

AI in Healthcare:

Navigating the Future with Governance and Innovation

Primary Author:

Ben Franjesevic, Director

Contributions from:

Stephanie Gonzales, Director

Steve Smith, Senior Manager

Tarun Dhawan, Associate Director

Allen Warner, Associate Director

The role of artificial intelligence (AI), which includes both machine learning (ML) and generative AI (GenAI), is becoming increasingly prominent in the business world. AI is challenging traditional models of business operations with its potential to radically transform processes, including those within finance and accounting, customer service, marketing and cybersecurity.

While the healthcare industry's widespread adoption of artificial intelligence (AI) is still in an early stage, many organizations are becoming increasingly aware of the potential benefits AI can offer. It is widely believed that AI will exponentially improve the healthcare industry's utilization of data to, for example, streamline how claims are handled, enhance the patient experience, and improve the detection and diagnosis of medical conditions.

However, the AI revolution has driven the realization that foundational governance and risk management processes need to be implemented to reap these benefits *responsibly*. Healthcare organizations must go beyond thinking only of the operational efficiencies or cost savings that AI can bring – they must place an emphasis on the potential implications to patient safety, ethical dilemmas, unintentional bias, data privacy, and more.

As healthcare organizations continue to further explore the use of AI, internal audit departments have a responsibility to help ensure their organizations implement a strong AI governance foundation. The following is a collection of AI governance initiatives and principles that Protiviti's experts have compiled through interactions with organizations, both inside of the healthcare industry and beyond, that should help equip organizations with the necessary knowledge to support them on their AI journeys.

Starting the journey: AI readiness

The readiness of an organization to implement or expand the adoption of AI, as well as a program to govern it, depends on the organization's specific needs. While specifics vary from organization to organization, several common factors should be considered:

- **Capabilities and resources:** The availability of people and financial/technological resources impacts the implementation of AI and corresponding governance models. While organizations likely have some existing capabilities that can be leveraged, there may be a need to invest in hiring new or training existing personnel, as well as enhancing the technology environment in support of AI adoption.
- **Existing policies and ethics guidelines:** Most organizations already have policies in place governing the implementation and use of technology, as well as providing ethical guidance in support of business objectives. While these policies and guidance may not be specific to AI, they provide the context within which AI should be implemented and used, and can serve as a foundation on which AI governance can be built.
- **Data governance and management:** AI relies on data during both the development and operation of a solution. As such, data governance and management strategies, policies and practices are important to consider when implementing AI and its governance structure. This includes how the organization sources data, manages its quality and protects it.
- **Risk management strategies:** The concept of risk is not unique to AI, and many organizations have already established strategies to manage risk more broadly. It is important to consider these strategies during the implementation and use of AI as well as when establishing AI-specific governance.
- **Security and privacy:** The security and privacy of data is an important consideration when implementing AI technologies. This involves not only securing the AI systems themselves but also managing the data lifecycle in a way that prioritizes security, confidentiality and integrity.

These factors can influence how AI is adopted and used and can inform how corresponding governance is designed and implemented. Given its positioning in the organization, internal audit should be involved in discussions about establishing governance around AI and can provide insight relative to these factors through a risk-and-control lens.

This concept of AI governance is key in the journey for responsible and ethical AI adoption, so it is important to explore exactly what is meant by AI governance.

Navigating innovation and control with AI governance

AI governance acts as a compass, orienting organizations throughout their AI journeys. It also provides a framework; establishes rules for using AI responsibly; and facilitates the ethical use of data and algorithms, compliance with regulations, alignment with business goals, and more. Governance includes the organizational structures, policies, procedures and controls that internally regulate the acquisition, development, deployment and use of AI systems.



Most healthcare organizations are familiar with the concept of governance outside of AI. Many have already established policies for technology that can serve as a guide and set the foundation for effective

AI governance. The lessons learned from establishing governance for other business aspects and technologies can act as an accelerator for developing an efficient AI governance approach and framework.

