Use Case Best Practice Guide

Clinical Decision Support

TABLE OF CONTENTS

Clinical Decision Support

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 clinical decision support (CDS) solution. The guide is organized by role (developer/implementer) and responsible AI principle areas where applicable (see figure below).

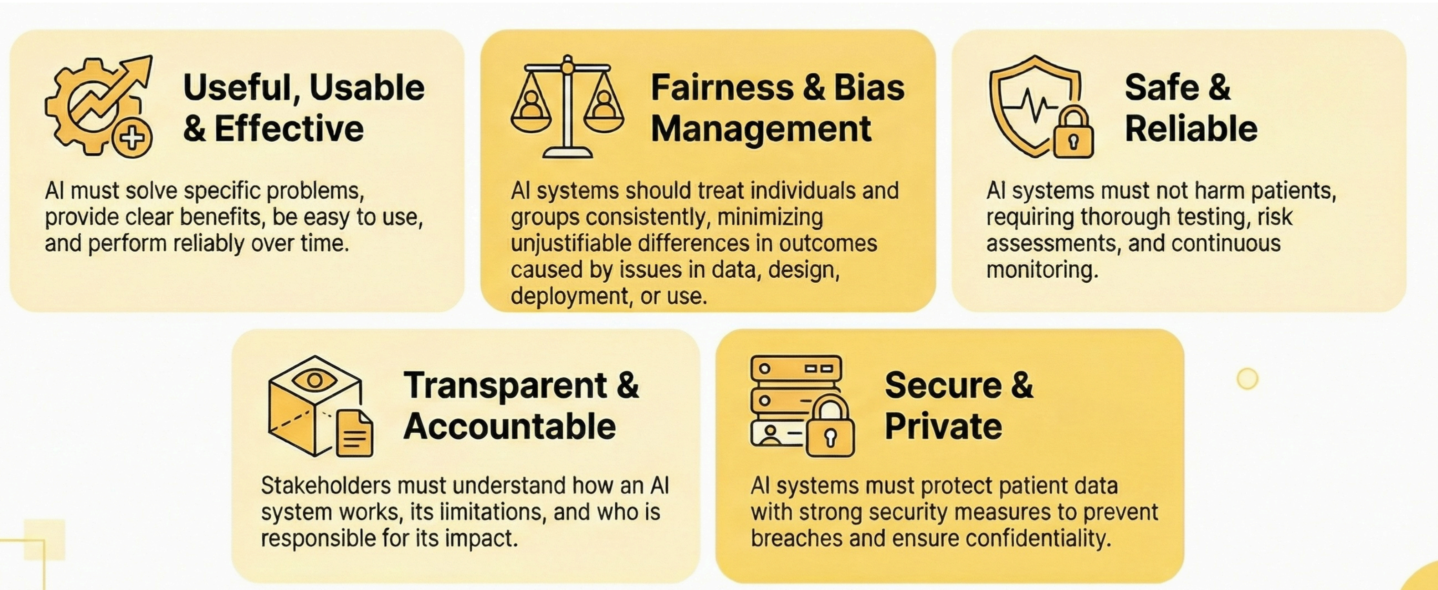

Use Case Description

While there are many different clinical decision support (CDS) use cases, this chosen focus area describes a CDS solution that leverages a large language model (LLM) using a retrieval-augmented generation (RAG) approach to process and deliver evidence-based medical information. By integrating with curated medical content, the AI system provides rapid and personalized clinical insights at the point of care. When a healthcare professional queries a clinical topic, the system generates an AI-driven response displayed alongside conventional search results. These responses reduce the need for clinicians to navigate through multiple sources. The system ensures transparency by displaying reference information alongside AI-generated responses, allowing users to assess and verify sources effectively. Ultimately, clinicians retain full responsibility for evaluating and integrating AI-provided insights into their decision-making process. While the content presented here focuses on this type of CDS use case but also discusses some aspects relevant to other CDS applications.

Primary End Users

- Healthcare professionals (anyone from students to experienced professionals), clinicians, and medical researchers seeking rapid access to clinical information.

Secondary End Users

- Clinical informaticians, hospital administrators, and medical librarians facilitating AI integration in healthcare settings.

