

White Paper

# AI Governance in Health Systems

Aligning Innovation, Accountability, and Trust



**Valerie J Parker, MSc-GH**  
Duke-Margolis Institute for Health Policy

**Nicoleta J Economou, PhD**  
Duke Health

**Christina Silcox, PhD**  
Duke-Margolis Institute for Health Policy

## Acknowledgements

The authors would like to thank several individuals for their contributions to this white paper. First, we are deeply grateful to each member of the working group for their expert perspectives, open discussion, and thoughtful feedback. We also thank the participants of our expert workshop (see [Appendix A](#)) for sharing their expertise and experiences, as well as the multiple other health system representatives that held individual informational calls with us. We also thank Rabail Baig of Duke Health for her work reviewing and designing this paper. Any opinions expressed in this paper are solely those of the authors and do not necessarily represent the views or policies of any other person or organization external to Duke-Margolis. This work was funded by Duke AI Health.

## Working Group

We are deeply grateful to each member of the working group for their expert perspectives, open discussion, and thoughtful feedback. Any opinions expressed in this paper are solely those of the authors and do not necessarily represent the views or policies of any person or organization external to Duke-Margolis.

### **Karen Habercoss**

VP, Chief Information Security and Privacy Officer  
The University of Chicago Medicine

### **Hailey Hildahl**

Sr. Digital Product Manager  
Mayo Clinic Platform

### **Nikesh Kotecha**

Director, Head of Data Science  
Stanford Healthcare

### **Mark Lifson**

Director, AI Systems Engineering  
Mayo Clinic Center for Digital Health

### **Michael Plesh**

Executive Director of Technology  
UNC Health

### **Anurang Revri**

Vice President, Chief Enterprise Architect  
Stanford Healthcare

### **Ram Rimal**

Manager of Data Science and AI  
UNC Health

### **Lauren Rost**

Senior AI/ML Engineer  
Mayo Clinic Center for Digital Health

### **Matthew D. Solomon**

Assistant Director  
Augmented Clinical Intelligence Program  
Kaiser Permanente Northern California

### **David Vidal**

Vice Chair, AI Enablement  
Mayo Clinic Center for Digital Health

### **Ellen Woo**

Executive Director  
AI and Emerging Technologies  
Kaiser Permanente

### **Daniel Yang**

Vice President of AI and Emerging Technologies  
Kaiser Permanente

## Summary

Tools enabled by artificial intelligence (AI) have the potential to transform patient outcomes and health system operations and are already having significant effects. AI applications have facilitated faster triage and diagnosis, enabled the anticipation of patient outcomes to create personalized treatment plans, and streamlined clinical operations, patient communication, and resource allocation. But while the integration of AI tools in healthcare systems offers immense potential, the use of AI in such a sensitive and critical sector also raises significant ethical, legal, and practical concerns.

A comprehensive governance system has multiple advantages, including ensuring patient safety, maintaining ethical standards, ensuring regulatory compliance, fostering trust through transparency and accountability, and managing privacy concerns and other legal issues. But AI governance is a relatively new concept for health systems, many of which have integrated only limited numbers of AI tools into their workflows. To better understand motivations and processes, the project team convened a working group of six health systems located across the United States who have established AI governance systems in the past several years. The project team also conducted individual interviews with other health systems to understand their approaches. Although there are important commonalities in the components of governance processes, there are different ways to accomplish these tasks. At the same time, we found that this is a resource-intensive process across the board.

In the following sections, we will walk through the main components of health system governance and explore how different health systems approach these components, as well as discussing how health systems can begin to set up their own governance systems. We will also offer recommendations for policy makers, health systems, and other stakeholders on how they can standardize and simplify these processes to democratize access to AI-enabled health tools and ensure the availability of technical expertise to help under-resourced health systems realize the benefits that AI tools may provide.

## AI Governance and Strategy

AI governance is the practice of reviewing, assessing, and evaluating individual AI tools to ensure that they can be used safely, responsibly, fairly, and effectively with the health system's patient population and in compliance with applicable laws. When designing a governance system, health systems should start with a clear articulation of the principles and goals of governance for that health system and identification of stakeholders who should be involved in the governance process. The benefits of a good governance system include visibility into the AI tools being used within the health system; predictability of the information needed to review, implement, and monitor an AI tool; transparency into the governance process; standardization of the procedures for evaluation, risk assessment, and risk mitigation; clear lines of accountability; and centralized and standardized documentation on each tool's assessment and testing.

