Joint Commission and the Coalition for Health AI (CHAI) Release Governance Frameworks to Address AI in Health Care

Insights 09.25.25

PROFESSIONAL



ERIC M. FISH
Partner
Washington, D.C.

On September 17, 2025, The Joint Commission and the Coalition for Health AI (CHAI) jointly released The Responsible Use of AI in Healthcare (RUAIH), a non-binding guidance document that provides a comprehensive framework to assist health care organizations in implementing AI technologies safely, effectively, and ethically. The RUAIH represents the first major output of the recent partnership between The Joint Commission and CHAI.

Although a number of health systems have already developed models for AI governance, standardization across health care remains inconsistent due to the lack of federal or state requirements, rapid innovations that are outpacing legal requirements, and varying levels of AI integration within health systems and with third party vendors. Inconsistent standardization not only creates gaps in accountability across health care, but also risks patient safety, data protection, and equitable access to the benefits of AI.

Key Provisions of the Responsible Use of AI in Healthcare

The RUAIH sets governance benchmarks that reflect the common themes identified by The Joint Commission and CHAI during conversations with stakeholders, as well as review of existing frameworks such as the National Academy of Medicine's AI Code of Conduct Principles, the NIST AI Risk Management Framework, and the Bipartisan House Task Force Report on Artificial Intelligence. Distillation of the information gathered during this multi-year process led to the creation of seven actionable steps for health care organizations that include: (1) establishment of AI policy and governance structures, (2) enhancement to policies designed to protect patient data, (3) modification of existing data use agreements, (4) implementation of processes to monitor and evaluate performance of AI tools, (5) institution of a voluntary reporting process to capture AI-related safety or performance issues, (6) identification and mitigation of risks and bias that may threaten patient safety or access to care, and (7) efforts to train and educate all clinicians and staff members about the use of AI tools.

1. AI Policy and Governance Structures

The RUAIH recommends that health care systems implement a formalized and systematic governance framework for the deployment and oversight of AI tools across all organizational functions, including direct patient care, operations, and administrative support. Governance

structures should be grounded in risk-based management principles and address the selection, implementation, and life cycle management of internally developed AI solutions and third-party vendor products. Organizations are advised to ensure that AI governance policies are harmonized with existing internal governance frameworks, informed by ethical standards, and subject to regular updates that reflect organizational strategy and evolving regulatory requirements. The RUAIH further recommends the establishment of a multidisciplinary governance team that incorporates technical experts with representation from executive leadership, data and cybersecurity specialists, frontline staff, providers, and patients. Additionally, the RUAIH encourages Boards of Directors and Boards of Trustees to incorporate regular review of AI governance structures, as well as associated use and outcome data, into their regular oversight activities to ensure compliance with their fiduciary responsibilities.

2. Patient Privacy and Transparency

The RUAIH recognizes that strong data privacy policies are essential to safeguard protected and other sensitive information, ensure regulatory compliance, and maintain patient trust. The guidelines call for the education of staff and patients on how data is collected, used, and repurposed into AI tools. The RUAIH directs organizations to invest in educational materials for patients and their family members that promote a full understanding of how information shared with or obtained with AI tools impacts the course of care and obtain consent when relevant.

3. Data Security and Data Use Protections

Health care organizations are already subject to established law and regulations, such as HIPAA's Privacy, Security, and Breach Notification Rules, which serve as a baseline foundation to the guidance presented in the RUAIH. The RUAIH encourages health care organizations to redouble existing compliance efforts and include additional AI-specific enhancements that include: (1) defining the permitted used of exported data and granting of rights tied to model outputs, performance data, and monetization, (2) ensuring that only the minimum amounts of data for a specific use are exported to third parties, (3) explicitly prohibiting the re-identification of previously de-identified data, and (4) strengthening third-party obligations and audit rights. Organizations are directed to review The Joint Commission's Responsible Use of Health Data framework and implement the framework as necessary.