Getting started with AI governance includes establishing an advisory board, aligning governance with existing policies, educating users, inventorying use cases and solutions, assessing risk and implementing controls.

Establishing an AI governance advisory board

One of the cornerstones of effective governance is a well-functioning oversight board. In addition to chartering an AI governance body with a clear mandate and vesting it with authority sufficient to fulfill its goal, senior management should appoint engaged leaders across a wide range of applicable disciplines to serve on it. The board should be comprised of key stakeholders, including those who have knowledge and expertise in key areas such as technology, internal audit, law, compliance and cybersecurity, as well as clinical and business representatives which may include the chief medical officer, chief nursing officer, pharmacy director, and representatives from areas such as clinical informatics, health information management and patient services.

An advisory board, sometimes called a steering committee, is primarily responsible for providing strategic direction and oversight, ensuring the organization's adoption of AI aligns with the organization's broader mission and objectives. The advisory board defines the vision and goals for deploying AI in the organization, prioritizes opportunities and risks, and facilitates adherence to ethical guidelines and regulatory compliance. It should ensure that aspects of the AI governance framework align with legal requirements, recognized frameworks such as the National Institute of Standards and Technology's (NIST) [AI Risk Management Framework](#), and the organization's ethical principles such as transparency, fairness and security.

Aligning AI governance with existing policies

As mentioned previously, most organizations operate within a framework of established policies and standards. Given this, it's important that the implementation of AI governance is not viewed as an isolated activity, but rather an evolution of the governance that's already in place and which builds on existing policies and standards to address new dimensions introduced by AI. For example:

- **Ethical guidelines:** Organizations often have already established core ethical guidelines and principles that can and should be extended to align with the characteristics of responsible AI, such as transparency, fairness, accountability and explainability.
- **Acceptable use:** Policies are often established to provide guidance on how personnel should treat data or use technologies. Many of these concepts relate closely to how AI should be used.
- **Third parties:** Organizations typically have defined policies for procurement and for evaluating third-party risk that include performing due diligence and understanding third-party controls around cybersecurity and data that can be expanded to also consider AI.

- **System development:** Many organizations have project management or system development lifecycle methodologies that define how technology solutions are evaluated and built, which can be extended to AI.
- **Data/information governance and privacy:** Organizations also often have policies around how data and information is classified, handled and used. These policies can also address data needs specific to AI while considering requirements more broadly (e.g., handling of protected health information, or PHI).


It's important to recognize that existing policies may not cover all aspects of AI adoption, necessitating new, AI-specific documentation. For example, formalizing the roles and responsibilities of the AI advisory board and introducing new monitoring requirements or net-new compliance requirements may require dedicated policies. As internal auditors consider how AI governance is being implemented, they should consider the organization's overall governance framework.

Educating users

As new policies are introduced or existing ones are updated to incorporate AI, educating users becomes key. This helps ensure that users not only understand what AI is but also helps them appreciate its potential benefits, limitations and risks. Users must also understand the organization's stance on how AI should or should not be used so they can interact with AI systems responsibly and safely.

Training programs should help foster awareness about the risks associated with AI systems. This will help users to identify risks such as data privacy breaches or biased decision making. Driving awareness of risks while emphasizing ethical considerations around transparency, fairness and security helps foster a culture of responsibility among users of AI systems.

As internal auditors evaluate the training that is provided to users, they should understand the training content and determine whether key risks are addressed and if the content caters to a broad audience, from clinical practitioners using AI systems in patient care to administrative personnel using them for



operational efficiency. Further, internal auditors should understand the training audience to assess whether the appropriate personnel across clinical and administrative operations are included.

Inventorying AI use cases and solutions

Without visibility into existing AI technologies and use cases, it is difficult to manage the corresponding risks. To drive visibility, organizations must identify existing technologies and develop an inventory that includes internally developed solutions as well as third-party applications that include AI capabilities. Often, organizations can leverage an existing application inventory or configuration management database to capture information about AI solutions. It's important to note that the identification of existing technologies is not a one-time activity; it requires periodic reevaluation to identify solutions where new AI capabilities are introduced.