It is important to separate the concepts of AI *strategy* from AI *governance*. An AI strategy involves a systematic consideration of how to prioritize the assessment and implementation of AI tools within the health system's overall mission. What are the available resources for IT implementation and worker training? Does the health system want to prioritize certain types of tools, such as system-wide operational efficiency tools, population health management tools to assist in accountable care programs, or specific clinical decision support tools? While governance significantly benefits from such a strategy having been clearly articulated by health system executive leadership, governance systems should not be responsible for establishing or updating AI strategies. However, a health system's AI strategy should prioritize establishing an AI governance process if one does not already exist.

## Government Efforts

Federal and State actions can incentivize responsible AI governance in multiple ways, and there have already been steps in this direction. California Attorney General Rob Bonta began to [send letters](#) to hospital CEOs in his state starting in September 2022, asking them to send a list of all commercial decision-making tools in current use for clinical decision support, population health management, operational optimization, or payment management. This list was meant to include the purpose of the tool, any policies or training around the tool, and contact information for the person(s) responsible for evaluating these tools for disparate impact. Although many health systems had been thinking about how to govern AI health tools, this state action made clear the need for a centralized inventory of such software tools and a standardized evaluation system. Meanwhile, the National Institute of Standards and Measures (NIST) was developing an [AI Risk Management Framework](#) to better manage risk during development, review, and operationalization. Drafts of the Framework were published in 2022, and the final version was released in early 2023, with multiple companion tools released over the rest of the 2023 and a Generative AI Risk Profile published in summer 2024. The Office of Civil Rights at HHS also released a draft rule in 2022 and finalized the rule in May 2024 regarding [Section 1557](#) of the Affordable Care Act, “which prohibits discrimination ... in covered health programs or activities.” Part of this new rule focuses on discrimination resulting from the use of patient care decision support tools. Health systems must make a “reasonable effort” to identify and mitigate the risk of discrimination or inequitable care resulting from the use of these tools. The rule specifically notes that investigations will review whether the health system has methods to review tools it adopts or uses, and whether the tool is being used as intended. More recently, FDA Commissioner Robert Califf discussed the need for health systems to “[step up](#)” governance of AI and remarked that “they’re going to end up holding the bag on liability when these algorithms go wrong.”

## Components of a Governance System

Through our interviews and discussions, we found that governance structures can vary significantly while remaining effective. It is important for health systems to right-size their AI governance to their resources. However, there were several commonalities that facilitated building effective governance processes.

### Organizational Alignment and Engagement

AI governance bodies should have an open line of communication with leadership and decision-makers. AI governance bodies must be empowered within their organizations to assess AI tools within their purview and those decision-makers need to take recommendations seriously. This decision-making may occur at the health system leadership level (C-suite) or by individual business owners making final purchase and implementation decisions. As mentioned above,

engaging with health system leadership also allows governance to be integrated into an overall AI strategy. We also found that governance teams were often multidisciplinary to ensure that governance is a holistic process.


### Identification and Registration of AI Tools

While methods can vary, AI governance bodies must have a system in place to identify AI tools under consideration for implementation at their organization, whether these tools are commercially available or were developed internally. Once identified, AI governance teams collect information on the AI tool that will be used, potentially along with findings from the review process below, to maintain an inventory of AI tools assessed or implemented in the health system. This inventory information may also modify the review process based on the perceived risk of the tool. Inventory processes also have the benefit of standardizing the type of information required for tool assessment and setting appropriate expectations on

information requirements when comparing potential tools to submit to the governance process. Some health systems we spoke with were also performing or considering a “look-back” process for registration of tools that had already been implemented before the governance system was running.

