ISSN: 2473-4901
v12 n6830

AI In Healthcare Practice And Information Systems: Challenges Faced And A Model Framework For The Legal And Ethical Management of AI In Healthcare Organizations

Mark Pisano
pisanom1@southernct.edu
Southern Connecticut State University
New Haven, CT 06516 USA

Kauther S. Badr
badrk1@southernct.edu
Southern Connecticut State University
New Haven, CT 06516 USA

Robert A. Smith
smithjrr1@southernct.edu
Southern Connecticut State University
New Haven, CT 06516 USA

Abstract

As the practice and business of healthcare continue to be disrupted by artificial intelligence (AI), concerns about protecting sensitive patient data continue to emerge. AI-enabled clinical and administrative applications increasingly rely on electronic health records and enterprise systems to support clinical services, decision-making, analytics, and operational efficiency. As such, governance and ethical challenges surrounding Protected Health Information (PHI) and electronic Protected Health Information (ePHI) have grown more complex. This paper examines the intersection of AI technologies, healthcare information systems, related regulations, and ethics, with emphasis on the Health Insurance Portability and Accountability Act (HIPAA). It also explores how AI systems within health organizations process, collect, and share PHI, which, we argue, calls on healthcare organizations to rethink and update compliance frameworks to address the nuanced vulnerabilities and compliance requirements arising from their use of AI. Key challenges discussed in this paper include the potential for data re-identification, secondary use of health information, algorithmic bias, and reliance on external vendors. The implications of emerging consumer privacy laws and governance gaps for health-related data generated outside clinical settings are also examined. Recommendations on policy, technical, and organizational approaches to ensure AI systems meet ethical and privacy standards are discussed.

Keywords: Healthcare, Artificial Intelligence, Federal Regulation, Compliance, HIPAA, HITECH

AI In Healthcare Practice And Information Systems: Challenges Faced And A Model Framework For The Legal And Ethical Management of AI In Healthcare Organizations

Mark Pisano, Kauther S. Badr and Robert A. Smith

1. INTRODUCTION

The rapid integration of artificial intelligence (AI) into healthcare delivery and healthcare business information systems in the United States continues to transform clinical decision-making, diagnostic processes, and operational management across the industry. Over the past five years, there has been an exponential increase in the number of healthcare organizations adopting or considering AI technologies, making this paper relevant and timely (Nguyen et al., 2025). Artificial intelligence has disrupted the healthcare industry by transforming how healthcare organizations provide services. Innovation in AI has increased its use in clinical practice, thereby altering traditional care delivery models, and it has also allowed healthcare organizations to realize process and administrative efficiencies. The capacity for such systems to "self-learn" from incoming data and adapt to changing clinical environments brings both promise and complexity to healthcare delivery (Magrabi et al., 2019). These advancements are set to revolutionize not only diagnosis but also operational processes across numerous clinical domains—including radiation oncology, robotic surgery, and organ allocation (Murdoch, 2021). AI systems have moved beyond simple recommendations to perform autonomous tasks such as patient triage and screening referrals. However, these innovations simultaneously create ongoing risks and vulnerabilities for healthcare organizations related to legal compliance and ethics.

This paper articulates and details use cases for AI in healthcare—clinical decision support, diagnostic imaging, and operations, and risks—data privacy and security breaches, algorithmic bias, employee overreliance, and reliance on external vendors, commonly associated with implementing AI in a healthcare organization so that they can be mapped to regulatory compliance mandates providing a foundation for a model implementable artificial intelligence compliance framework for healthcare organizations.

2. COMMON USE CASES OF ARTIFICIAL INTELLIGENCE IN THE HEALTHCARE INDUSTRY

The past decade has witnessed an exponential rise in the application of AI across various facets of healthcare. AI-driven systems have expanded beyond simple diagnostic tools into comprehensive clinical decision support systems (CDSS), predictive models, and even autonomous systems that facilitate tasks such as patient triage and scheduling. The enhanced capabilities of AI tools are evident in several areas, including clinical decision support, diagnostic imaging, and autonomous systems, across diverse healthcare domains. These advancements have led to improved patient outcomes, reduced operational costs, and accelerated research, marking a transformative era in healthcare.

