VIEWPOINT

AI IN MEDICINE

New Guidance on Responsible Use of AI

Sofia Palmieri, PhD; Christopher T. Robertson, JD, PhD; I. Glenn Cohen, JD

Artificial intelligence (AI) has moved rapidly into clinical application, such as imaging and other diagnostics, ambient scribing, and decision support. It has brought both potential benefits and important risks. In the US, it is unclear who has the responsibility to set the rules, validate the models, and, ultimately, pay when things go wrong. Against this backdrop, the Joint Commission (TJC, the primary US hospital accreditor) has partnered with the Coalition for Health AI (CHAI, a clinician-led group that includes industry and stakeholders). Together, they have issued new guidance on responsible use of AI in health care,[1] which builds on prior statements from many groups, including the American Medical Association.

Core recommendations include establishing multidisciplinary AI governance committees, validating models on local patient data and workflows before deployment, and instituting continuous postmarket monitoring for drift, performance degradation, or bias[1] (**Table**).

The CHAI and TJC guidance leaves many details to health care organizations to operationalize. For example, it seems to take a surprisingly hard line on engaging with patients, advising that "[w]hen appropriate, patients should be notified when AI directly impacts their care.... Where and when relevant, consent should be obtained."[1] Physicians typically do not disclose or ask for consent for each technologic device that they may deploy (eg, this scalpel vs that one). Is AI different in kind? Does every radiologist report need a disclaimer about AI use or even preconsent for that use? What are patients to make of such disclosures?[2] The hard questions persist in the guideline's words *appropriate* and *relevant*.

For health care institutions, compliance with the new guidance will require governance structures, monitoring dashboards, incident reporting mechanisms, and documentation.[1] These measures parallel established quality and safety frameworks (such as accreditation requirements, patient safety event reporting, and risk management programs) so that AI oversight can be integrated into structures already familiar to hospitals.

However, compliance will plausibly require substantial new investments. TJC and CHAI call on each hospital to set up a formal governance structure for AI, with policies and procedures for selecting, implementing, and overseeing AI tools. This function will require executives with appropriate AI expertise to assemble and lead a multidisciplinary AI committee. Perhaps one of the most resource-intensive elements is continuous monitoring of AI performance and outcomes.[3] The hospital must measure accuracy, errors or adverse events associated with AI recommendations, and performance over time, including equity and bias across patient populations. Although big academic centers have the capacity to rigorously validate AI tools, nearly two-thirds of US hospitals have fewer than 100 beds, and many operate in rural settings with resource constraints (eg, lack of expanded Medicaid coverage).[4] A smaller community hospital is more likely to procure off-the-shelf technology and trust vendor assurances. Conducting an independent bias audit and risk analysis might be beyond their expertise and resources.

These challenges explain why research shows that although 61% of hospitals using AI had conducted a local validation for accuracy, only 44% evaluated their models for bias by using their own data.[4] For many hospitals, the CHAI and TJC guidance may be aspirational.

What are the consequences of noncompliance? TJC has not (yet) incorporated these standards into its accreditation requirements, which are a precondition for billing public and private insurers. For now, these guidelines will plausibly affect liability. Courts have held

Table. Summary of TJC and CHAI Guidelines

Quoted guideline heading	Paraphrased key requirements
AI policies and governance structures	Require expertise and cross-functional representation (leadership, IT, compliance, clinicians, safety, privacy)
	Oversee selection, life cycle, compliance, and risk management of AI tools
	Report AI activities and adverse events to the governing board
Patient privacy and transparency	Establish data access, use, and protection policies consistent with HIPAA and state law
	Disclose and educate patients on AI use affecting their care; obtain consent when relevant
Data security and data use protections	Encrypt data in transit and at rest
	Apply strict access controls and audit logs
	Perform regular security assessments and maintain an incident response plan
Ongoing quality monitoring	Require vendors to supply validation and bias-testing data
	Continuously test, validate, and monitor AI performance and bias after deployment
	Scale monitoring frequency to risk proximity to patient care
Voluntary, blinded reporting of AI safety-related events	Implement confidential, deidentified reporting of AI safety events to independent entities
	Incorporate AI incidents into existing safety reporting systems
Risk and bias assessment	Identify, document, and mitigate risks and biases before and after deployment
	Request vendor bias data and model cards
	Use representative training and validation datasets; test on local data
	Perform regular bias audits
Education and training	Provide role-specific AI training for clinicians and staff
	Ensure documentation and guidance on each AI tool

Abbreviations: AI, artificial intelligence; CHAI, Coalition for Health AI; HIPAA, Health Insurance Portability and Accountability Act; IT, information technology; TJC, the Joint Commission.

physicians and health care organizations to a reasonable care standard, in part informed by persuasive guidance from reputable organizations, such as TJC and CHAI. A physician defendant could use these guidelines to emphasize that the responsibility is on health care organizations, thus reducing or avoiding liability for the physician.[3] On the other hand, hospitals are in a bind if they are unable to comply with onerous guidelines. Still, courts may take resource limitations into account when assigning liability for negligence.