### Review and Assessment

The main function of AI governance is to evaluate AI tools to be used in the health system. These review processes vary by organization and may include multiple domains spanning tool performance, privacy, compliance, legal, patient safety (including bias evaluations), clinical integration, IT integration, and others (Figure 1). AI governance teams may solicit information about the product from the developer directly or engage the internal champion for the tool to supply the relevant information. For some tools, validation with internal data may also be done, either retrospectively or prospectively. Typically, AI governance groups convene to discuss relevant



*Figure 1 - Potential areas of assessment in health system governance of AI tools. Health systems review and assess AI tools in multiple domains. Review processes and domains vary by organization and organizations may add or change components as their governance system process matures.*

findings before issuing a recommendation to relevant decision makers within the organization. Different organizations may limit the AI governance group to issuing recommendations on concerns or mitigation suggestions, while other may have authority to make implementation or veto decisions.

### Monitoring and Surveillance

Depending on the assessed risk of the AI tool, different types of monitoring may be required. Currently, this is often concentrated during pilot phases or immediately post-implementation. However, many people we spoke with acknowledged that ongoing monitoring is needed and are working to establish more standardized surveillance processes, as well as methods for users to communicate any concerns they may have or any perceived performance changes over time. Some systems require that every implemented tool have a prespecified person who is responsible for monitoring the use of the tool. Other systems have created a schedule for governance-led audits of tools, which allows tools that have been or become less useful than expected to be updated or decommissioned.

Although this listing of components may imply the process is fairly linear, it is often iterative in practice. For instance, it is not unusual for the AI governance committee to gather additional information from the developer throughout the process, should new questions or workflows be discovered. The new information would inform the governance review and assist committee members in making a recommendation. Another example of this iterative process could be if the AI governance body issues a recommendation against implementation of a tool due to specific risks. Should the decision-makers decide the tool is high priority, the AI governance body could be tasked with coordinating either with the developer or internal process owner to develop mitigation measures to address the identified risks to meet health system needs and meet quality and ethical standards.

These AI governance components often proceed through multiple iterative phases. For example, a tool may go through an initial assessment that results in a pilot study being recommended. Afterward, the AI governance committee may reassess the tool, applying

learnings from the pilot program, to determine a recommendation on a wider implementation. Health systems that are developing AI tools in-house will often have multiple iterative review processes to move on to the next phase of research and development. When tools developed in-house are at the point of being piloted or fully implemented, they are subject to similar governance processes as commercial tools.

## Tailoring an AI Governance Approach

Within the commonalities discussed above, each health care system we spoke with had tailored certain aspects to establish an AI governance system that best fit the needs and resources of their organization (Figure 2). One of the most interesting differences is that some systems rolled their AI governance into the existing general governance around software tools, taking more of an educational approach to ensure traditional governance entities were able to ask the right questions about AI tools, while other health systems pulled AI tools into a fully separate governance process. Still others took a hybrid approach, with some of the review components integrated into pre-existing governance processes and other components being considered separately. These approaches were not consistent based on resource availability, although health systems with fewer resources available for AI governance may be more likely to rely on additional training for existing governance system participants.

Organizations differed regarding whether a more centralized or a federated approach was a better fit for their health care system. Smaller health systems or systems with less AI expertise available generally took a more centralized approach to assessments. Very large systems' approaches were more variable. While some were centralized, others used a more federated approach to allow for more flexibility between geographic sites or regions. In some cases, the tool review and assessment were done by a more centralized team who made detailed recommendations but the ultimate decision-making was federated. In other cases, some aspects of review such as legal and

## People-Focused Variations



### Decision-making Authority

A critical piece of governance design is identifying who has the ultimate decision-making authority on whether a given AI tool will be implemented or decommissioned. This authority varied by organization. Some gave this authority to the person who allocated budget funds for the AI tool, and the review process is meant to guide this decision. Other health systems favored a more centralized decision process, where the review team or a larger governance group make the final decision. Still other health systems placed some or all decision-making with executive leadership, who rely on recommendations from the review process. This can ensure AI tool selection is consistent with the overall AI standards and strategy.



### Governance Committee Composition

Many AI governance committees are interdisciplinary, with members from disciplines such as IT, clinical care, informatics, legal, privacy, ethics, compliance, human resources, patient engagement, DEI, and finance. Some have relevant background in the AI, but others may need additional training on the implications of AI within their area of expertise. Organizations that opt to integrate AI governance into their traditional governance for other technologies also provided training on how to effectively assess AI tools. In these cases, adding a reviewer with AI technical expertise may be necessary. Some health systems had small governance teams that consisted of one to three individuals. In these cases, assessments were sometimes more focused on developer-reported performance, IT integration, privacy, and legal compliance, while local performance was assessed through qualitative pilots or monitored in post-implementation reviews.