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)

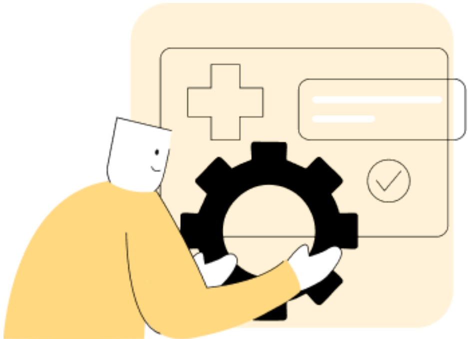

Implementers: individual(s) responsible for the procurement, deployment, and/or overall realization of a system or component in accordance with a specified design (derived from IEEE 829 and IEEE 730)

Listening In: A Summary of Challenges & Insights

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

Challenge #1: Fragmented and incomplete patient data reduces accuracy and meaningfulness.

- External data (e.g., cross-hospital patient histories) may be incomplete, leading to skewed predictions. AI-enabled CDS solutions often rely on fragmented patient data confined to a single hospital or health system. When a patient receives care across multiple facilities, the lack of comprehensive data (e.g., external EHR records, lab results, or pharmacy claims) can reduce the accuracy of predictions or recommendations.
- Patient context—including unstructured data like notes—is critical for meaningful decision support. Reliance on structured data alone risks missing key social or clinical signals (e.g., housing instability affecting diabetes care)
- Social determinants of health (SDOH) significantly influence patient outcomes and must be considered in model development.
- Model predictions trained on biased or incomplete data may overlook social risk or underrepresent safety-net populations.

Challenge #2: Healthcare professionals lack time to understand complex AI explainability techniques during clinical workflows.

- Clinicians tend to respond better to systems that are explainable and interactive. They prefer traceability from recommendations back to patient data or evidence sources, and prefer natural language, intuitive interfaces over rigid alerting systems.
- Confidence scores are often difficult to interpret without context.
- Clinicians need clarity that AI outputs are guidance cues, not final answers.
- Adoption of clinical AI tools often reveals gaps in transparency around data lineage, training sets, and vendor practices.

Challenge #3: LLM hallucinations and unbounded inputs can produce fabricated or unsafe clinical details without proper constraints.

- “Prompt and pray” approaches—dropping full records into general LLMs—can yield erratic, unsafe outputs. Safe AI-enabled CDS requires structured grounding (e.g., RAG + validated guideline libraries) and plausibility checks to reduce spurious reasoning
- Clinical notes contain variability and noise (e.g., incomplete details). AI-enabled CDS solutions cannot guarantee truth of underlying documentation. These solutions should avoid unsupported predictions and focus on transparent presentation.
- Generative AI introduces risks of hallucination and bias. Clinicians emphasize that outputs must be verifiable and grounded in real evidence, especially when patient context intersects with inequities.
- AI-enabled CDS solutions may generate outputs even in inappropriate contexts, such as outside the intended patient population or with insufficient input data.
- LLM hallucinations can occur if patient context is too broad or insufficiently constrained. LLMs may “hallucinate” by producing fabricated or irrelevant clinical details if their input context is not carefully bounded.

> **Note**
> It is important to measure the efficacy of retrieved information, and in particular to have metrics that can capture evidence concordance and hallucination rates. See accompanying CDS testing and evaluation framework for a list of relevant metrics.

Challenge #4: Healthcare professionals may misinterpret or misuse AI outputs when systems lack clear intended use boundaries and validation pathways.

- AI outputs that lack appropriate human validation steps (e.g., clinician review) can pose safety risks if blindly trusted, especially when recommendations appear authoritative or urgent.
- Human review remains essential. Even high-quality automated matches need physician oversight, both to validate correctness and to catch automation bias where clinicians may over-trust the AI's outputs.
- Clinicians and patients may be unaware of how AI-generated risk scores—particularly those not tied to a diagnosis—are recorded or used in downstream clinical decision-making processes.
- Clinicians often misunderstand or misapply AI tools due to limited clarity around intended use, target population, or model design choices.

Insight #1: Healthcare Professionals typically prefer AI-enabled CDS solutions that make documentation usable by surfacing longitudinal trends and contextual signals.

- Clinicians need AI solutions that make fragmented documentation usable — surfacing longitudinal trends from multiple notes in a way that saves time and prompts better questioning.
- Patient context—including unstructured data like notes—is critical for meaningful decision support. Reliance on structured data alone risks missing key social or clinical signals (e.g., housing instability affecting diabetes care)
- Looking at data across large patient groups can help identify differences in health needs (for example, more financial strain or higher rates of depression in certain communities). But these analyses must be handled carefully. If test data only reflects certain kinds of query types or certain demographics, the results may be biased.

