

# AI Catalyst

July 2024

## Five Lessons for Structured, Impactful AI Governance

*Across the last six months, The Health Management Academy conducted research and interviewed more than 100 health care leaders about AI governance. These leaders included C-Suite executives, informatics officers, innovation leads, and data experts across Leading Health Systems (LHS). The following insights are the net result of this research and interactions. AI Catalyst will continue to provide additional guidance and discussion opportunities on AI governance as technology advances.*

---

Once you get past all the jargon, when Leading Health Systems (LHS) talk about AI governance, what they're really looking for is a way to make sense of and reassert control over an unwieldy, volatile, and evolving technological trend. And that's fair. Whether we talk about ideal committee structures, parameters for AI pilot life cycles, or methods for data bias mitigation, discussion around AI governance is grounded in a desire for clarity, actionability, and accountability.

IBM defines AI governance as the ability to **direct, manage, and monitor** AI activities across your organization. Simple words that underline a far-from-simple task. That's why THMA has put together five action steps for structured, impactful AI governance:

1. Understand what you're trying to govern: Boost AI literacy through training, AI awareness through organizational use-case audits
2. Don't reinvent the wheel: Build AI governance on your existing procurement and governance frameworks.
3. Build awareness of AI-specific ethical challenges into training, pilot design, and evaluation processes
4. Don't ignore risk mitigation strategies: Prepare to navigate worst-case AI scenarios
5. Orient your governance structure to rapid evolution in the AI space

More information on each step is below.

### 1. Understand what you're trying to govern: Boost AI literacy through training, AI awareness through organizational use-case audits

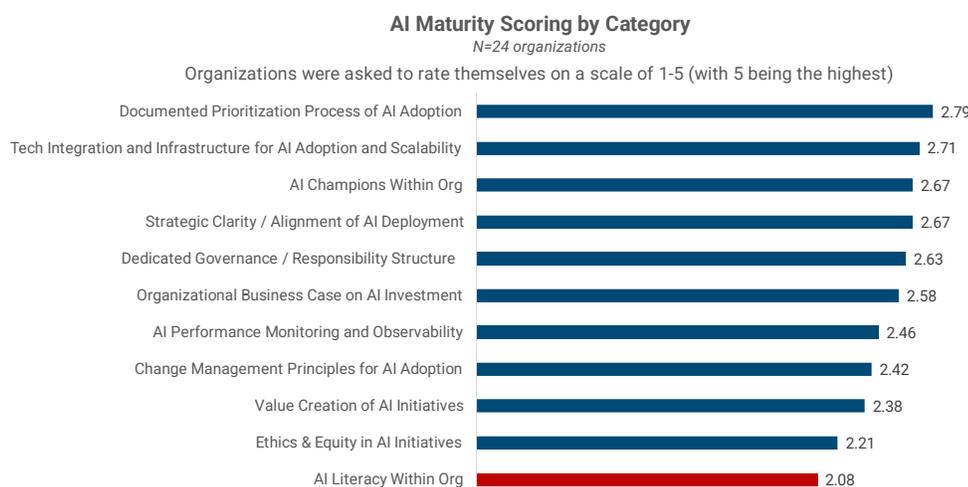
You can't govern what you don't understand or fail to recognize; that's why the lack of AI literacy at the organizational level precludes effective governance. It is difficult to leverage existing investment evaluation frameworks (such as for software or novel tech investment) if colleagues can't even agree what processes fall under the AI umbrella. And when it comes to identifying the ROI of AI, many executives have told us that competing definitions of AI make it hard to pin down just what the initial "investment" is.

The biggest frustration that many healthcare leaders express when it comes to building effective AI governance is defining just what it is that they are trying to govern. Whether it's peers who don't know

the difference between AI and automation, clinicians who don't realize that they're already using AI to assist with diagnostics, or IT leaders struggling to monitor model performance across an ever-increasing list of algorithms, the governance challenge comes back to a lack of organizational awareness and literacy around AI.

In fact, our research suggests that AI literacy – rather than forming the basis of AI maturity efforts across LHS – is often overlooked.

## AI Prioritization Processes and Tech Integration Rank as Most Advanced



Source: THMA research and analysis.

As the graph above shows, a survey of informatics and data leads across 24 LHS asked them to rate their organizations' AI maturity across 11 categories. When averaged, "AI literacy within the organization" was the lowest ranked of all categories.

