technology supports clinicians in their natural workflow.

### Pilot, Evaluate, and Benchmark Responsibly

- Pilot retrieval solutions with representative end users to ensure the length, format, and presentation of retrieved content fit their workflow and support safe, timely decision-making.

> **Example**
>
> *Before wide deployment, test the tool with clinicians across specialties (e.g., hospitalists, cardiologists, nurses). Measure how long, or how many clicks it takes them to find key information with AI retrieved results versus workflow as usual.*

- Set evaluation standards based on non-inferiority to existing workflows, rather than perfection. Benchmark against current EHR search and patient interview workflows to ensure solutions are at least "no worse" and ideally more efficient.

### Embed Feedback Loops for Continuous Improvement

- Deploy EHR information retrieval solutions with built-in clinician feedback loops. Enable users to flag errors, omissions, or irrelevant outputs, and ensure these inputs feed into iterative system updates. If purchasing from a third-party vendor, request a similar clinician feedback loop functionality.
- Incorporate lightweight feedback mechanisms (e.g., flagging or quick edits) within the AI solution to track changes for model refinement.

## Fairness & Bias Management

### Monitor Retrieval Performance for Bias and Coverage Gaps

- Set up regular reviews of retrieval outputs. These reviews should check whether certain patient groups, clinical specialties, or EHR setups (such as Epic, Oracle/Cerner, or Allscripts instances with different local customizations, code sets, or document structures) are missing information or poorly represented. Retrieval outputs are the information that an AI system pulls from the EHR and presents back to users (e.g., patient summaries, medication lists, lab results, or imaging reports).
- Log unrecognized or mismatched data during retrieval and periodically review those logs (e.g., by patient group, department, or facility) to spot systematic gaps or inequities that could bias results.

**Example**

*If a retrieval system cannot parse certain lab results because one clinic stores them under custom codes rather than LOINC, each failed match should be recorded with metadata such as the source site, data type, and (if available) patient demographics.*

## Safety & Reliability

**Maintain Data Quality & Input Monitoring**

- Deploy (semi-)automated quality checks for identifying incorrect or incomplete data. Verify schema structure and required fields like patient_id, date_of_birth, etc.
- Establish a process to track expected data inputs, such as the anticipated counts and types of structured elements (labs, medications, diagnoses) and unstructured documents (notes, scanned forms). Use spot checks, dashboards, or threshold alerts to identify anomalies before AI-supported use.
- Scope early deployments to use cases with relatively normalized or high-quality data. For example, start with domains with consistently reliable data before expanding.
- Work with EHR vendors to enforce stricter code set usage, data schema alignment, and documentation practices. Prioritize semantic interoperability, not just data transport.

**Examples**

The Kahn Data Quality Framework (Harmonized Data Quality Assessment Terminology & Framework), and associated extensions provide terminology, methods, and metrics to help manage data quality in single, reusable data assets. See the accompanying EHR information retrieval testing and evaluation framework (T&E) for a mapping of some of the accuracy and bias metrics to Kahn framework.

There is also an increasing need for methods and metrics assessing data quality in federated networks with multiple data assets that require compatibility with each other (see link for a comparison of data quality between centralized and federated approaches to data quality) . It is important to know which factors may impact the correct use of data assets in a federated learning environment. Examples of such factors might include things like the uniqueness of concept(s) to a specific data asset.

**Monitor Model Performance**

- Monitor performance degradation due to data inconsistencies (e.g., missing fields, low-quality notes, or unstandardized terminologies) and adjust use accordingly. Treat data quality as a dynamic risk—adjust thresholds or retrain models as needed.

- Maintain version control and audit trails for AI performance to track regressions and unexpected behavior. Supports transparency and accountability across updates.

**Establish Governance, Safety, and Escalation Protocols**

- Establish clear governance policies on error handling and clinician notifications. Define when and how retrospective corrections will be surfaced, alerts triggered, and re-contact with patients initiated.

