Responsible Use of AI in Healthcare and Public Health: From Innovation to Implementation

Merage Ghane, PhD

CHAI

AI THAT SERVES ALL OF US.

Learning Objectives:

- Value add and interest in AI
- Define Responsible AI in the context of healthcare
- Identify potential legal, ethical, and safety risks
- Learn about available frameworks, tools, and resources
- Discuss governance strategies
- See examples of responsible AI principles applied to real-world use cases.

CHAI

The Opportunity:

In population health and value-based care, AI can:

- Reduce clinician, administrative, and operational burden, protecting workforce
- Help identify rising-risk patients
- Help optimize resource allocation
- Support timely & personalized patient engagement to support care coordination and preventative interventions
- Help with quality reporting
- Help with surveillance and outbreak detection
- Help with identifying environmental and non-medical determinants of health

CHAI

The Challenges:

- AI’s potential to scale existing baises or population differences, or introduce new ones
- Change management challenges with new technologies
- Resource differences across organizations
- Desire to innovate and stay current, but also minimize harm.
- Not knowing where to start.

CHAI

WHY GOVERNANCE MATTERS

- Aligns AI use with mission, ethics, and compliance
- Minimizes harm, inconsistencies, and system-level risks
- Builds trust, clarity, and accountability across teams
- Ensures that money and time spent is intentional and not wasteful

THE AI LIFECYCLE

The AI lifecycle is central to understanding and implementing CHAI's Responsible AI Guidance in healthcare. The six-step lifecycle outlines the essential stages and processes involved in developing, deploying, and maintaining AI systems.

By systematically addressing each phase of the lifecycle, the framework ensures that AI systems adhere to the highest standards of safety, efficacy, fairness, transparency, and security. This structured approach supports risk mitigation, managing model inconsistencies, and promotes accountability and trustworthiness in AI applications.

1

- Engage stakeholders to define the problem & perform root-cause analysis
- Identify solution & plan future state
- Gather business requirements
- Assess feasibility, potential for impact, & prioritization
- Make procure/build/partner decision

2

- Select/understand model task & architecture
- Capture design & technical requirements or determine best solution to meet business requirements
- Design solution application & system workflow according to human-centered design principles
- Design deployment strategy with end users
- Design risk management, monitoring & reporting plan

3

- Access data
- Prepare data
- Develop data management plan
- Train & tune model

4

- Conduct installation qualification (when applicable)
- Validate local system performance (when applicable)
- Execute prospective, silent evaluation
- Establish risk management plan
- Train end users
- Test usefulness
- Ensure compliance with applicable healthcare regulations & standards

5

- Implement small-scale pilot to assess real-world impact
- Execute and update risk management plan
- Educate & train users on AI application reporting
- Assess usefulness and adoption

6

- Deploy at a larger scale on a general population
- Audit AI system to inform whether to maintain, refine or sunset
- Conduct ongoing risk management

CHAI

RESPONSIBLE AI PRINCIPLES

At its core, governance is to help define, document, and execute processes to ensure that these principles are being upheld at every stage of the AI lifecycle.

- Usefulness, Usability, & Efficacy
- Fairness
- Transparency
- Security & Privacy
- Safety

CHAI

BACKGROUND

- CHAI is developing a series of governance playbooks that will help organizations of all sizes/resourcing approach their health AI governance processes (Release date: June 2026)
- We released joint guidance that describes high-level foundational principles for responsible use of AI in healthcare.
- Information presented here reviews key components necessary for effective health AI governance
- Note that while less resourced healthcare and public health settings face unique challenges toward implementing AI governance, even larger health systems find this challenging. Working together can help us alleviate some of these burdens.