That's why it's imperative that LHS foster a common understanding of AI use at their organization. **LHS should incorporate learning about AI into continuing education for leaders**, building on existing compliance-based training around data use in healthcare settings. Leaders can then socialize this knowledge and cascade down through the organization, building on tried-and-true change management best-practices seen during EMR adoption, DEI/health equity initiatives, and other strategic endeavors rolled out by health systems. AI Catalyst stands ready to help, with our [AI Bootcamp Series](#) designed to provide leaders with the scoped healthcare-contextualized knowledge and know-how to navigate AI adoption across their organization and specifically within their functional area.

But effective governance requires more than baseline AI literacy: You need baseline awareness of how AI is currently being used at your organization. **Deploy regular, anonymous audits to reveal the full scope of AI adoption, including innovative unofficial uses, and identify high-risk applications requiring immediate intervention.** Consider creating a "no-fault" disclosure period to build trust between IT/compliance teams and frontline staff and provide your governance leads with valuable insights into why employees are turning to unauthorized tools. Additionally, use these routine audits to ensure that leaders of functional areas understand the penetration and proliferation of AI use within their through-lines, helping to boost oversight of AI adoption.

## 2. Don't reinvent the wheel: Build AI governance off existing procurement and governance frameworks at your organization

AI may come across as belonging to its own class of technological innovation, but in practice, it is inextricably linked to many other novel technologies, such as process automation, cloud computing, or the internet of things. In healthcare applications, AI is heavily reliant on data as found within the EMR or within claims databases. Altogether, this means that **AI is not a stand-alone technology—and it cannot be governed as a stand-alone technology**. What's more, just as AI use is inextricably linked to other innovations that healthcare has embraced, AI governance is naturally aligned with many best-practices already in place to govern these already-adopted innovations.

**Undertake an audit of investment frameworks, steering groups, and standing committees, both clinical and non-clinical, to see what governance infrastructures exist at your organization.** Examples of existing governance structures include any bodies dealing with the following areas:

- Data governance
- Novel technology adoption (such as EMR)
- Major capital expenditures (clinical and non-clinical)
- Digital health
- Social determinants of health
- Health equity
- DEI
- Research and research testing ethics

That doesn't mean you *shouldn't* create an AI-focused steering group for governance purposes. AI raises new considerations around risk and reward that may well require dedicated oversight. But any new, AI-specific governance structure should have a mandate or process that expands, rather than re-creates, any existing data or clinical governance bodies at your organization.

*What AI Catalyst members are saying:*

One academic medical center created councils spanning education, research, and clinical departments to steer AI governance decisions, including algorithm monitoring and pilot scaling. An executive explained, "The governance is handled through faculty participation in these three key areas. In IT, we play a role of facilitation and supporting the steering from the choices that get made by these councils."

Another approach is to build AI due diligence work across existing IT and business governance committees. Speaking about how to make the decision of whether to scale an AI pilot, one CIO shared, "We run all of these sorts of decisions through our different IT business governance committees. We have a calendar for planning, a calendar for implementation. Pilots and proof of concepts happen throughout the year, but the go or no-go decision really indexes back to our existing investment governance process."

## 3. Build awareness of AI-specific ethical challenges into training, pilot design, and evaluation processes

Because AI adoption in healthcare is predicated on the use of patient data, it shares many of the same risks and ethical concerns that surround data use in general. Patient privacy, bias management, and cybersecurity are all important considerations when it comes to governing AI use. Yet, the way AI works makes many of these concerns especially acute.

Three risks inherent to AI use create particularly significant ethical challenges:

1. Many AI models function as unexplainable “black boxes,” meaning that it is difficult for end users to understand the process by which an AI algorithm has arrived at a given output.
2. AI is susceptible to data bias. Other data-based algorithms share this vulnerability – but because AI operates as a data enhancer, it can enhance/perpetuate existing biases indefinitely.
3. AI sometimes “hallucinates,” creating outputs that are nonsensical or simply inaccurate. These errors can be caused by a variety of factors, including insufficient training data, incorrect assumptions made by the model, or biases in the data used to train the model

**Because a system-wide governance body cannot be expected to anticipate every manifestation of AI risk, it’s crucial to spread awareness and education of these dangers throughout your system.**

Healthcare leaders should identify concrete examples of how these three risks may surface within their specific functional areas and use these examples as the basis of further training. LHS should broaden information security and data privacy training to include specific focus on AI. Annual compliance training for employees around data should be expanded to address the use of data within AI.

Bias mitigation, particularly from a health equity perspective, is imperative. For effective AI governance of pilot use cases, disaggregate data by key demographic factors from the start when evaluating model performance. **Leverage specialized equity subject matter experts throughout your day-to-day governance process.** For example, one CMIO we spoke with mentioned that they have incorporated experts from their Center for Health Equity and Social Justice directly into the evaluation process of new AI use cases. For systems without such centers, involving experts in public health and health policy research is suggested.