Insight #2: Conservative safety thresholds and recurring validation (e.g., silent evaluation, calibration checks) are important when LLM outputs influence clinical decisions.

- Clinical AI systems should adhere to conservative safety thresholds, reflecting the low tolerance for error in healthcare settings. Even rare misclassifications can carry significant consequences, so acceptable error rates are often stricter than in other domains.
- Prospective silent validation and recurring calibration checks are important to maintaining safety.
- As models evolve rapidly, organizations will benefit from strong version controls to balance innovation, stability, and cost in production environments.
- External reviewers reduce bias, but disagreements are common; scalable approaches must combine human subject matter expert (SME) judgment with automation.

Insight #3: Meaningful AI-enabled CDS impact comes from actionable, evidence-grounded interventions, not static predictions.

- AI-enabled CDS solutions should focus on actionable interventions rather than static predictions (e.g., predicting readmissions without suggesting modifiable steps has limited value).

- Clinical guidelines often use vague terms like “moderate risk” or “low oxygen” that different clinicians interpret in different ways. AI can end up copying these inconsistencies, leading to uneven or unsafe recommendations.

Best Practices for Developers

Usefulness, Usability, & Efficacy

Ingest, Interpret, and Surface Contextual Clinical Data

- Build AI solutions that have the ability to aggregate across provider notes and visualize data in a format clinicians can review and easily understand.
- Design AI-enabled CDS systems to ingest and interpret both structured (e.g., FHIR) and unstructured data sources. Build pipelines that surface context-sensitive recommendations and flag when contextual nuances (e.g., social determinants) alter applicability of guidelines. Of note, unstructured data will not be required across all pipelines - determine where unstructured data could be useful and consider how to use.

Deliver Predictive and Descriptive Outputs That Are Clinically Actionable

- Design AI-enabled CDS systems that connect predictive outputs to clinically meaningful or operationally actionable insights—whether at the individual or population level. Where appropriate, surface modifiable or evidence-linked factors (e.g., medication adherence, follow-up visits) to support shared decision-making and quality improvement, while maintaining transparency about model rationale and avoiding automated intervention recommendations.

Build Layered, Clinically Relevant Explainability

- Design explainability features using a layered approach—providing clear, clinically relevant outputs by default, with optional access to detailed model information for users who want deeper transparency. For example, rather than showing “elevated red cell distribution contributed to the sepsis prediction,” display: “This patient has a 58% chance of developing sepsis in the next 24 hours," with a "Learn More" option linking to contributing factors or supporting documentation for more tech-savvy users.

Standardize Output Formats for Reliable Integration

- Deliver AI outputs in predictable formats (JSON/XML) so CDS solutions can parse and connect data reliably. For example, return clinical recommendations as JSON objects with coded fields (e.g., "suggested_action," "evidence_source") for smooth integration with an EHR dashboard.

Pilot Safely Before Deployment

- Pilot new features in development sandboxes before clinician/patient-facing release ("fail fast, fail cheap"). For example, test an AI-enabled CDS alert for abnormal labs in a training environment with synthetic records before deploying to live patients.

Benchmark Against Current Clinical Practice

- Define performance benchmarks relative to current clinical practice. Use these benchmarks to guide iterative refinement.

Fairness & Bias Management

Incorporate SDOH Thoughtfully to Improve Fairness Without Amplifying Bias

- Incorporate SDOH features where available, but ensure these do not amplify bias against vulnerable populations (e.g., Medicaid patients, low-income communities); Including SDOH improves the real-world applicability and fairness of CDS models by identifying factors that may affect patient care planning (e.g., access to transportation for appointments, medication affordability).

Build Test Datasets With Explicit Demographic and Clinical Diversity

- When building test datasets, include a balanced mix of clinical specialties, demographics, and query types (e.g., diagnosis questions, treatment questions). Clearly label (e.g., metadata tagging) these categories so fairness checks can be completed.

Safety & Reliability

Suppress Unsafe Outputs by Embedding Guardrails and Population-Specific Exclusion Logic

- Design CDS models with embedded guardrails, and/or incorporate exclusion criteria directly into the model logic, that suppress or silence recommendations in cases where the model is known to be unreliable, unvalidated, or operating outside of its intended population; for example, an AI-enabled CDS tool that predicts risk for opioid misuse is trained only on patients without a history of opioid use disorder (OUD) or overdose. Instead of relying on a clinician to know this limitation, the model should automatically suppress its output when a patient's chart includes an OUD diagnosis or recent overdose event.

