


Is your AI governance enabling innovation?

Five questions for healthcare leaders

As AI becomes embedded in care delivery, these five questions
can help leaders ensure governance enables safe innovation
rather than constraining it.



Jack Kain, PharmD, Principal Architect
and Clinical Consultant, Micromedex



Inside this whitepaper

- Introduction: When AI meets clinical reality
- Governance as an enabler of clinical AI
- Practical guidance: Five questions that turn governance into enablement
- From AI governance to product enablement

When AI meets clinical reality

A clinician needs to check the evidence to support a treatment decision. They use an AI tool and the response is fast, confident, and cited.

But what happens when the answer isn't straightforward?

Two guidelines conflict. The evidence changed last month. The patient has renal impairment. The lab result cannot support the inference being requested. The citation technically supports the sentence but not the clinical conclusion. A dose calculation produces a single answer when multiple clinically defensible methods exist. The clinician clicks accept without reviewing the underlying evidence.

An AI enabled diagnostic reduces unnecessary biopsies and lowers health system costs but also reduces procedure volume and revenue for the surgeon expected to adopt it, creating an incentive conflict.

These are just some of the situations clinical AI governance will increasingly need to address.

That is where governance becomes real.

And where governance can become an enabler of innovation rather than a constraint on it.

Governance as an enabler of clinical AI

Health systems are moving beyond the question of whether they will use AI. In a national survey of 2,174 U.S. hospitals published in 2025, 31.5% reported already using generative AI integrated with their EHR, while another 24.7% planned to implement it within the next year.¹


The harder questions are increasingly about where AI belongs, what organizations should expect from it, and how leaders can expand use without introducing unacceptable clinical, operational, or organizational risk.

Those decisions cannot be answered by the performance of AI models alone.

Clinical AI enters an environment already shaped by competing evidence, complex patients, established workflows, economic incentives, clinician judgment, and accountability for patient care. Governance provides a structure for navigating that complexity.

In a recent case study involving 15 health systems, the Coalition for Health AI (CHAI) reported that evaluating a single AI solution can take two weeks to two months, illustrating how greater transparency could reduce governance friction.² Done well, governance should not simply tell an organization where AI must stop. It should help leaders determine where AI can responsibly go next.

U.S. HOSPITALS SURVEYED ON THEIR UPTAKE OF GENERATIVE AI INTEGRATED WITHIN THE EHR¹



Evaluating a single AI solution can take two weeks to two months.²

Five questions that turn governance into enablement

Clinical AI governance becomes more actionable when leaders move beyond broad principles and ask specific questions about the products entering clinical workflows.

Together, the following five questions help healthcare organizations scale clinical AI responsibly.

1. Evidence governance: What evidence should clinicians trust?

An AI model's ability to answer a question does not establish that a health system should trust the answer.

For health system leaders, one of the most important questions is therefore not simply which model is being used, but what evidence environment sits behind it.

What sources can the system access? How are clinical guidelines, regulatory information, primary literature, and other evidence prioritized? What happens when authoritative sources disagree? How quickly does new evidence enter the system? Who is responsible for determining whether a source remains appropriate for clinical use?

Retrieval-augmented generation (RAG) provides one mechanism for governing the evidence environment by grounding AI responses in selected sources rather than relying solely on a model's pre-trained knowledge.

Emerging evidence supports the role of RAG in reducing hallucinations in healthcare AI. A systematic review of 44 studies found that RAG was the most frequently evaluated mitigation strategy, assessed in 18 of the 44 included studies, with RAG approaches demonstrating 30–50% reductions in factual errors compared with base models without retrieval augmentation where quantitative

data were available. Knowledge graph approaches were evaluated in 12 studies, while combined approaches, such as retrieval plus verification, generally outperformed single-method controls.³

But RAG alone does not establish clinical authority. What information is retrieved, how sources are prioritized, how quickly evidence is updated, and who determines which sources are appropriate for clinical use remain governance decisions. Evidence hierarchy, source freshness, editorial oversight, and the management of conflicting recommendations remain essential.

This distinction matters when evaluating clinical AI vendors as well as internally developed systems. A sophisticated model connected to an inadequately governed evidence environment can still produce a sophisticated wrong answer.

Governance starts before the answer is generated.

2. Provenance: Can clinicians see what supports the answer?

Health systems should expect more from transparency than a reference list.

An AI response may contain credible citations while leaving the clinician unable to determine which evidence supports a specific clinical assertion. A source may even support an individual sentence without supporting the broader clinical conclusion being presented.

Sentence-level provenance provides a stronger standard.

Emerging evidence also suggests that granular provenance can be practical for clinicians. In a 2026 formative usability study involving 46 clinician interactions, a click-to-inspect interface connecting individual AI-generated statements to supporting source material achieved a mean System Usability Scale score of 86.25 out of 100. Participants identified rapid access to supporting evidence as important for calibrating trust during initial chart review.⁴

Instead of forcing a busy clinician to guess which reference backs up an answer, clinical AI must adopt an academic publishing standard. By attaching footnotes to nearly every single sentence generated and anchoring clear links to the underlying monographs or studies, clinicians can audit the response exactly like a peer-reviewed paper.

Furthermore, true clinical utility requires streamlining the workflow within the response itself. If a drug is mentioned multiple times, it shouldn't require hunting down a single hyperlink; the text must actively maintain a direct link to the source right where it repeats in the body of the text.

Clinicians should be able to connect important clinical assertions to the evidence supporting them without reconstructing an AI system's reasoning from a collection of references.