Besides inventorying existing AI technologies, organizations should also identify and catalog new AI use cases. The advisory board plays a key role in identifying and evaluating these use cases, ensuring each one has customized controls based on its unique needs, benefits and risks. Governance components and controls should be flexible and easy to adapt to the specific AI use case, acknowledging that different use cases, such as billing and coding versus clinical decision support, will require distinct controls and considerations.

When internal auditors begin to evaluate the processes and controls that management has put into place around the organization's use of AI, some of the first things that should be reviewed include the inventory of AI solutions as well as the backlog of AI use cases. It is important for internal audit to understand the data (including risk information) that has been collected to inform the audit approach.

Assessing risk and implementing controls

Organizations should continuously assess risks associated with all AI solutions, whether existing or proposed, and should design and implement corresponding controls. This includes both broader, enterprise-level risks as well as risks specific to individual use cases. Examples of relevant risks that should be considered as part of the AI governance framework, along with corresponding control activities, include the following:

1. **Data security** involves protecting electronic data from unauthorized access, while **data privacy** involves protecting the confidentiality of personal information. Risks include exposure of sensitive patient data, reduced data integrity, and AI systems inferring sensitive details. Controls include firewalls, data anonymization and security over model and training data.
2. **Bias** refers to systematic errors resulting in unfair outcomes. Bias can be data driven (e.g., from societal biases in historical data), algorithmic (e.g., a model favors patterns or characteristics), or driven by how users interact with AI (e.g., introduced through the user interface). Risks include incorrect patient care plans, violation of anti-discrimination laws and healthcare regulations, and reduced model performance. Controls involve methodologies to prevent bias, algorithms to detect and combat bias, and transparency in AI decisions.
3. **Explainability**, or the "black box problem," relates to the frequently opaque nature of AI models, which can make it difficult to understand the rationale for their outputs and decisions. Risks include clinical errors from reliance on unexplainable outputs and regulatory issues due to nontransparent decisions. Mitigation tactics include employing explainable AI (XAI) for decision insights, providing clear descriptions of model components, and maintaining comprehensive documentation on the model design, criteria, limitations, validations and performance.
4. **Data and model integrity** refers to the accuracy, reliability and consistency of data and AI models. Risks involve inaccurate data influencing a model, unexpected outcomes due to changes to model inputs or assumptions, and irreproducible results. Controls include quality checks on data, adjusting models for changing patterns, and human oversight and intervention.

5. **Clinical safety** requires that AI technologies don't endanger patients. Risks involve models overlooking factors that a human wouldn't, incorrect diagnoses, patient harm or suboptimal care, and incorrect patient prioritization in triage. Controls include thorough pre-deployment testing and validation, regular recalibration based on new clinical evidence, and creation of triage-specific performance metrics.

As an organization establishes and enhances its AI governance framework, it's important that the framework is designed to consider risks relative to the enterprise as well as to specific uses of AI. Internal auditors should understand how the organization is identifying, assessing and mitigating AI risk both broadly as well as on a use-case by use-case basis. This includes how controls are applied to a use case based on the unique risks associated with it.

Regulatory compliance and AI governance

As one of the key drivers of AI governance is compliance, internal audit should consider the organization's AI governance framework within the context of the healthcare industry's complex and ever-evolving legal and regulatory landscape. AI technologies introduce new challenges from a compliance perspective. This includes diagnostic and treatment errors, Health Insurance Portability and Accountability Act (HIPAA) breaches, and discrimination or civil rights violations. Compliance departments should be proactive by gaining understanding of the business's current and planned uses of AI – and their associated risks and how to mitigate them – to help ensure compliance with both existing and emerging laws and regulations.