In addition to distributional concerns that cut across health care organizations, there are also broader efficiency concerns. Does it make sense for each of the thousands of health care organizations to do all this work independently? On the contrary, a rational system could have much of this work done once authoritatively. Likewise, we do not expect each hospital to perform a de novo evaluation of the competence of every physician; instead, hospital staff committees rely on state licensure agencies. Yet no such agency presently exists for AI.

TJC and CHAI guidelines attempt to make the best of a world that is already fragmented, in which the US federal government has not stepped forward to provide a comprehensive mechanism to assess AI services that are used in interstate commerce and funded through federal health care programs.[5] To its credit, CHAI has also proposed to fill this gap by creating a network of private assurance laboratories to vet algorithms so that health care organizations can rely on them. The Biden administration was aligned with that approach, but the Trump administration opposes it. The concern is that a private, industry-heavy coalition could entrench dominant vendors and act as a de facto regulator.[6]

In part, the federal gap is created by a US Food and Drug Administration (FDA) decision to focus its enforcement on automated AI technologies when there is not a physician in the loop. A much larger FDA budget would be necessary to reach all AI health care technologies, and industry would have to invest heavily in meeting the demands for science that FDA would require. Such investments might save money for the health care system overall by reducing local redundancies and improving overall performance. In late September 2025, the FDA requested public comment on a potential new approach to "practical approaches to measuring and evaluating the performance of AI-enabled medical devices in the field, including strategies for identifying and managing performance drift, such as detecting changes in input and output."[7]

Ultimately TJC and CHAI have created an important framework aiming for enhanced patient safety, improved patient outcomes, stronger data protection, increased trust, and greater operational efficiency. Without a unified system for evaluating AI products, hospitals will continue to shoulder fragmented responsibilities, creating inefficiencies and widening disparities in safety, quality, and potentially liability.

ARTICLE INFORMATION

Author Affiliations: Harvard Law School, Cambridge, Massachusetts (Palmieri, Robertson, Cohen); Boston University School of Law, Boston, Massachusetts (Robertson).

Corresponding Author: Christopher T. Robertson, JD, PhD, Boston University School of Law, 765 Commonwealth Ave, Boston, MA 02215 (ctr00@bu.edu).

Published Online: December 3, 2025. doi:10.1001/jama.2025.23059

Conflict of Interest Disclosures: Mr Cohen is a member of the Bayer Bioethics Council; a bioethics advisor for Bexorg; an advisor for World Class Health and Manhattan Neuroscience LLP; recently concluded service as the chair of the ethics advisory board for Illumina; was compensated for speaking at events organized by Philips with the *Washington Post*, as well as the Doctors Company, attending the Transformational Therapeutics Leadership Forum organized by Galen Atlantica; and was retained as an expert in health privacy, gender-affirming care, and reproductive technology lawsuits. No other disclosures were reported.

Funding/Support: Dr Palmieri and Mr Cohen were supported in part by a Novo Nordisk Foundation grant for a scientifically independent International Collaborative Bioscience Innovation & Law (Inter-CeBIL) Programme (grant NNF23SA0087056).

Role of the Funder/Sponsor: The funder had no role in the preparation, review, or approval of the manuscript or decision to submit the manuscript for publication.

REFERENCES

1. Joint Commission. Joint Commission and Coalition for Health AI (CHAI) responsible use of AI in healthcare (RUAIH). Accessed October 31, 2025. https://digitalassets.jointcommission.org/api/public/content/dcfcf4f1a0cc45cdb526b3cb034c68c2?v=3edb8a95

2. Robertson C, Woods A, Bergstrand K, Findley J, Balser C, Slepian MJ. Diverse patients' attitudes towards artificial intelligence (AI) in diagnosis. *PLOS Digit Health*. 2023;2(5):e0000237. doi:10.1371/journal.pdig.0000237

3. Price WN II, Gerke S, Cohen IG. Potential liability for physicians using artificial intelligence. *JAMA*. 2019;322(18):1765-1766. doi:10.1001/jama.2019.15064

4. Nong P, Adler-Milstein J, Apathy NC, Holmgren AJ, Everson J. Current use and evaluation of artificial intelligence and predictive models in US hospitals. *Health Aff (Millwood)*. 2025;44(1):90-98. doi:10.1377/hlthaff.2024.00842

5. Babic B, Gerke S, Evgeniou T, Cohen IG. Algorithms on regulatory lockdown in medicine. *Science*. 2019;366(6470):1202-1204. doi:10.1126/science.aay9547

6. Reader R. Trump to the health and tech giants: AI is not your "cartel". Politico. Published October 1, 2025. Accessed October 31, 2025. https://www.politico.com/news/2025/10/01/trump-ai-artificial-intelligence-regulation-hhs-00590902

7. Request for public comment: measuring and evaluating artificial intelligence–enabled medical device performance in the real-world. US Food and Drug Administration. Updated September 30, 2025. Accessed October 31, 2025. https://www.fda.gov/medical-devices/digital-health-center-excellence/request-public-comment-measuring-and-evaluating-artificial-intelligence-enabled-medical-device