## Transparency

**Evaluate Trasparency & Accuracy Before Adoption**

- Request vendors disclose how their systems perform on *factuality benchmarks* (e.g., percentage of outputs that match verified source data) and how hallucination risks are mitigated. This sets the foundation for trust before deploying AI solutions.

**Flag & Contextualize Incomplete or Risk-Prone Outputs**

- When interpreting AI outputs, flag risk scores derived from patients with incomplete data or limited engagement history. Incomplete data can skew outputs, especially for “temporary” patients. Systems should surface these limitations clearly.

**Enable Source Verification at the Point of Use**

- Provide user-interface (UI) elements (e.g., clickable links or highlighting) to allow users to quickly trace summary statements back to original records. This supports real-time verification and promotes clinician trust

## Security & Privacy

- For Implementers, use standards-based protocols with strong authentication and access control to ensure retrieval logic is consistent across systems and not dependent on a single vendor’s proprietary interface.

### Examples

*Examples for how to use standards-based protocols include: adopting emerging interoperability approaches such as the Model Context Protocol (MCP) or Agent-to-Agent (A2A) frameworks. When configuring AI-supported retrieval features, prefer calling standardized, protocol-compliant endpoints rather than hard-coding to a vendor-specific API. Such implementations should also align with established security and privacy standards (e.g., ISO/IEC 27001 for information security, ISO/IEC 27701 for privacy information management, and ISO 27799 for health data protection) and follow U.S. Health IT test method requirements for authentication, access control, and authorization*


# Best Practices for Developers & Implementers

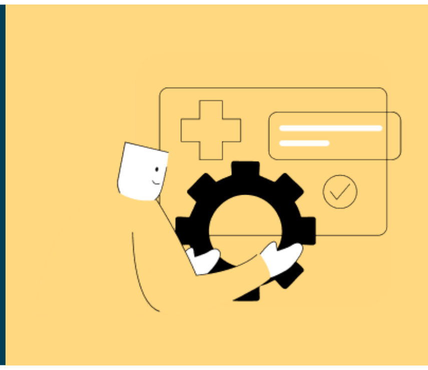

## Safety & Reliability

**Ensure research data models align with interoperability standards and address gaps with translation strategies.**

- For Implementers and Developers, proactively assess and document how research data models (e.g., OMOP, PCORnet) align with regulatory interoperability frameworks such as USCDI and HTI-1. Where misalignments exist, define mitigation strategies and interim solutions to ensure AI systems remain viable and compliant as standards evolve.

### Example

If a model trained on OMOP data lacks mappings to USCDI-defined elements such as clinical notes or social determinants of health, implementers and developers should collaborate and establish a translation layer or mapping protocol to maintain interoperability and support real-world deployment.

- PCORnet
- USCDI v3
- OMOP
- HTI-1

Transparency

**Know that Mapping Local Codes to Standards Take Time**

- For Developers and Implementers, during implementation phase, build in time for mapping local/proprietary codes to standard terminologies. It is important for Developers to be transparent in what type of information is important to have standardized in order to help Implementers focus their own standardization efforts.

# Appendix 1: Consensus Method

Best practice statements are collected from work group presentations and discussions. To ensure alignment across stakeholders, CHAI uses a three-phase consensus process for Best Practice Statements (BPS) generated through Work Group activities:

**Phase 1: Initial Consensus Check**

- **Purpose**: Gauge initial agreement on each draft BPS.
- **Voting Options**:
- *Include / Include Contextually / Exclude / Abstain*
- **Decision Rules**:
  - If $\geq 2/3$ vote "Include" → **Consensus achieved** (no further action).
  - If $< 2/3$ "Include", but $\geq 2/3$ combined "Include" + "Include Contextually" → **Flagged for Phase 2**.
  - If $\geq 25\%$ vote "Exclude" → **Flagged for Phase 3**.
  - If $\geq 2/3$ vote "Exclude" → **Automatically excluded**

**Phase 2: Revote with Revisions**

- **Purpose**: Re-evaluate BPS that did not reach consensus in Phase 1 (but had $< 25\%$ "Exclude").
- **Format**: Original and revised BPS shown side-by-side (based on Phase 1 feedback).
- **Voting Options**: *Include / Exclude / Abstain*, with an optional comment field.
- **Outcome**: Results used to determine final inclusion or exclusion.