Minimize Hallucinations

- Design systems that minimize hallucinations by constraining model inputs to relevant, validated patient data when feasible, and by transparently communicating data sources, limitations, and uncertainty to users. Techniques such as using scoped vector databases or embedding queries can help narrow model context and reduce irrelevant or fabricated outputs. Systems should also surface when insufficient information prevents reliable responses and clearly distinguish verified facts from model-generated content.

Ground AI-enabled CDS Recommendations in Trusted Evidence

- Use retrieval-augmented generation (RAG) or model context protocol (MCP) approaches that ground recommendations in trusted medical guidelines and structured evidence. Plausibility filters can be included as a step before the AI does deep analysis, where obvious mismatches or impossible options are removed (e.g., ruling out pediatric criteria for an adult patient). This reduces wasted computation and helps prevent nonsensical outputs. Employ plausibility filters to screen out irrelevant criteria before evaluation.

Ensure accurate and high-quality data capture

- Ensure accurate and high-quality data capture from diverse clinical documents by selecting appropriate processing approaches—such as traditional ML services, generative AI models, or a combination of both—based on data type and use case. When confidence is low, route extracted information for human review to safeguard data quality and clinical reliability. For example, a developer may use a document-intelligence API to parse scanned notes and a generative model to cross-check medication fields, with uncertain entries flagged for clinician validation.

Use Interoperability Standards to Expand Reliable Data Coverage

- Use data exchange protocols or interoperability standards to integrate external patient data where feasible; Data exchange protocols may include ONC's Trusted Exchange Framework and Common Agreement (TEFCA) - designed to facilitate nationwide interoperability among health information networks (HINs) for broader data exchange; Interoperability standards may include HL7 Fast Healthcare Interoperability Resources (FHIR) - adopted for modern APIs, enabling standardized exchange of clinical data.

Prioritize Data Quality by Avoiding Premature Prediction

- Do not prematurely move to predictive analytics on top of noisy documentation. Instead, focus on string-level extraction and transparent surfacing of source text and references (e.g., clinical documentation excerpts, guideline citations) so clinicians can validate outputs.

Define Safety-Aligned Thresholds

- Set operating thresholds based on a formal risk assessment that reflects the intended use, patient population, and clinical context—rather than developer preference. Collaborate with clinicians and risk managers to define what "safe" performance means for each application, prioritizing patient safety even if it reduces model responsiveness or recall in lower-risk, non-clinical workflows. For example, an AI solution that flags potential strokes should be tuned to avoid missing even a single true case, while a scheduling optimization tool may tolerate more false alerts since it doesn't directly impact patient safety.

Validate Prospectively Using Silent Mode and Pilot Testing Before Live Deployment

- Validate prospectively before live deployment. Run models in "silent mode" for extended periods to confirm stability. Silent mode means the AI runs in the background, making predictions without influencing real clinical decisions. This allows teams to check accuracy and stability safely before putting it into actual use.

Adopt Controlled Versioning

- Adopt develop→test→production promotion for models; lock production solutions to approved version IDs. For example, maintain clear version tags for each CDS solution, test new versions in a staging environment with representative cases, and promote only when internal safety criteria are met — keeping production pinned to the approved version until validation is complete.

Implement Guideline Refresh Cycles

- Implement refresh cycles or automated triggers so vector databases stay current with clinical guideline updates. For example, re-index stored guidelines whenever new evidence is published or institutional protocols change.

Conduct Ongoing Calibration and Post-Deployment Monitoring

- Conduct prospective validation and recurring calibration checks before and after deployment to confirm model stability, accuracy, and fairness under real-world conditions. Use approaches such as running models in "silent mode" or within limited pilot settings to observe predictions without influencing clinical decisions, capturing evidence of model drift or data distribution shifts. The duration, metrics, and evaluation cadence should be defined collaboratively with customers and clinical stakeholders based on risk and intended use. For example, an AI-supported sepsis prediction tool might run silently for 3 months to compare predicted alerts against confirmed cases, enabling recalibration if performance degrades across different patient populations.