### Including the Patient Voice

Many health systems want to include the patient voice in decisions on AI tools that may affect care. However, they reported legal and logistical challenges in allowing individuals who are not health system employees to have visibility into the full review process. Finding patients with the relevant expertise or providing adequate training on AI concepts to allow for informed involvement is also difficult. In the meantime, some health systems have brought in ethics professionals with expertise in patient opinions to help fill that perspective gap, while other health systems have consulted with pre-existing patient committees as appropriate. As patients traditionally have not been involved in technology selection and implementation, more work is needed on best practices in this space.

## Process-Focused Variations



### Governance Scope

AI is a broad term, and governance systems need to clarify the scope of tools within their purview. Some organizations focused on a range of AI tools, while others focused on machine-learning enabled tools only. Others only reviewed enterprise tools. Some took a risk-based approach to different types of AI. For example, an organization may only require registration for AI tools used for billing or business purposes but perform more in-depth reviews on AI tools that directly affect patient care. Others may have different processes for tools that have been authorized by the FDA. Governance scope was determined by several factors, including the resources available, and the scope may change as a governance system matures.



### Tool Identification

Health systems must ensure they are aware when AI tools are being considered in order to bring them into the governance process. There were a variety of strategies for this, including general informational campaigns, directed conversations with individuals involved in purchase decisions, and training with internal AI developers on how and when to engage with the governance committee. Some groups built in processes to “catch” tools within scope, often in connection with IT and procurement offices. There is not a perfect process and tools can slip through cracks at times. It can also be difficult to identify when existing tools are upgraded with AI-enabled software options and when already implemented AI tools have significant updates that may require additional governance actions.

**Figure 2 - Tailoring a Governance Approach**

privacy were centralized, but other aspects of review, such as performance and clinical integration, were federated. Still other systems centralized governance of enterprise software but federated governance of department-specific AI tools or centralized the registration process but federated the review and monitoring processes. Many larger systems, even if their processes were federated, mentioned that they were working to ensure open communication between locations to reduce or prevent repetition of work.

The project team also found interesting differences in some of the logistics of governance. However, all of the health systems emphasized that they considered their governance programs to be an evolving work in progress and anticipated there could be changes in processes and scope. For example, some governance systems started with inventories and preliminary evaluations, but are now expanding to include aspects of monitoring operational integration and performance. Other systems plan to expand the types of tools that their governance will oversee over time.

## Democratizing AI Across Health Systems

Health care in the United States already has significant equity challenges. There is concern that AI tools could worsen these inequities either because the tools may replicate and scale existing biases in care if not designed and tested carefully, or because the AI tools are effective but only highly resourced health systems can safely deploy the tools. If this happens, it will greatly diminish the impact that AI could have, especially in addressing persistent problems in healthcare such as access issues and diagnostic excellence.

We heard that government and health system leadership cannot continue to rely on volunteer efforts for sustainable governance of AI. Health systems are concerned that they lack the resources to bring in staff or train existing staff, or to build the infrastructure needed for effective and ethical governance. There is a critical need for ways to scale and propagate internal AI expertise as well as templates and best practices for

governance processes as health systems begin to deploy these tools.

In the sections below, we will walk through some recommendations on how different stakeholders can help democratize safe and efficient implementation of AI tools through effective governance.

### Government Actions

Federal and state governments have multiple options to incentivize and to support effective governance of AI tools in health. The OCR Section 1557 rule described earlier is one such example. The Office of the National Coordinator of Health Information Technology (ONC) also released a final rule on [transparency requirements](#) for certain types of predictive decision support intervention tools. This latter rule ensures that health systems and users will have information about those specific tools but also sets a baseline standard for what information health systems should know about before implementing these types of care tools. Similarly, the FDA recently put out [guiding principles](#) around transparency for AI/ML devices that have significant overlap with ONC requirements.