*What AI Catalyst members are saying:*

“Prospective tracking of potential health inequities is a key tool to prevent inequities pre-AI and during AI use. It’s easy to say, harder to get done, but crucial,” emphasized a healthcare executive at our July governance challenge huddle.

#### **4. Don’t ignore risk mitigation strategies: Prepare to navigate worst-case AI scenarios**

It’s important to recognize that the use of AI is not free from error, bias, or adverse outcomes. AI often takes the form of internet-connected software – which makes it susceptible to cybercrime. In healthcare, AI often builds off patient data – which means there are privacy concerns about its use. And incorrect deployment of or sub-optimal performance by AI tools can adversely impact a patient outcome – which opens up an organization to claims of medical malpractice.

No LHS, despite its best efforts and due diligence, can expect to remain totally immune to these types of worst-case scenarios. That’s why stratifying AI use cases by risk level is a powerful, proactive strategy that any health system can take. **Dedicated risk mitigation plans should be put into place for AI-related mishaps that impact patient data, diagnosis, or treatment.**

Governance stakeholders should identify AI use case liabilities that parallel risks found in other major healthcare initiatives for which they already have existing communication and crisis management plans. Examples include:

- **Regulatory and legal risks:** The regulatory landscape for AI in healthcare is rapidly evolving, creating uncertainty around compliance requirements and potential legal liabilities.
- **Vendor risks:** Using external AI vendors creates risks related to interoperability, vendor lock-in, data ownership, and intellectual property. Managing these risks requires careful vendor selection, contract negotiation, and contingency planning.
- **Investment risks:** AI can be costly, and ROI can be difficult to predict given the rapid evolution of technology – creating the danger of costly misfires. In making investment decisions, health systems must assess the potential for both “hard” ROI (e.g., cost savings) and “soft” ROI (e.g., improved patient outcomes).
- **Reputational risks:** High-profile AI failures or misuse could damage a health system's reputation and erode patient trust.
- **Timing risks:** Health systems face risks associated with both acting too quickly (before technology is ready) and too slowly (after you've lost ground to competitors).
- **Workforce risks:** As AI systems automate certain tasks, healthcare professionals may fear job displacement or changes to their roles.

*What AI Catalyst members are saying:*

As one CMIO put it, “We're thinking about how to stratify the risk of these different solutions. A solution to summarize meeting notes has a very different risk profile than a solution that's doing diagnostic support. We need to make sure we're focusing our attention in the right places at the right level.”

## 5. Orient your governance structures to respond to rapid evolution in the AI space

One of the largest challenges with governing AI is keeping up with its rapid evolution. Since 2017, the number of companies dealing in AI products in the US has [doubled](#). That meteoric growth stems from the fact that AI is deployed in only 25% of US companies, suggesting massive room to expand. But beyond safeguarding AI investments against technological obsolescence, effective AI governance efforts must respond – resiliently and agilely – to three growing imperatives.

1. *Keeping up to date with legal and policy challenges:* The pace of legislative changes, particularly at the state level, is ramping up considerably, with 2025 likely to be a banner year for enactment of new AI policies. Government relations officers, legal counsel, and regulatory experts should elevate healthcare specific uses of AI to focus areas when meeting with legislators. Just as boosting AI literacy and awareness among healthcare leaders is vital, translating this education into digestible talking points for legislators is a no-regret move that you can make today. As legislators look at AI regulation, helping them understand healthcare specific use-cases and governance challenges will ensure that one-size-fits-all approaches don't extinguish innovation within health systems.
2. *Addressing growing workforce mistrust of AI:* As union concerns grow over the use of AI, healthcare leaders should educate their workforce about current and proposed future use cases of AI, taking care to address the impact of AI on specific job responsibilities. Engaging union leaders directly into proposals for AI adoption (and boosting their AI literacy) can help reframe the conversation away from job losses to reducing burnout and enabling workers to operate at the top of their licenses.

3. *Ensuring effective partnership with vendors to boost model performance:* Lastly, AI governance processes should ensure that vendors play their part in keeping up to date with rapid evolution in the AI ecosystem. Inform investment decisions by asking whether your vendors are willing to collaborate long-term, not only on model performance, but on AI governance work itself. As new regulations emerge, ensure that vendors are offering support in ensuring that model use remains compliant.

*What AI Catalyst members are saying:*

As one executive noted, “We wanted to have some very strong partners in this space who would be willing to do things like explore the possibilities with us, build workflows that would solve problems we were trying to solve under particular use cases, help us with a governance structure, and do research to see not only if our use cases work, but if they really make a difference.”