THE CORE COMPONENTS OF EFFECTIVE AI GOVERNANCE

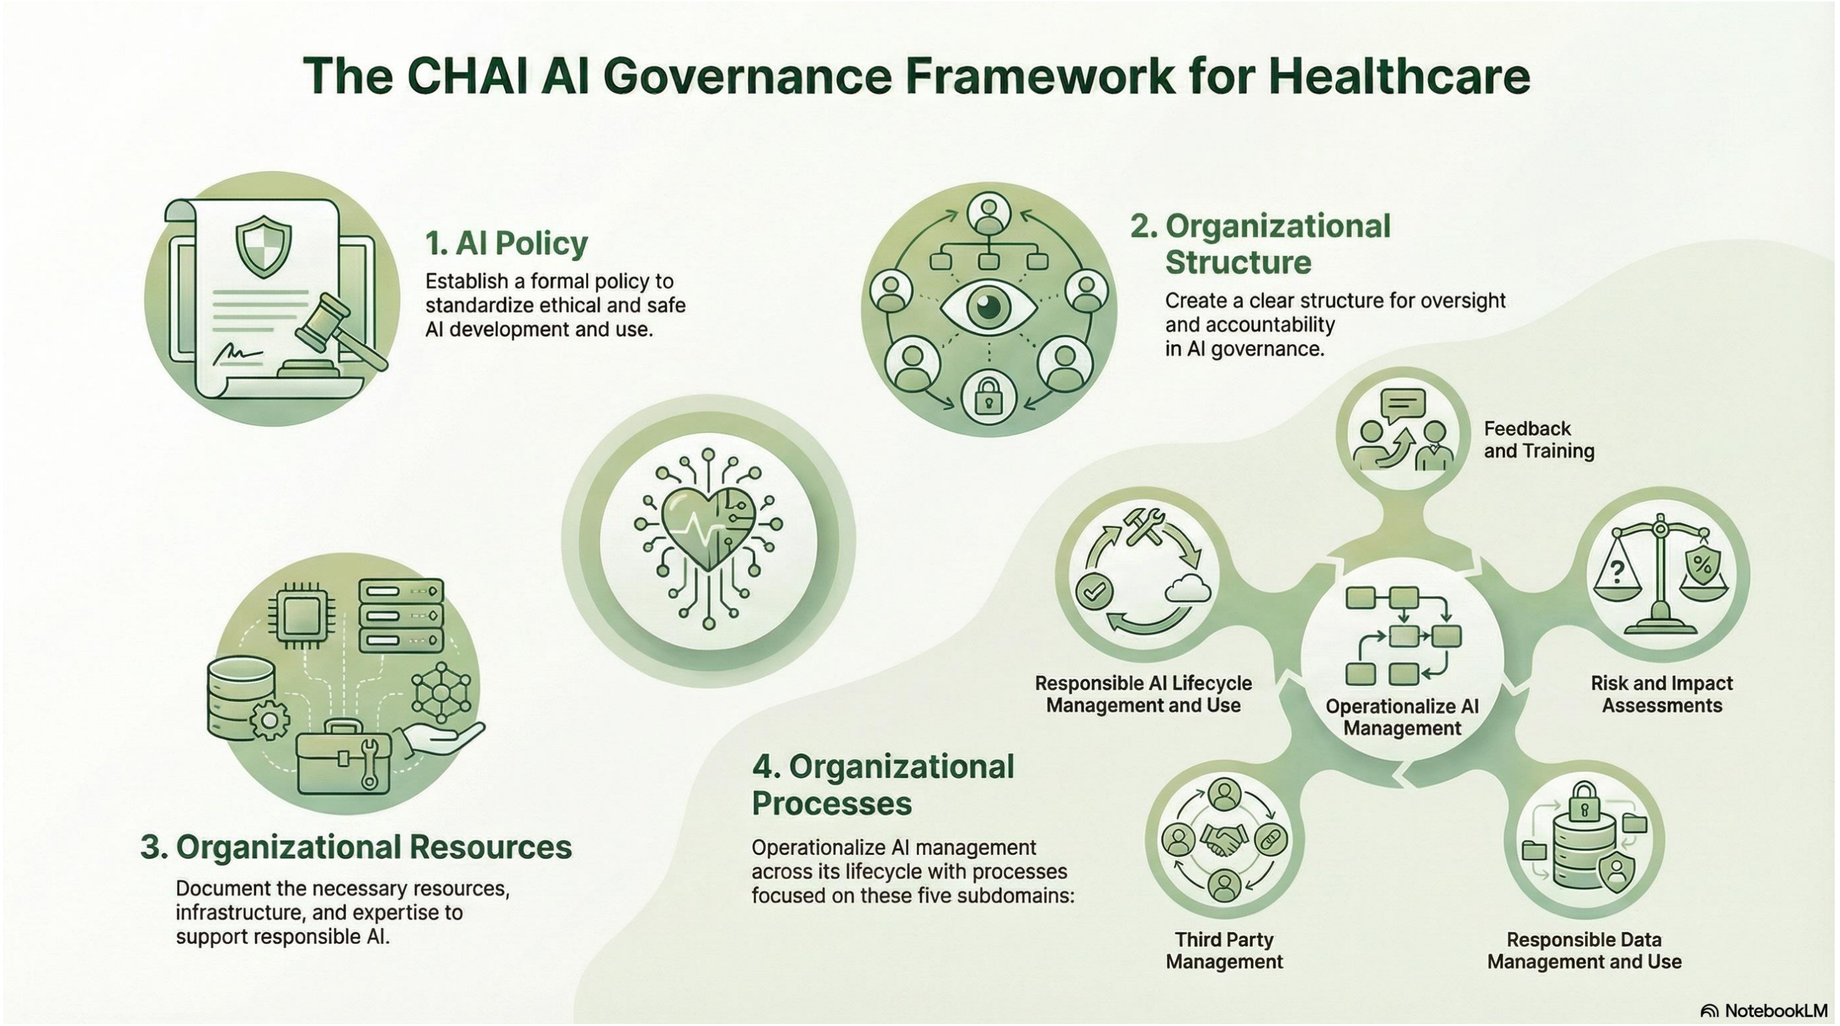

POLICIES

What:

- AI Policy: Formal documentation outlining scope, principles, and approval processes for AI use.
- Policy Alignment: Ensures consistency with broader legal, regulatory, and internal frameworks.
- Policy Review: Periodic review and revision cycle triggered by new risks or regulations.

How:

- Maintain version-controlled, leadership-approved policies.
- Cross-walk AI policy with privacy, security, and clinical safety policies.
- Schedule annual reviews and log change rationales.

POLICIES

Key Information to capture in an AI policy:

- Introduction & Purpose
- Scope & Applicability: Description of when/which tools require governance (internally developed, vendor supplied, platform embedded, administrative, clinical, etc.) and when they do not, as well as a description of permitted and prohibited uses (e.g. prohibition of use of PHI in public AI tools)
- Relevant Definitions & Terminology
- Governance structure: Who is involved in the governance committee, how is it structured, how does reporting occur, and how often do they meet?
- AI System Lifecycle Management: Description of the risk-based governance procedure & processes across the lifecycle

ORGANIZATIONAL STRUCTURE

What:

- Context: Defines governance scope, goals, and readiness.
- Roles & Responsibilities: Assign duties across key domains.
- Concern Reporting: Anonymous AI risk reporting protocols.

How:

- Create a governance charter and appoint cross-functional leads.
- Hold documented governance meetings regularly.
- Establish secure reporting channels and escalation paths.

CHAI

ORGANIZATIONAL STRUCTURE

Centralized: A centralized AI governance structure concentrates decision-making authority, standards, and oversight within a single enterprise-level body or function (e.g., a central AI governance office or committee).

Decentralized: A decentralized AI governance structure places primary responsibility for AI oversight within individual departments, service lines, or business units, with minimal central coordination.

Distributed: A distributed AI governance structure combines central standards and guardrails with local execution and accountability.

CHAI

ORGANIZATION RESOURCES

What & How:

- Resource Documentation: Inventory AI tools and systems.
- Data Resources: Document data assets and uses
- Tooling/Computing: What do you have and what do you need?
- Human Resources: Staff roles and competencies.