4. Ongoing Quality Monitoring

During the procurement process, organizations should request detailed information about the testing and validation of AI tools, the evaluation of bias, and whether the willingness of the vendor to refine the tool once deployed in the organization's ecosystem. The guidance also encourages proactive monitoring that includes regular validation and testing, comparing AI outputs to known sets of performance data, assessing use-case relevant outcomes, and user confidence in the tools. Further, the guidance suggests that an organization tailor its governance structure to monitor the risks associated with the tool, with patient-facing, clinical tools assessed more frequently and, in more detail, than administrative tools.

5. Voluntary, Blinded Reporting of AI Safety Related Events

The RUAIH recognizes that active self-regulation can facilitate the concept of keeping humans-in-the-loop and is critical to encourage innovation and the safe integration of AI tools into health care. The guidance treats AI-related safety events as analogous to patient safety events and encourages organizations to capture the details about the incident in internal reporting systems and share de-identified details with patient safety organizations or other relevant organizations. If the AI tool has been classified as a regulated device by the FDA, organizations should account for the incident internally as well as follow the FDA reporting pathway. Governance policies should encourage staff members at all levels to voluntarily report incidents, near-misses, and other performance issues. The RUAIH suggests the adoption of strong reporting practices may insulate health care organizations from changes in law and regulation that may stifle developing AI programs and limit the ability of organizations to take advantage of future innovations.

6. Risk and Bias Assessment

The RUAIH instructs that AI models trained on data that is neither complete nor reflective of the patient population served creates additional risks of harm, misdiagnosis, and diminished reliability. The RUAIH notes that a lack of diversity is especially concerning health care, where individual patient characteristics are often determinative to treatment plans and eventual

outcomes. Accordingly, the RUAIH encourages health care organizations to systematically evaluate whether datasets and model outputs may suffer from deficiencies in training data and appropriately mitigate any bias to prevent disproportionate or unsafe outcomes. Mitigation steps include a detailed review of training data, questioning whether the tools were subject to bias detection assessment during development, and inquiry into whether the models are sufficiently tuned to local data and data specific to the patient populations served to mitigate the effects of bias in the output.

7. Education and Training

The RUAIH's final recommendation emphasizes the importance of including education and training in any governance framework. Organizations should define and document authorization protocols specifying which individuals may use each AI tool and implement role-specific training to ensure appropriate use. All users of AI tools, including providers, must have access to relevant information about each tool and must demonstrate familiarity with the organization's applicable policies and procedures governing AI. Furthermore, organizations should implement enterprise-wide educational programs to promote common terminology and an understanding of organizational policies to help ensure a consistent foundation of knowledge across the workforce.

Actionable Insights for Your Health Care Organization

Irrespective of an organization's size or maturity of governance structures, health care organizations should treat the development and continual refinement of AI governance structures as a strategic priority to ensure appropriate oversight and risk management.

As an initial step, health care organizations should conduct a systematic review of existing contracts with third-party vendors. Existing agreements will lack provisions addressing emerging standards regarding data access, ownership, or permissible use in connection with AI technologies. Many of these agreements may also preclude vendor accountability through legal disclaimers, suffer from limited representations and warranties and severe limitations on vendor liability, and have other limitations on available remedies. Early identification and remediation of existing contractual gaps will mitigate legal and operational exposure and better position an organization for future negotiation.

At the same time, organizations should prioritize educational initiatives for both staff and patients. Investing in education represents a comparatively low-cost, but high-impact, measure that establishes the foundation for broader adoption and sustained risk mitigation. Organizations should design training and educational materials specific to the particular AI use cases deployed within the organization. Incorporating the CHAI Applied Model Card to internal resources and procurement documentation provides staff transparency on the creation and contents of the algorithm. These model cards are commonly referred to as "AI Nutrition Labels" and include information on known risks and limitations, the data used to train algorithms, bias mitigation approaches, and any ongoing maintenance and improvements. Additionally, providing of physicians the ability to review the model cards associated with tools they are using aligns with recommendations found in the FSMB's 2024 report, *Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice* , and further supports the responsible use of AI.