At the state level, Colorado recently [passed a law](#) that will require deployers of AI systems where outputs are a “substantial factor” in decisions regarding the provision, denial, cost, or terms of health care services to implement a risk-management system, conduct impact assessments, do annual reviews, and report any discoveries of algorithmic discrimination. Although the specifics of these laws and rules can be debated, there are clear options for government bodies to incentivize good governance. The government can also create positive incentives around governance such as safe harbors for health system deployers that employ best practices to reduce some of the risk in deploying AI tools, such as liability. For example, Colorado built an affirmative defense into their law for deployers of AI tools that could show that they had complied with specific national or international AI risk management frameworks.

The government can also prioritize funding research to simplify governance and make it more efficient. This may include creating research funding priorities around



governance best practices, maturity models, and infrastructure to make monitoring for performance draft and bias and general surveillance more efficient. Government could also fund development of open-source tools such as inventory systems and testing tools to make the governance process less burdensome. A recent [journal perspective](#) suggested that the government could help build a registry of AI tools similar to ClinicalTrials.gov that would also have a federated component linking to health system assessments of that tool. A system like this could also provide HHS a central location for users and patients to report safety concerns about AI tools, which is a task required in the 2023 Executive Order on AI. Finally, the government could consider establishing and funding Health AI Technical Centers of Excellence to provide training modules for staffing governance teams and act as an expert resource for under-resourced health systems, as well as general workforce development around AI literacy.

### Developer Actions

Developers are a significant source of AI expertise. Although they cannot be considered impartial, they do have an interest in increasing trust in AI tools and ensuring that they are being implemented and used correctly. The project team spoke with several commercial developers of tools and heard concerns that many health systems are not asking enough questions about their products. Developers and health systems should be working together to create standardized checklists of information for different types of AI tools, to set appropriate expectations and increase transparency. This would also allow developers to create a standard information disclosure form that could be shared with governance teams that could reduce the amount of back-and-forth communication between developers and governance teams, increasing governance efficiencies. Health systems that the project team spoke with frequently mentioned that transparency around health system data movement and how that data is used is especially important to them.

Developers should also work to foster collaboration and trust with health systems. Aligning on expectations

early and improving understanding of health system legal compliance requirements would be helpful. One example of this involved product updates. Health systems felt there were often significant mismatches on what constitutes a “substantial change” that would require more active alerts to health systems to allow for governance review. This was often around data security and privacy, but also when a product that previously did not use AI/ML in its software was updated to include AI/ML components.

Developers can also create tools to facilitate local governance. One company recently announced that they would provide kits to simplify local tuning and testing of their products. This aligns with the previously mentioned FDA [transparency principles](#) stating that it would be helpful for developers to provide information on “how to conduct local site-specific acceptance testing or validation” and “plans for ongoing performance monitoring.” Another company we spoke with described tools that would be able to automate monitoring for performance drift.

### Health System Actions

Health system leaders should prioritize AI governance now and seek learning from early adopters, assessing what is the right-sized approach for their specific circumstances. However, health systems that have already built governance systems or have significant expertise in AI (such as academic health systems) also have a role in democratizing AI across settings. These systems should [share documentation](#) on how their governance systems work, [including tools](#) such as registration/information intake forms and surveillance procedures, and consider [partnering](#) with other health systems, especially those with fewer resources. At the same time, all health systems need to work to diffuse knowledge about AI and responsible AI implementation throughout their workforce.

To accomplish both tasks, health systems can create peer-to-peer learning spaces to educate, share and support each other implementing best practices in AI governance. For example, the ECHO Institute New Mexico Hub is starting a [community of practice](#) for providers around implementing AI in medicine. A recent pilot project called the [Practice Network](#) also

just launched, where participants will receive “access to one-on-one guidance, expert consultation, peer learning community, educational materials, and other resources” from more experienced health systems partners to increase equitable use of AI in healthcare. Health systems with less AI expertise have found that working with more local or smaller AI developers can be helpful in building a more trusting relationship and allow the developer to better understand and then design toward the needs of a specific health system.

### Other Stakeholder Actions

Other stakeholders also have important roles. Entities such as clinical societies, public-private partnerships, and standards groups should focus on creating guidance in this space. Examples include the National Academies’ [AI Code of Conduct](#) and the Coalition for Health AI’s [Assurance Guides](#). We also expect that third parties will enter this space to help provide [assurance reviews](#) and provide contracted governance services or commercial governance software tools. One early example is Dandelion Health, which provides a [free validation service](#) for certain types of health AI tools through a grant from the Gordon and Betty Moore Foundation. Outsourcing some of these tasks to reusable validation platforms could create efficiencies of scale and reduce overall costs.