2.1 AI In Clinical Decision Support

Clinical Decision Support Systems (CDSS) have evolved from early rule-based expert systems to sophisticated machine learning models that guide diagnosis and treatment decisions (Ogut, 2025). CDSS now play a crucial role in augmenting healthcare providers' abilities by incorporating patient-specific data, clinical knowledge, and other health information into the decision-making process (Ogut, 2025; Sutton, et al., 2020). These systems assist in various functions, including reducing medication errors, improving patient safety, providing clinical management, and enhancing care coordination—all vital aspects when addressing complex patient cases (Sutton, et al., 2020). AI-driven CDSS can analyze combinations of symptoms to diagnose rare diseases or to personalize treatment suggestions based on big data from similar patients (Abbas et al., 2025; Ogut, 2025). They also show promise in predicting adverse events and identifying occult conditions from tests, suggesting that AI can augment human decision-making (Ogut, 2025). Despite these advancements, challenges remain in ensuring data quality,

addressing algorithm transparency, and integrating CDSS seamlessly into clinical workflows to maximize their utility without overburdening healthcare providers.

2.2 AI Diagnostic Imaging

The integration of AI into diagnostic imaging is among the most dynamic areas of healthcare innovation. AI algorithms can rapidly interpret diagnostic images, such as chest X-rays, to identify a multitude of pathologies within mere seconds (Arzikulov & Komiljonov, 2025; Murdoch, 2021). AI-powered systems enhance the accuracy, speed, and consistency of image interpretation, transforming radiological practice into a data-driven and efficient discipline (Arzikulov & Komiljonov, 2025). Deep learning models, particularly Convolutional Neural Networks (CNNs), can detect early-stage pathologies and differentiate complex tissue structures with precision comparable to that of experienced radiologists (Arzikulov & Komiljonov, 2025). For instance, the Food and Drug Administration (FDA) has approved machine-learning software to detect diabetic retinopathy from diagnostic imaging (Murdoch, 2021; Saeidnia et al., 2025)—a breakthrough that underscores the potential for faster and more accurate diagnostics in ophthalmology (Murdoch, 2021). AI algorithms can quickly analyze vast amounts of imaging data, identify abnormalities, and reduce radiologists' workload, leading to earlier and more accurate diagnoses (Miyoshi, 2025).

2.3 Operations

AI is making inroads in a variety of other healthcare domains. These include operational management in healthcare facilities, optimizing scheduling, improving resource allocation, and even supporting research in precision medicine. The cumulative effect is a healthcare ecosystem in which AI serves multiple roles, thereby increasing reliance on digital data flows and further emphasizing the need for robust data governance. The rapid proliferation of AI in healthcare is not limited to clinical applications alone. Business information systems within healthcare organizations increasingly rely on AI for data analytics, performance monitoring, and operational decision-making. This intersection of clinical and business intelligence creates an ecosystem in which patient care and organizational effectiveness are closely intertwined (Karunanayake, 2025).

3. COMMON RISKS ASSOCIATED WITH THE USE OF AI IN THE HEALTHCARE INDUSTRY

As AI systems become more embedded in healthcare infrastructure, understanding the common risks and challenges flowing from the adoption and deployment of AI becomes critical. This section discusses the most common risks and challenges a healthcare organization is likely susceptible to when adopting and deploying AI-powered tools or processes. The integration of AI in healthcare not only highlights the need for data protection but also exposes specific challenges that must be addressed to ensure equitable and secure practices. These challenges include reliance on external vendors and the implications of evolving consumer privacy laws.

3.1 Data Privacy and Security Breaches

According to IBM’s Annual Cost of a Data Breach Report, healthcare organizations experience the highest costs related to data breaches, and the data that healthcare organizations collect remains highly valuable for the criminally inclined. (IBM, 2025). Transmitting sensitive patient data to external servers, even in de-identified or anonymized formats, increases the risk of data breaches and unauthorized access. Closed AI models are "black boxes," making it difficult to ascertain how data is processed or stored. As such, ensuring compliance with regulations like HIPAA or GDPR becomes more challenging when data is outside an institution's direct control (Dennstädt et al., 2025). AI and machine learning can, however, be leveraged to enhance breach detection systems, optimize data encryption algorithms, and automate compliance monitoring (Conduah et al., 2025). Beyond security and privacy concerns, implementing AI, especially open LLMs, requires high-performance computational hardware and specialized expertise for customization and maintenance. This can be a significant barrier for healthcare facilities with limited resources (Dennstädt et al., 2025). Relying on external vendors for closed AI solutions can limit healthcare facilities' control over their data and AI infrastructure, leading to challenges related to data sovereignty and potential "vendor lock-in" if policies, pricing, or service availability change (Dennstädt et al., 2025). Given these inherent and potentially significant risks,

healthcare organizations must include regulatory compliance, transparency, security, and long-term viability and sustainability in their assessments of AI models rather than focusing solely on performance