ORGANIZATIONAL PROCESSES

What:

Define objectives and operationalize processes for:

- Responsible AI system Lifecycle Management and Use
- Impact Assessment: Risk & Benefit
- Responsible Data Use and Management for AI Systems
- Third-Party Management
- Feedback and Training

PROCESS: RESPONSIBLE LIFECYCLE MANAGEMENT AND USE

What:

- Operationalize the objectives & processes for responsible use (e.g. how would you align with principles of responsible AI?)
- Lifecycle Stages: Define requirements for vendors and organization from development/purchasing to monitoring.
- Logging Events: Monitor for data and system drift or failure.

How:

- Maintain documentation per stage (e.g., validation reports).
- Log key model performance and use events.
- Review logs or set alerts for anomalies.

CHAI

PROCESS: RISK & IMPACT ASSESSMENT

What:

- Risk Categorization: Scores AI system's potential to cause harm.
- Risk Assessment: Evaluate high risk operational, clinical, and ethical risks.
- Create a Risk/Impact Profile: Track and mitigate risks and benefits

(e.g. IEEE 7003:2024).

How:

- Assess risk before and after deployment.
- Maintain Impact profiles and update throughout AI lifecycle.
- Document assessment results and use in approvals.

CHAI

Title: Risk Categorization Tool v2.0
Risk Assessment Domain: Life and Patient Safety & Data & Technology Risk
Primary Audience: Health systems (any size) and teams responsible for pre-deployment risk review.

Risk to Life & Safety

Risk Modifiers: (Low, Med, High)

- Proximity of Impact on Patients
- Human in the Loop
- Consequences of Failure
- Patient Population Vulnerability
- Level of Difficulty Monitoring Solution Output
- Data Transparency
- Solution's Clinical Level of Care
- Time to Intervention (if output is wrong)
- Breadth of Potential Harm
- Integrated Error Propagation Risk
- Population Sensitivity or Outcome Variability Risk

CHAI

Use Case Example: AI-assisted Patient Scheduling software (e.g., scheduling chatbot) - for outpatient, primary care clinics - simplifies booking, rescheduling, and managing patient appointments. These products allow patients to select appointment times and providers to manage their calendars, reduce no-shows, and optimize clinic workflows. Typical features include automated reminders, real-time availability updates, and integration with EHRs. Human confirms appointment once scheduled

Modifier: Distance from Patient
How physically or operationally close is the AI solution to the patient?

Low: No direct impact on individual patient care, supports back-end functions such as back office, administrative tasks, population health analysis, or workflow optimization

Medium: Indirect impact on patient care, access to care, or informational use, such as scheduling, transportation, non-clinical informational chatbots

High: AI solution has a semi-direct impact on patient care, such as, used by a healthcare professional as part of a broader clinical judgment; OR AI solution is directly involved in patient care/patient interaction

CHAI

Modifier: Human in the Loop
The extent to which human oversight is involved in reviewing, verifying, or overriding the AI solution outputs before they affect patient care

Low: AI solution output is always reviewed by a relevant expert before any action is taken

Medium: AI solution output has optional human in loop review by relevant expert before any action is taken

High: AI solution output is never reviewed by provider before an action is taken

CHAI

Modifier: Data Transparency

The clarity, completeness, and accessibility of the data sources and datasets used to train, test, and validate the AI solution.

Low: health system has complete access to training data of the underlying model(s) for the AI solution; lowest level of detail for the data/datasets are shared and available (e.g. AI solution is developed internally)

Medium: partial access to training data of the underlying model(s) for the AI solution; OR some level of detail for the data/datasets are shared and available.

High: no access to training data OR no components of the data/datasets are shared or available (e.g.. Data provenance and data catalog/dictionary unavailable

??