Medical, nursing, and other clinical professional schools and training programs also should develop curricula on best practices in assessment and using AI tools while clinical societies and other organizations should establish [continuing education courses](#) on responsible governance and use of AI.

**About Duke Health.** Duke Health is committed to advancing health and transforming lives through clinical care, medical education, and innovative research. Duke Health’s comprehensive network of hospitals, outpatient clinics, and specialty centers, serving diverse populations across NC and beyond. Its mission emphasizes training the next generation of healthcare leaders and advancing cutting-edge research in areas such as precision medicine, AI-driven healthcare, and population health. Duke Health is a committed to equity, innovation, and continuous improvement, ensuring that care is not only safe and effective but also equitable and responsive to the needs of all patients. Through collaborative efforts across its academic and clinical arms, Duke Health remains at the forefront of addressing the most complex healthcare challenges, fostering a healthier future for individuals and communities worldwide.

## Conclusion

AI tools present an extraordinary opportunity to transform health care, but establishing a robust AI governance framework is essential to ensure that these tools are deployed safely, ethically, and in compliance with regulatory standards. Governance systems not only protect patient safety and foster trust; they also facilitate innovation by providing clear guidelines and processes for assessing and implementing AI technologies. The diverse strategies employed by different health systems highlight the continued need for flexibility in governance approaches, factoring in health systems’ specific considerations around resources and processes. However, the number of commonalities found when exploring the different governance processes suggests that health systems should make use of published frameworks and guidance as they create their own processes.

Work is also needed to ensure that the safe and effective use of AI tools can be democratized across all health systems. Widespread implementation of AI governance in healthcare hinges on addressing key challenges such as funding, staffing, and training. This will involve targeted actions among government, developers, health systems and other stakeholders.

Ensuring that AI tools can be safely used in all health systems will be a challenging and ongoing task but must be accomplished if health AI is to fulfill its potential to improve health outcomes, reduce costs, enhance the clinical experience for both patients and providers, and advance health equity.

**About Duke-Margolis.** The Robert J. Margolis, MD, Institute for Health Policy at Duke University is directed by Mark McClellan, MD, PhD, and brings together expertise from the Washington, DC, policy community, Duke University, and Duke Health to address the most pressing issues in health policy. The mission of Duke-Margolis is to improve health, health equity and the value of health care through practical, innovative, and evidence-based policy solutions. Duke-Margolis catalyzes Duke University’s leading capabilities, including interdisciplinary academic research and capacity for education and engagement, to inform policymaking and implementation for better health and health care. For more information, visit [healthpolicy.duke.edu](http://healthpolicy.duke.edu).

# Appendix A: Expert Workshop Participant List

## Health System Governance of AI Tools

Virtual Expert Workshop | June 6, 2024

**Laura Adams**

National Academy of Medicine

**Brian Anderson**

Coalition for Health AI

**Allie DeLonay**

SAS

**Nicoleta J Economou**

Duke Health

**Tom Ferrone**

Tempus AI

**James Gaston**

Parkland Health & Hospital System

**Mallory Gibreal**

Bryan Health

**Karen Habercross**

University of Chicago

**Marianne Hamilton Lopez**

Duke-Margolis Institute for Health Policy

**James Leo**

MemorialCare Health System

**Brett Moran**

Parkland Health & Hospital System

**Zachary Lipton**

Abridge

**Valerie Parker**

Duke-Margolis Institute for Health Policy

**Anurang Revri**

Stanford Healthcare

**Ram Rimal**

University of North Carolina Chapel Hill

**Brian Scarpelli**

Connected Health Initiative

**Christina Silcox**

Duke-Margolis Institute for Health Policy

**Jennifer Stoll**

OCHIN, Inc.

**Christine Swisher**

Oracle Health

**Sylvia Trujillo**

OCHIN, Inc.

**David Vidal**

Mayo Clinic

**Celena Wheeler**

Oracle Health

**Ellen Woo**

Kaiser Permanente

**Daniel Yang**

Kaiser Permanente