3.2 Algorithmic Bias

AI models may inherit biases related to gender or ethnicity from their training data, or generate clinically harmful recommendations (Dennstädt et al., 2025). Generative AI can produce inaccurate statements that appear factual, a phenomenon known as "AI hallucinations," necessitating cautious implementation and human oversight in healthcare settings (Dennstädt et al., 2025). Patients are concerned that AI may prioritize commercial interests over patients' welfare, particularly by private companies that may not fully understand the healthcare system or might seek to leverage patient data for their own interests (Ossa et al., 2025). With the implementation of appropriate governance frameworks, healthcare organizations can foster transparency, more accurately monitor AI systems for bias and accuracy, establish accountability for AI-assisted decisions, and ensure that patient data are used only for authorized purposes. Healthcare organizations that implement effective governance measures will be in a better position to ensure that AI supports the integrity of patient care rather than compromising it.

3.3 Over-reliance by Medical Staff

Many patients are concerned that the increasing utilization of AI will lead to a lack individualized care. They fear that AI may not capture nonverbal cues or provide emotional support, risking the loss of personal interaction, which is considered essential in healthcare (Gundlack et al., 2025). These concerns may be exacerbated when medical professionals over-rely on AI systems, potentially overlooking critical details or relying on AI-generated recommendations without proper independent verification, particularly in understaffed healthcare settings (Gundlack et al., 2025).

If medical staff over-rely on AI, this can significantly diminish patients' ability to participate meaningfully in AI-supported care due to perceived complexity and a lack of education about their rights and AI's capabilities. This could promote passive acceptance of AI-generated medical decisions and minimize the patient's role in the healthcare process (Ossa et al., 2025). Concerns also exist regarding accountability if AI makes an error. Patients generally place greater trust in physicians than in AI, particularly when diagnoses conflict, underscoring the importance of proper medical oversight and establishing clear accountability and transparency when using AI (Gundlack et al., 2025).

Healthcare organizations must adopt clear protocols that require independent human oversight and verification of AI-generated decisions and recommendations. AI should supplement rather than supplant communication between medical professionals and patients. Healthcare organizations can reduce the risks associated with the over-reliance on AI while preserving accountability and patient participation.

3.4 Reliance On External Vendors

The use of external vendors in healthcare AI creates regulatory vulnerabilities that are often overlooked. While data collected by AI or robots operated by HIPAA-covered entities will in most cases be protected by the HIPAA Privacy, Security, and Breach Notification rules, which are discussed further below, AI or robots operated by non-covered entities are subject to a far less protective regulatory regime (Terry, 2019). This asymmetry creates opportunities for regulatory arbitrage, where entities structure their relationships to avoid HIPAA coverage altogether.

This arbitrage is facilitated by the fact that data analytics entities are often third parties with whom the patient has no direct relationship and against whom the patient cannot assert meaningful data protection rights (Terry, 2019). Data acquired from another entity—such as a supermarket's record of a patient's over-the-counter purchases—can be used to build healthcare-relevant profiles entirely outside the HIPAA-regulated zone (Terry, 2019). Because this activity falls outside existing regulatory boundaries, it is largely invisible to the oversight mechanisms designed to protect patients. Given these regulatory gaps, litigation is likely to emerge as a significant driver of accountability in the healthcare AI vendor ecosystem, filling the void left by an underinclusive regulatory framework.

Litigation involving healthcare AI is likely to raise unresolved questions of legal agency and accountability—including whether AI systems can bear a legal relationship to patients, who bears

ultimate decision-making responsibility between clinician and AI, and whether hospitals' growing investment in AI will expose them to direct or corporate liability claims traditionally reserved for independent contractor negligence (Terry, 2019).

These combined regulatory and litigation dynamics suggest that the vendor-driven structure of healthcare AI is creating accountability gaps that neither existing privacy law nor traditional liability doctrine was designed to address. The law is likely to evolve, and we will continue to see changes in the regulatory environment accordingly.

4. IMPLICATIONS OF EVOLVING CONSUMER PRIVACY LAWS

The protection of both individual and societal interests from surveillance and datafication requires a modern, non-domain-specific regulatory system—one that moves beyond siloed, sector-by-sector rules and instead embodies Fair Information Practice Principles across the AI lifecycle (Terry, 2019). Emerging privacy-respecting technologies, including federated learning, differential privacy, and homomorphic encryption, can preserve many of AI's benefits while protecting the individuals whose data underlie these systems. However, as discussed further below, these technical solutions must be complemented by regulatory frameworks that establish minimum data protection standards; technology alone cannot substitute for enforceable legal accountability.