It is important to acknowledge that there is no universal regulatory framework specifically designed for AI applications in the United States, though some federal departments have published guidance on AI, such as the U.S. Department of Justice's updates to its *Evaluation of Corporate Compliance Programs* policy, and some states have begun passing legislation targeting AI, such as Colorado's Consumer Protections for Artificial Intelligence Act. However, healthcare organizations are subject to legal and regulatory considerations covering data more broadly, as well as healthcare-specific laws and regulations that impact how an organization uses AI. Some examples that should be considered include:



- The U.S. Food and Drug Administration (FDA) regulates medical devices under which some AI-powered solutions fall. The FDA published material guiding the appropriate use of AI for medical devices as early as 2019. As recently as March 2024, the agency released additional [guidance on AI regulation for biologics, drugs, devices and combination products](#).
- New industry-specific guidance such as the American Medical Association's (AMA) "Principles for Augmented Intelligence Development, Deployment, and Use," released in November 2023, will need to be considered for AI systems being deployed and developed.

- AI often relies on large datasets for training and operation, and these datasets may include PHI or other sensitive data. Regulations, such as HIPAA's Final Privacy and Security Rules, regulate how patient information is protected. Moreover, individual states have begun to enforce their own set of regulatory requirements when it comes to consumer data, including PHI. For example, the [California Privacy Rights Act](#) (CPRA) introduces protections for sensitive personal information. Currently, 15 states have enacted their own consumer data privacy laws.
- As internal auditors evaluate the AI governance frameworks that are being implemented, it's important to consider how regulatory compliance has been incorporated as a foundational component of the framework. This includes considering how:
 - Relevant legal and regulatory requirements at the national, state and local levels are identified, inventoried and monitored, as well as how changes to these requirements are identified and communicated to stakeholders.
 - Policies are developed or enhanced to align with legal/regulatory requirements, considering data privacy, security and ethical considerations for AI use and patient rights.
 - The organization's personnel are trained on expectations related to responsible AI use as established both by policies as well as laws and regulations.
 - AI activities are monitored for adherence to requirements to verify responsible, ethical and compliant use.

The benefits and importance of AI governance

Realizing AI's opportunities, measuring and managing both benefits and risks, becomes difficult without appropriate and effective guardrails in place. Such guardrails are important — not to slow down AI adoption, but rather to accelerate and optimize its responsible and effective use throughout the enterprise while allowing for new opportunities to emerge.

By implementing effective governance frameworks and policies, organizations can efficiently recognize, comprehend and measure the potential risks and opportunities associated with AI. This enables them to align their governance practices with the use cases they have identified and make well-informed decisions regarding the utilization of AI.

The benefits provided by a strong governance program should be communicated to stakeholders to help drive understanding. Some of these benefits include:

- **Maximizing value and impact:** Effective AI governance helps align AI adoption with organizational strategy, maximizing business value. It directs investments toward beneficial projects, driving innovation and sharing benefits among stakeholders.
- **Risk management:** AI systems can introduce risks like biases, privacy concerns and misuse. Effective governance frameworks can help to identify, assess and mitigate these risks, driving

safety and fairness. This approach helps build trust among stakeholders by demonstrating a commitment to ethical and responsible AI use.

- **Regulatory compliance:** AI governance can help organizations remain compliant with existing laws and prepare for new regulations. Building in compliance considerations as part of the AI governance program helps avoid legal penalties and build trust related to the use and security of patient data.
- **Ethical AI use and building trust:** Clear governance frameworks and structures can foster trust by making AI processes understandable and fair. Aligning AI systems with ethical standards is key for maintaining public trust and delivering ethical patient care.
- **Sustainable innovation:** Effective AI governance enables responsible development and deployment of AI technologies, balancing rapid advancement with ethical considerations and long-term viability. This allows healthcare organizations to optimize AI use while maintaining public trust and creating new opportunities.



How internal audit can use AI

While the focus of this paper has been on how internal audit can approach a healthcare organization's implementation of AI and associated governance, it's also important to acknowledge that this technology can also revolutionize the internal audit department itself. Some examples of how AI can be used in internal audit include:

- **Risk assessment and audit planning:** AI solutions can help evaluate risk more broadly. In planning for an audit, AI can be used to help identify potential areas of risk and corresponding control activities to be evaluated for a given audit area. AI can also be used for developing evidence requests and testing procedures based on the audit scope. In an ongoing capacity, AI may be leveraged to further enhance true dynamic risk assessment practices such as monitoring in-the-news data sources about the organization (e.g., physicians, in-network payers, key suppliers, key payers, industry peers) to identify areas of potential risk to inform the audit plan.