Transparency

Communicate Model(s) Purpose, Scope, and Assumptions in Plain Language

- Ensure model labeling clearly communicates inclusion and exclusion criteria (e.g., cohort design, target outcome) in plain language for clinical end-users. Avoid technical jargon and explain model intent at the point of use.
- Build "context libraries" that spell out clear definitions and edge-case rules for ambiguous terms. Flag these areas as high-risk and require human review to prevent the AI from making unsafe assumptions.

Make Evidence, Data Inputs, and Guideline Application Fully Verifiable

- Incorporate transparent sourcing (e.g., links to guidelines, citations, or underlying patient record sections) in CDS outputs. Build user interfaces that support natural language queries and interaction rather than relying solely on alerts or complex menus.
- Implement mechanisms for verifiability. Recommendations should always include the source of supporting evidence and the specific patient data points used. Where guideline application could amplify inequity, embed discrepancy detection or context filters (e.g., avoid diet advice for patients flagged as homeless).

Provide Lifecycle Transparency Through Model Documentation

- Include details within model cards (or similar) such as data retention, de-identification, and retraining practices so downstream users can answer regulatory questions.

Communicate Uncertainty Clearly and Tie It to User Action

- Provide interpretable confidence indicators that explain the source of uncertainty (e.g., low confidence due to unfamiliar input patterns) and tie these to user actions (e.g., consider secondary review).

Best Practices for Implementers

Usefulness, Usability, & Efficacy

Select and Evaluate AI Solutions Based on Organizational Value and Comparative Performance

- Develop clear criteria for selecting AI solutions that provide the most benefit and align with organizational priorities, while avoiding fragmented or misaligned adoption. As AI-supported CDS solutions evolve, expect to evaluate and integrate multiple specialized modules rather than a single all-in-one platform.
- When evaluating AI solutions, compare performance to existing workflows and outcomes; establish processes for continuous monitoring and improvement.

Deliver Clear, Contextualized On-Screen Guidance to Support Correct Interpretation

- Provide brief, contextualized tooltips or labels within the EHR interface (e.g., "Model trained for patients ⩾18 years old with no prior opioid overdose") to help clinicians correctly interpret model outputs without relying on back-end documentation.

Present Succinct, Actionable Outputs That Minimize Cognitive Burden

- Display EHR interfaces with succinct, actionable outputs that fit seamlessly into the clinical workflow. Avoid requiring users to navigate to secondary pages or interpret probability distributions; consider incorporating “summary boxes” or alert banners that present a risk category (e.g., “High Risk”) alongside a concise, clinically relevant recommendation.

Embed Alerts Where Users Can Act and Minimize Workflow Disruption

- Integrate AI-enabled CDS alerts into workflows where decision-makers (clinicians, case managers) can act on the suggested interventions; e.g., avoid sending alerts if the risk factors are non-modifiable.
- Prioritize workflow embedding (e.g., single sign-on, minimal clicks) so clinicians don’t need to step outside of their existing processes to access AI support.

Fairness & Bias Management

Ensure Training Data and Model Performance Reflect the Populations Served

- Verify that an AI solution's training data and performance characteristics reflect the demographics and clinical profiles of the populations the organization serves. Models trained on one population (e.g., adults) should not be applied to substantially different groups (e.g., pediatrics) without additional validation, if at all. Pay particular attention to whether social risk factors or underrepresented populations are adequately represented.

Encourage Users to Treat AI Outputs as Prompts for Dialogue, Not Clinical Conclusions

- Support users in treating AI outputs as discussion prompts with patients, colleagues, and other stakeholders - rather than confirmed findings. Reinforce this through training and deployment messaging.

Monitor Override and Usage Patterns to Detect Systematic Bias Across Populations

- Analyze override and usage data by population to surface systematic bias. For example, review whether clinicians override AI-generated screening prompts more often for certain patient groups.

Safety & Reliability

Conduct Formal Review and Documentation of Safety Constraints Before Deployment

- Establish a formal process requiring that all AI modules—including updates or changes from existing vendors—are reviewed prior to deployment, regardless of vendor reputation/relationship.

- Request vendors document exclusion criteria and validate that the model is silent in contexts where it lacks reliable inference capacity; include this criterion in procurement checklists or AI governance review.