**Phase 3: Live Discussion and Vote**

- **Purpose**: Address BPS with $\geq 25\%$ "Exclude" in Phase 1 (strong disagreement).
- **Steps**:
  - Facilitated group discussion of flagged BPS.
  - Live revote during the meeting + optional 1-week offline voting.
- **Voting Options**: *Include / Exclude / Abstain*
- **Outcome**: Final decision made based on discussion and revote results.

# Appendix 2: Thank You and Contributors

We want to thank every individual who showed interest, participated, listened, and came along with us in the early stages of our work. CHAI is at its core, a convener, and a member-driven non-profit. We are so grateful to be on this journey with you towards responsible AI in health for all. Your experiences, your feedback, your contributions, all make us who we are and help bring us to where we need to be.

For those who want to be credited directly by name, please reach out to us at program-management@chai.org to request contribution credit for the DTC: General Health Advice Chatbot Work Group. Below is a list of organizations who had at least one (many times more) individuals who showed interest and/or participated in this work group.

If you want to learn more about our work groups (current and future), or have feedback on CHAI work groups, products, or services, please contact our Director of Responsible AI: merage@chai.org and Program Manager: anthony@chai.org.

**Workgroup Leads**

| Name | Organization |
|---|---|
| Erin Palm | Infinitus Systems |
| Paru Srinivasan | MCG Health |
| Vince Hartman | Abstractive Health |
| Aaron Boussina | University of California San Diego |
| Jackie Martinez | Regard |
| Sara Murray | University of California San Francisco |
| Shalmali Joshi | Columbia |
| Micah Sheller | Intel |
| Rebecca Mishuris | Mass General Brigham |

**Participating Organizations**

Individual Contributor, Dorcas Yao, MD

101 GenAI

Abstractive Health, Vince Hartman

AdventHealth

Airia

Alameda Health System

Altarum

Ambience Healthcare

Amruta Inc

Apricity Innovations

Ascension

athenahealth

Atria Health

Atropos Health

Belmont University College of Pharmacy, Jarod Parrish, Pharm.D.

Boston Children's

Brooke & Associates

Clinia, Daniel Buades Marcos

Cognome

Columbia University Irving Medical Center

Darena Solutions/Prompt Opinion

Department of Health

Duke

DxTx

Elsevier

Ensemble Health Partners

Essentia Health

Evidentli

Greater National Advocates

Health Data Analytics Institute (HDAI)

Healthcare Performance Group, Joe Derenzo

Healthcare Performance Group, Joe Vanderbeek

HealthLeap

HonorHealth, Craig Norquist, MD

Hyro

Innovaccer

Intermountain Health

Iowa State University

John Snow Labs

Lucem Health

Marshall Health Network

Mayo Clinic, Shiba Kuanar, PhD

MCG Health

Medical University of South Carolina

Mercy

MGB

Mount Sinai Health System

Munson Healthcare

National Association of Community Health Centers

National Institutes of Health (NIH), National Center for Complementary and Integrative Health (NCCIH), Qilu Yu, PhD

New York-Presbyterian

NJ Department of Health

Nordic Global

NYU Langone

Optum

Oracle Health, Madhvi Sharma

Orion Health

PointClickCare

Practice Fusion

Product Advisory Collective

Providence, Tristan Markwell

Quest Diagnostics

RAAPID

RAND Corporation

Regard

Roche

Rush University Medical Center

Shriners Children's

Steer Health

Suki, Sudha Jayaraman, MD, MSc, FACS

Sutter Health

Texas Children's Hospital

ThetaRho

UC San Diego

UC San Francisco

UCI Health

University of Alabama at Birmingham

University of Missouri

University of Pittsburgh

University of Rochester Medical Center

UT Southwestern

Vanderbilt Department of Biomedical Informatics

Viz.ai

WellSky

Xsolis, Heather Bassett

ZGM Advisors, Tanya A. Zucconi, MBA