CHAI

Risk modifiers for Data and Technology Risk

- Use of Sensitive Data
- Accuracy, Completeness, and Veracity of Data Used for Model Training and Operation
- Sufficiency & Representativeness of Data Used for AI Model Training and Operation
- AI Model Security Vulnerabilities
- Third-Party AI components and Vendor Risk
- AI Model Lifecycle Management and Updates
- AI Monitoring, Incident Detection, & Response
- AI Detection & Traceability

PROCESS: THIRD-PARTY MANAGEMENT

What:

- Responsibility Allocation: Document **shared** duties and risks.
- Supplier Management: Enforce transparency and auditability.
- Customer Considerations: Ensure usability and trustworthiness.

How:

- Require model cards or equivalent and evaluation methods from vendors.
- Develop **shared** deployment plans and audit protocols.
- Define clear terms in contracts for incident response and feedback.

CHAI

Tool: Applied Model Card

Applied Model Card Template

Name: **Developer:**		**Inquires or to report an issue:** abc@abc.com or +1 (999) 999- 9999
Release Stage: **Global Availability:**	**Release Date:** **Regulatory Approval, If applicable:**	**Version:**
Summary: **Keywords:**		**Uses and Directions:** • **Intended use and workflow:** • **Primary intended users:** • **How to use:** • **Targeted patient population:** • **Cautioned out-of-scope settings and use cases:**

Warnings

- **Known risks and limitations:**
- **Known biases or ethical considerations:**
- **Clinical risk level:**

Trust Ingredients

AI System Facts:

- **Outcome(s) and output(s):**
- **Model type:**
- **Foundation models used in application, if applicable:**
- **Input data source:**
- **Output/Input data type:**
- **Development data characterization:**
- **Bias mitigation approaches:**
- **Ongoing Maintenance:**
- **Security and compliance environment practices or accreditations, if applicable:**
- **Transparency, Intelligibility, and Accountability mechanisms, if applicable:**

Transparency Information:

- **Funding source of the technical implementation:**
- **3rd Party Information, If Applicable:**
- **Stakeholders consulted during design of intervention (e.g. patients, providers):**

Key Metrics

Completed Example

Tool: Applied Model Card

Key Metrics					
Usefulness, Usability, and Efficacy		**Fairness and Equity**		**Safety and Reliability**	
Goal of metric(s):		**Goal of metric(s):**		**Goal of metric(s):**	
Result:	**Interpretation:**	**Result:**	**Interpretation:**	**Result:**	**Interpretation:**
Test Type:		**Test Type:**		**Test Type:**	
Testing Data Description:		**Testing Data Description:**		**Testing Data Description:**	
Validation Process and Justification:		**Validation Process and Justification:**		**Validation Process and Justification:**	

Resources

- **Evaluation References, If Available:**
- **Clinical Trial, If Available:**
- **Peer Reviewed Publication(s):**
- **Reimbursement status, if applicable:**
- **Patient consent or disclosure required or suggested:**

CHAI

PROCESS: DATA USE & MANAGEMENT

What:

- Data Use & Enhancement: Define training, validation, and deployment uses.
- Data Acquisition & Quality: Prioritize fairness, completeness, validity.
- Data Provenance & Preparation: Document transformations and sources.

How:

- Enforce data governance SOPs and source tracking.
- With third-party vendors: know where your data is going, what form it needs to be in, how it will be used, etc.
- Remember that useful AI needs quality data for training, tuning, validation, and use—data is an ASSET that can bring solutions
- Audit data for missingness and schema issues.

CHAI

PROCESS: DATA USE & MANAGEMENT

Steps:

- Identify data types
  - Protected Health Information (PHI) → BAA or HIPPA exception
  - Limited Data Set (LDS)- de-identified except for things like city, date, Zip-code → DUA or HIPPA exception
  - De-identified data (DID)- safe harbor or expert determination → DUA/ Other

PROCESS: DATA USE & MANAGEMENT

BAA: Creates, receives, maintains, or transmits PHI

- **Permitted Uses and Disclosures:** Must be limited to defined functions (e.g., data hosting, analytics).
- **Safeguards:** Administrative, physical, and technical measures consistent with the HIPAA Security Rule.
- **Breach Notification:** Timely reporting (typically within 10–15 days) of any impermissible use, disclosure, or security incident.
- **Subcontractors:** Flow-down of all BAA obligations to any subcontractors.
- **Return or Destruction of PHI:** Required upon contract termination.
- **Indemnification and Liability Caps:** Important for risk allocation (often negotiated heavily with vendors).