Ensure Human-in-the-Loop Clinical Oversight and Guard Against Automation Bias

- Deploy AI-enabled CDS solutions with human-in-the-loop review protocols. Ensure that users can easily approve, edit, or reject AI-generated suggestions, and monitor usage data for automation bias (e.g., over-reliance without adequate review).
- Keep humans in charge of clinical approvals/denials; use AI only to pre-populate or suggest, not to finalize. For example, allow AI to draft suggested differential diagnoses, but require a clinician to confirm before adding them to the patient record.

Combine Human Oversight and Automated Verification During Pilots

- Combine human oversight and automated verification tools during the piloting and deployment stages of CDS solutions to track and mitigate hallucinations. Pilot stage monitoring: treat early piloting as a "learning period" where hallucinations are logged, categorized, and analyzed for root causes (e.g., over-broad queries or uncurated free-text notes); Secondary AI tools ("AI Judges"): use a second LLM or rules engine to cross-check primary AI outputs against verified EHR data.

Establish Continuous Monitoring and Scalable Deployment Pathways

- Provide intake/feedback channels so users can flag drift or stale data quickly. For example, maintain a shared chat space or ticket queue where staff report if guideline-based reminders seem out of date.
- Establish end-to-end deployment pathways, including workflow integration and post-deployment monitoring, before scaling AI-enabled CDS solutions.

Transparency

Communicate Limitations Due to Incomplete Data

- Clearly communicate the limitations of predictions when data is incomplete (e.g., no history from external facilities). AI-enabled CDS solutions often operate on partial datasets due to limitations in interoperability, access to external EHR systems, or the absence of historical data from other care providers. Implementers (e.g., health IT teams, system integrators) need to set clear expectations with clinicians and end-users about these data gaps.

Establish Clear Responsibility Guidelines and Labels

- Establish clear clinical responsibility guidelines and validation expectations for end users, such as ensuring clinicians understand which AI outputs are preliminary or unvalidated and which are fully reviewed; e.g., labeling: "Not yet reviewed by clinician" or "AI-generated preliminary alert—requires clinical validation"); or, training end users during pilot on interpretative responsibility (e.g., how and when to trust, question, or override AI recommendations based on clinical judgment).

Ensure Transparency Around Data Use, Risk Scores, and Vendor Practices

- Establish data use protections to prevent the misapplication of AI-generated risk scores, especially when those scores are not clinically confirmed diagnoses; implementers should ensure that risk scores are not treated as structured diagnostic data in the EHR unless clinically validated; implementers should also include prohibitions in data sharing agreements to prevent risk scores from being shared with payers, insurers, or third parties unless explicitly authorized and justified for care coordination.
- Request transparency from vendors via model card (or similar). Before adoption, clarify how vendor systems were trained, whether patient data are retained, and how secondary uses (e.g., retraining) are handled.

Info

Many areas of healthcare rely on visible vigilance processes that allow clinicians or end-users to report suspected errors or adverse events, and for those reports to accumulate in a way that supports shared learning and timely mitigation as facilitated by transparency. Examples of this might include Adverse Drug Event or Reaction reporting processes for medications. The group discussed the value of real-world reporting and feedback processes for stakeholders once CDS tools are deployed, with an emphasis on transparency for such reports. Such transparency is likely to increase in importance when the tool spreads across multiple sites.

For non-FDA regulated CDS solutions, it may be helpful to discuss and identify what post-deployment issue/feedback handling processes are between implementing and developing organizations (e.g. between procurers and vendors, or between development teams and implementing/user teams). It may also be helpful to identify ways of transparently sharing learnings with other sites currently deploying the same solution.

Note that for FDA regulated CDS solutions, there are required post-market reporting guidelines and transparency processes (See FDA SAMD and AI related guidance and regulations, including [Medical Device Reporting (MDRs)] and the [Medical Device and User Experience (MAUDE) database] for further information).

Best Practices for Developers & Implementers

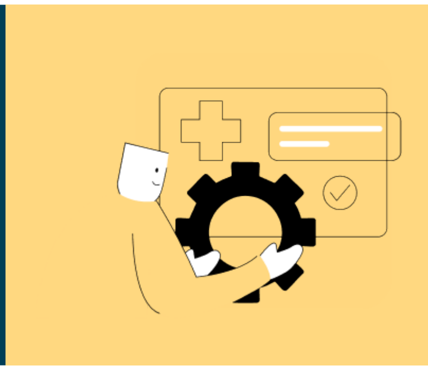

Security & Privacy

- Implement multi-layered controls – Inference, Use Case, and IT Controls;
  - (1) Inference Controls: e.g., monitor AI inputs and outputs for anomalies, test for adversarial prompts or data poisoning, and ensure outputs are verifiable against structured data
  - (2) Use Case Controls: e.g., continuously validate whether the AI-enabled CDS solution is meeting hospital goals (e.g., improving workflow efficiency, improving sepsis prediction in patient population) and align model outputs with real clinical needs and
  - (3) IT Controls: e.g., ensure robust cybersecurity, role-based access management, encryption, and prevention of unintended data sharing or leakage.