For health system leaders, this is more than a user experience feature. It affects clinical trust, reviewability, accountability, and the organization's ability to understand how an AI-enabled decision was reached. Provenance shows what evidence supports an AI-generated assertion. Auditability goes one step further by allowing an organization to reconstruct what the system produced, what evidence was available, and how the output was reviewed or used. As AI becomes more embedded in clinical workflows, both become important foundations for accountability.

Sentence-level provenance provides a stronger standard.

3. Uncertainty: Can the system recognize when medicine does not have one answer?

AI products are often evaluated on their ability to provide answers. Clinical medicine also requires knowing when a single answer is not justified.

Guidelines may conflict. Evidence may differ in strength. Patient characteristics may change the interpretation. Several accepted methodologies may produce different results. Sometimes the available information simply cannot support the inference being requested.

Consider a clinical dose calculation for which multiple defensible methodologies exist. Returning one precise number without acknowledging those alternatives may create the appearance of certainty where medicine itself contains uncertainty.

Health systems should therefore evaluate not only whether an AI system can answer a clinical question, but whether it can appropriately communicate ranges, alternatives, limitations, conditional answers, and circumstances in which it should defer.

Good governance governs confidence as well as content.

4. Human oversight: Where does clinical judgment add value?

“Human in the loop” is increasingly treated as a universal answer to AI risk. For health systems, that phrase deserves greater scrutiny.

A clinician clicking accept on an AI generated output does not necessarily constitute meaningful oversight.

The appropriate level of human involvement should depend on what the system is doing and the consequence of being wrong. Summarizing evidence presents a different risk than recommending therapy. Supporting a calculation presents a different risk than autonomously executing a clinical action.

This suggests a more useful model for health systems: risk calibrated oversight.

The executive question becomes not simply whether a human is involved, but where human judgment provides meaningful protection, what the clinician is expected to verify, and how that responsibility is incorporated into workflow.

Those expectations may also change as AI performance improves. Governance must be capable of changing with them.

5. Continuous governance: What happens after implementation?

Approval is not the end of governance.

The clinical environment that justified an AI product at deployment will continue to change. Evidence changes. Guidelines change. Models and underlying technology change. Clinical workflows evolve. Users find new applications for products that their developers may never have anticipated.

Health systems therefore need governance that extends beyond procurement and initial validation. The need for continued oversight is already apparent. In a 2025 national study, only 53.3% of hospitals using predictive AI reported that most or all of their models were monitored after deployment.¹

That does not mean repeatedly stopping innovation for another review. It means establishing mechanisms for monitoring performance, maintaining evidence currency, evaluating meaningful changes, identifying emerging failure patterns, and determining when additional review is warranted.

Governance must operate at the speed of the products it governs.

Only 53.3% of hospitals using predictive AI in study reported that most or all of their models were monitored after deployment.¹

53.3%



From AI governance to product enablement

For health system leaders, these five governance principles ultimately translate into capabilities that determine whether clinical AI can safely deliver value:

1. Evidence governance determines what a clinical AI product can responsibly know.
2. Provenance determines whether clinicians can understand what supports its assertions.
3. Uncertainty management determines whether the product recognizes the limits of its own answers.
4. Human oversight determines where clinical judgment should remain decisive.
5. Continuous governance determines whether those assurances survive after deployment.

These are not simply controls placed around innovation. They are part of the infrastructure required to scale it.

The health systems that succeed with clinical AI may not be those that adopt the most tools or move the fastest. They will be those that develop the governance capabilities to distinguish where AI can create meaningful clinical and operational value, where additional safeguards are necessary, and where the technology is not yet ready.

Governance, in that sense, is not the opposite of innovation.

It is what allows innovation to move from promising technology to trusted clinical infrastructure.

WHAT'S NEXT?

[Learn more about the Micromedex approach to AI here.](#)

About Micromedex

Micromedex by Merative delivers award-winning clinical decision support to professionals across the healthcare ecosystem. Named 2026 Best in KLAS point-of-care drug reference, Micromedex is the market leader and now includes AI-powered search to accelerate workflows. The drug, toxicology and patient education content is trusted for its expert curation and ease of use, supporting informed, evidence-based decisions globally.

Learn more at micromedex.com

References

1. Everson J, Nong P, Richwine C. Uptake of Generative AI Integrated With Electronic Health Records in US Hospitals. *JAMA Netw Open*. 2025;8(12):e2549463. doi:10.1001/jamanetworkopen.2025.49463.
2. Coalition for Health AI. Case Study: How Transparency Can Accelerate Responsible AI Adoption – CHAI's Model Card Registry. Published August 27, 2026.
3. Basu S, Huynh B. Mitigating hallucinations in healthcare AI: a systematic review of evidence-based strategies. *BMC Health Serv Res*. 2026;26:1115. doi:10.1186/s12913-026-14851-1.
4. Parambath A, Pulpo G, Hartman V. Sentence-Level Provenance for AI Medical Record Summarization: Formative Usability Evaluation of a Click-to-Inspect Interface. *JMIR Hum Factors*. Published online July 30, 2026. doi:10.2196/95644.

© Merative US L.P. 2026. All Rights Reserved.

Produced in the United States of America
September 2026

Merative and the Merative logo are trademarks of Merative US L.P. Other product and service names might be trademarks of Merative or other companies.

The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on Merative's product plans and strategy as of the date of this publication, which are subject to change by Merative without notice. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from Merative, or stating or implying that any activities undertaken by you will result in any specific performance results. Merative products are warranted according to the terms and conditions of the agreements under which they are provided.

MDX-000000088 Rev 1.0