DUA: Limited Data

- **Permitted Uses and Disclosures:** Only for research, public health, or health care operations.
- **Recipient Obligations:**
  - No re-identification or contact of individuals.
  - Use only for purposes specified in the DUA.
  - Implement safeguards to prevent unauthorized use or disclosure.
- **Reporting Obligations:** Must notify the covered entity of any misuse.
- **Data Flow Control:** Subcontractors or third parties must be bound by equivalent DUA terms.

CHAI

PROCESS: DATA USE & MANAGEMENT

De-identified Data Considerations

a. Contractual Controls

Include use and re-identification prohibitions even when HIPAA doesn't require them.

Define ownership and derivative rights: Who owns new datasets, trained models, or insights created using the de-identified data?

Require non-re-identification warranties and monitoring/audit rights.

b. Re-identification Risk

Prioritize data minimization

c. State and Other Laws

Some states (e.g., California, Washington, Colorado) impose privacy obligations even for de-identified or pseudonymized data under their consumer privacy laws (CCPA/CPRA, MHMDA, CPA).

Ensure contracts reflect both federal and state privacy regimes.

Guidance for data management & use outside of HIPAA

- Encryption in transit and at rest when possible
- Access controls
- Regular security assessments
- Incident response plan
- Effective Data Use Agreements:
  - Permitted uses
  - Data minimization
  - Explicit Prohibition of reidentification
  - Third-party obligations that align with internal policies (e.g. encryption, regular security assessments, etc.)
  - Audit rights

Upcoming tool: Technology & Data Risk Tool

CHAI

PROCESS: FEEDBACK & TRAINING

What:

Define and document:

- Processes for user feedback and incident reporting
- Key AI-related incident types
- Reporting pathways and responsible parties
- Processes for how users are trained

How:

- Update existing clinical and operational safety reporting processes
- Ensure the right people know the right information about the AI system
- Train relevant staff on how to use and monitor the AI solution effectively
- Consider broader AI Literacy education and change management

CHAI

GETTING STARTED

- Draft or refine an AI policy
- Form a governance group across compliance, IT, clinical, care management, patient advocate, human factors, and admin leads.
- Define a risk-based governance process for AI solutions
- Start a basic registry of existing AI solutions and readiness inventory.
- Pilot a governance process on a single low-risk or high-priority AI use case.
- Don’t reinvent the wheel. Find existing processes and adapt them.
- Identify which of these processes you can streamline at the network level, vs. approach at the organization-level.

CHAI Public Health Effort: Call to Action

CHAI is leading a **public health effort**. We have negotiated 2000 free licenses from ChatGPT/Claude and we are looking for public health officials to join communities of practice around 5 pilot use cases.

- We are ironing out the workflow to get maximum data protections and minimum bureaucracy (e.g. a self-serve, automated BAA, for individual ChatGPT Health licenses)
- We have tracked down multiple use cases and are collating into a long list of priority use cases to focus on

CALL TO ACTION: We need support in

a) voting on the use cases
b) disseminating the launch of this effort to apply to use cases
c) Email: lucy@chai.org if you or your networks are interested in participating

CHAI

About Coalition for Health AI (CHAI)

The CHAI (Coalition for Health AI) mission is to be the trusted source of guidelines for responsible AI in health that serves all. It aims to ensure high-quality care, foster trust among users, and meet the growing healthcare needs. CHAI membership is open and rapidly expanding with nearly 3000 organizations including health systems, patient advocacy groups, and a wide range of industry leaders and start-ups across the healthcare and technology ecosystems. CHAI is committed to convening and dialogue to achieve consensus. There is no limit to who can join and participate. Learn more about a CHAI membership here.

To learn more about CHAI or to inquire about membership:

www.chai.org

admin@chai.org

Questions about this presentation?

merage@chai.org

CHAI