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 ⩾ 2/3 vote "Include" → **Consensus achieved** (no further action).
  - If < 2/3 "Include", but ⩾ 2/3 combined "Include" + "Include Contextually" → **Flagged for Phase 2**.
  - If ⩾ 25% vote "Exclude" → **Flagged for Phase 3**.
  - If ⩾ 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 ⩾ 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 start by thanking 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 Clinical Decision Support Work Group. Below is a list of organizations who had at least one individual who showed interest and/or participated in our Clinical Decision Support 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
Han-Chin Shing	Amazon
Geralyn Miller	Microsoft
Raj Ratwani	Medstar Health
Jessica Handley	Medstar Health
Katherine Eisenberg	EBSCO/DynaAI
Ben Hollis	EBSCO/DynaAI
Pawan Jindal	Darena Solutions
Howard Strasberg	Wolters Kluwer
Sheila Bond	Wolters Kluwer
Dennis Shung	Mayo Clinic
Amber Featherstone-Uwague	Elsevier

Participating Organizations

- 101 GenAI
- AdventHealth
- Agency for Healthcare Research and Quality (AHRQ)
- AI VOBIS
- Airia
- Alameda Health System
- ALIGNMT AI
- Ambience Healthcare
- American Academy of Family Physicians (AAFP)
- Amruta Inc
- Ascension
- Asher Orion Group
- athenahealth, Marisa Sileo, MSN, RN, NI-BC
- Atria Health
- Atropos Health
- AWS
- BCBS Minnesota
- Belmont University
- Boston Children's
- Brooke & Associates
- Clinia, Daniel Buades Marcos, MS
- Cofactor Labs, Arfa Rehman
- Columbia University Irving Medical Center
- Cooperative Benefits Group, Mark L. Fowler
- Darena Solutions/Prompt Opinion
- Dartmouth Health
- Department of Health
- Duke
- DxTx
- EBSCO/DynaAI, Katherine Eisenberg, MD
- Elsevier
- Endeavor Health
- Ensemble Health Partners
- Essentia Health
- Federation of State Medical Boards (FSMB)
- Ash Hopkins, Flinders University
- Grand Prairie Services, Leon M. Huddleston, MD, JD
- Greater National Advocates
- Hartford HealthCare
- HCA Healthcare
- Health Data Analytics Institute (HDAI)
- HealthPartners
- iBPM
- Key Software

- Mass General Brigham
- Mayo Clinic, Joe Melnick
- Mayo Clinic, Dennis L. Shung, MD, PhD
- MCG Health, Jason Gillman, MD, FIDSA
- Medstar
- Memorial Sloan Kettering Cancer Center
- Mercy
- MettaHealth Partners
- Microsoft
- Mount Sinai Health System
- Kleber Araujo, MV Saúde Digital
- Jeremy Attermann, National Council for Mental Wellbeing
- National Association of Community Health Centers (NACHC)
- National Institutes of Health (NIH), National Center for Complementary and Integrative Health (NCCIH), Qilu Yu, PhD
- New York-Presbyterian
- Nordic Global
- North Carolina Community Health Center Association
- Northern University
- Northwell Health
- Optum
- Oracle
- Orion Health
- Parkland Health
- PointClickCare
- Product Advisory Collective
- Providence
- Quest Diagnostics
- RAAPID
- RAND Corporation
- Roche
- Rush University Medical Center
- San Joaquin General Hospital
- Sayvant
- Scale LLP
- Sharp HealthCare
- South Carolina Primary Health Care Association
- Suki, Sudha Jayaraman
- Sutter Health
- Trimesters AI
- UC Irvine Health
- Unified Women's Healthcare
- UnitedHealth Group/Optum
- University of Alabama at Birmingham
- University of Missouri
- University of Pittsburgh

- University of Vermont
- Vivyo Group
- Viz.ai
- WellSky
- Wolters Kluwer