Organizations should communicate content designed for patients in plain, non-technical language, deliberately avoiding legalistic or overly technical terminology, to maximize comprehension, encourage engagement, and facilitate acceptance of AI-enabled tools. By thoughtfully addressing governance in the still early stages of AI adoption, health care organizations can promote the safe, effective, and ethical use of AI as a tool to enhance care and maintain accountability and trust in clinical care.

Undertaking AI compliance efforts now may reduce liability risk from AI related litigation, which is anticipated to increase as adoption of technology increased. These efforts will also better position organization to respond more nimbly to potential regulations from federal and state actors that will impact plans to integrate AI into all aspects of their business. Increasingly, legislators and policymakers are signaling interest in establishing some standards that both safeguard patients and regulate provider use of AI, while carefully balancing these protections against the need to sustain innovation. State-level experiments are already serving as precursors to this anticipated regulatory trajectory. For instance, HB1915, introduced in Oklahoma

during the last legislative session, proposed a comprehensive framework requiring health care organizations to implement quality assurance programs for AI devices, establish governance bodies to oversee AI adoption, and conduct ongoing performance evaluations tied to patient outcomes. Such efforts, coupled with greater legislative interest in establishing AI guardrails, foreshadow the likely introduction of similar bills in 2026.

A long-term challenge for organizations implementing any governance framework will be addressing the practical considerations that extend beyond established guidance. Organizations must carefully design the structure and functions of a multidisciplinary governance team and ensure that AI-specific governance does not devolve into an additional administrative burden for staff and clinicians who may already be working under significant constraints. Furthermore, organizations must determine appropriate criteria for evaluating AI models, identify metrics to monitor performance degradation, and establish protocols for responding effectively when failures occur. These tasks are inherently complex for a rapidly developing area of health care and will require strong commitments from those involved. However, emerging initiatives, such as those from Stanford's Institute for Human-Centered Artificial Intelligence (HAI), have begun to develop evaluation tools that assess the ability of AI systems to carry out clinically relevant tasks, including retrieving patient data, ordering diagnostic tests, and prescribing medications. Existing policies and procedures for reporting and quality monitoring, such as the long-standing FDA requirements for post-clearance surveillance monitoring and reporting of adverse drug and device events. Review of current processes for parallels and engagement with those creating new tools to review and validate AI will be beneficial steps as organizations look at ways to operationalize governance policies.

Key Actions to Consider Now

• Create or review existing governance frameworks to ensure they are AI-ready
• Review existing data contracts with third party vendors and revise as necessary to protect against unauthorized use of sensitive information and limit future exposure
• Make 'nutrition labels' –such as the CHAI Model Card—available for review by physicians to enhance transparency and support responsible use of AI tools
• Engage patients and create educational materials to ensure informed consent and a greater understanding of how AI will impact the delivery of care

What to Expect in the Future

The RUAIH is the first in a series of projects between The Joint Commission and CHAI, designed to help the industry align on practices that protect patients from AI-related risks while also improving administrative, operational, and clinical outcomes through effective AI use. Although created in part by the Joint Commission, implementation of the recommendations does not impact any accreditation or certifications issued by the organization. However, adoption of the recommendations may reduce an organization's exposure to claims of negligent or improper incorporation of AI into the delivery of care, as well as correspond to maintenance of best practices required by applicable insurance policies. It is foreseeable that insurers may require or provide better terms to insureds who adopt robust governance frameworks and engage in measures that would reduce litigation exposure.

Additional guidance documents are in production and expected for release later in 2025 and into 2026. CHAI and The Joint Commission are developing governance playbooks informed by a series of workshops designed to capture input from hospitals and health systems of varying sizes, locations, and capabilities. The playbooks will build on the initial guidance by incorporating community feedback and providing more practical, implementation-focused direction. The Joint Commission intends to develop a voluntary AI certification program derived from the finalized playbooks. The certification program will be made available to The Joint Commission's network of more than 22,000 accredited and certified health care organizations nationwide.

RELATED CAPABILITIES

Digital Health

Government Relations and Public Policy

Health Information Privacy and Security