- **Audit execution:** AI can be embedded throughout the execution of an audit to help prepare questions for meetings, capture meeting notes, develop analytic audit procedures or extract information from unstructured data in support of testing.
- **Audit reporting:** AI can be used in support of reporting the results of an audit, from the drafting of initial issue and observation language to summarizing findings to be presented to senior leadership and the audit committee.
- **Continuous monitoring:** AI can be incorporated into a continuous monitoring program to identify risks and control gaps in areas such as:
 - *Payer claims management:* Identifying and flagging suspicious claims patterns early to prevent fraudulent payouts
 - *Provider revenue cycle:* Identifying negative payer trends and potentially predicting denials and/or underpayments and compliance coding issues before they occur or before they are billed
 - *Provider prescribing:* Utilizing historical data to predict abusive prescribing of controlled substances or conflicts of interest related to prescribing
 - *Provider pharmacy monitoring:* Leveraging Natural Language Processing (NLP) to categorize free-text blind count discrepancy reasons and trends across the population

Internal audit's ability to perform dynamic risk assessments across the organization is critical in healthcare, and nothing is more suited to facilitate the dynamic assessment process than AI. The opportunity to enhance internal audit's ability to monitor and evaluate an organization's risk exposure is here, and the possibilities are endless.

As healthcare organizations continue their AI journeys, it's clear that this technology holds unprecedented potential for transforming healthcare delivery. From improving patient experiences to streamlining operational processes, AI is poised to play an integral role in shaping the future of healthcare.

However, as with any powerful technology, the adoption of AI needs to be managed responsibly. By aligning governance practices with specific use cases, organizations can effectively navigate potential risks while maximizing the benefits associated with AI. Moreover, by adopting a proactive approach toward ethical considerations and risk management, healthcare organizations can drive sustainable innovation that respects both regulatory and industry standards as well as the public's trust.

For internal auditors in particular, AI presents an opportunity not just to enhance audit efficiency but also to provide strategic value within the organization. Internal audit is uniquely positioned in the organization and can play an important role in providing risk and control guidance as the organization undertakes its AI journey. In doing so, it can assist the organization by helping maximize AI's transformative potential, drive operational efficiency, and improve health service quality and consistency.

About AHIA

The Association of Healthcare Internal Auditors (AHIA) is a network of experienced healthcare internal auditing professionals who come together to share tools, knowledge, and insight on how to assess and evaluate risk within a complex and dynamic healthcare environment. AHIA is an advocate for the profession, continuing to elevate and champion the strategic importance of healthcare internal auditors with executive management and the Board. If you have a stake in healthcare governance, risk management and internal controls, AHIA is your one-stop resource. Explore our website for more information. If you are not a member, please join our network, www.ahia.org. AHIA white papers provide healthcare internal audit practitioners with non-mandatory professional guidance on important topics. By providing healthcare specific information and education, white papers can help practitioners evaluate risks, develop priorities, and design audit approaches. It is meant to help readers understand an issue, solve a problem, or make a decision. AHIA welcomes papers aimed at beginner to expert level practitioners. This includes original content clearly related to healthcare internal auditing that does not promote commercial products or services. **Interested? Contact a member of the AHIA White Paper Subcommittee.**

Subcommittee:

Alan Henton, White Paper Chair
Vanderbilt University Medical Center
alan.p.henton@vumc.org

Valerie Mattas
Sharp Healthcare
valerie.mattas@sharp.com

Debi Weatherford
Piedmont Healthcare
debi.weatherford@piedmont.org

Laura L. Sak-Castellano
Advocate Aurora Health
Laura.Sak-Castellano@aah.org

Megan DeVries, AHIA Board Liaison
CoreWell Health
megan.devries@corewellhealth.org