While Congress and industry work toward a federal privacy law they can both live with, privacy advocates are increasingly concerned that any such legislation will be relatively weak and primarily directed at preempting more robust, emerging state laws. This concern is not unfounded: companies generally prefer a single, more permissive federal standard to a patchwork of stricter state regimes, giving them a strong incentive to support federal legislation mainly for its preemptive effect. The risk is that federal preemption would eliminate the state-level innovation that has driven privacy reform in recent years while establishing only minimal federal protections in its place (Terry, 2019).

This risk is compounded by the underlying structure of U.S. privacy law itself. The U.S. debate about stronger data protection remains nascent compared to the European context, largely because the U.S. continues to rely on a sectoral approach—separate, industry-specific laws such as HIPAA for health data and GLBA for financial data—rather than a comprehensive framework. This approach, along with outdated downstream protection models and a proliferation of domain-specific regulators, leaves significant coverage gaps and makes coordinated reform difficult.

Even as U.S. debates continue, the EU's broader regulatory framework reveals its own unresolved tensions. The GDPR's requirements of accountability, transparency, purpose, time limitations, and data minimization are arguably antithetical to the training of AI, its black-box algorithms, and the business models of AI and Big Data companies (Terry, 2019). As such, even a comprehensive, well-established legal framework for data privacy struggles to reconcile its core principles with the technical and commercial realities of modern AI development. Nonetheless, healthcare organizations in the United States must currently comply with both state and federal regulations governing their industry. Part of that compliance framework must include continuous monitoring and adaptation to remain compliant with the evolving legal landscape.

5. MAPPING COMMON USE AND COMMON RISKS OF AI IN THE HEALTHCARE INDUSTRY TO FEDERAL LAW AND REGULATORY MANDATES

Healthcare organizations are one of the most regulated industries in the United States. Being able to identify and understand the relationships among technology use, resulting organizational risks, and the regulatory environment allows an organization to prepare for and mitigate said risks by ensuring that the corresponding legal requirements are met.

Healthcare organizations must account for any business associates with whom they share PHI. A business associate is defined by law and is intended to cover any third-party entity with which a covered entity shares PHI (45 C.F.R. § 160.103). A "[b]usiness associate includes: (i) A Health Information Organization, E-prescribing Gateway, or other person that provides data transmission services with respect to protected health information to a covered entity and that requires access on a routine basis to such protected health information. (ii) A person that offers a personal health record to one or more

individuals on behalf of a covered entity. (iii) A subcontractor that creates, receives, maintains, or transmits protected health information on behalf of the business associate" (45 C.F.R. § 160.103). Business associates are held to the same security and privacy standards as covered entities (42 U.S.C. § 17931).

Healthcare organizations and their business associates risk civil and monetary penalties that could rise into the millions if they are found in violation of the law (45 C.F.R. § 160.404). Healthcare organizations also risk being held liable for violations of the law by their business associates, and vice versa (45 C.F.R. § 160.402). These exposures to liability underscore the need for healthcare organizations to understand and identify where their risks are, how they can be controlled and mitigated, and to be intentional in the continuous management of AI use.

5.1 Mapping Common Use

AI-powered tools used for clinical decision support (CDS), diagnostic imaging, and healthcare operations each involve the collection, generation, or transmission of Protected Health Information (PHI), and each therefore triggers overlapping HIPAA obligations. HIPAA's Privacy Rule (45 C.F.R. §164.500-524) prohibits a covered entity from using or disclosing PHI without patient consent except under recognized exceptions (45 C.F.R. §164.502). The broadest such exception, "routine use" for treatment, payment, and health care operations (45 C.F.R. §164.506) covers most clinical and administrative AI functions, including quality assessment, staffing, scheduling, fraud monitoring, and coding support (*Citizens for Health v. Leavitt*, 2005). Regardless of use case, organizations must disclose only the minimum PHI necessary to accomplish the tool's purpose (45 C.F.R. §164.502(b); §164.514(d)), and any AI tool hosted or provided by a third-party vendor makes that vendor a business associate, triggering the Business Associate Rule's contracting and safeguarding requirements (45 C.F.R. §164.502(e)). If unsecured PHI is breached in any of these contexts, HITECH's reporting requirements are triggered (42 U.S.C.S. §17932). While these regulatory triggers are common across use cases, each domain presents distinct compliance considerations:

Clinical Decision Support (CDS) tools that assist in diagnosis, triage, medication review, or care management risk violating the Minimum Necessary Rule if they collect or generate data beyond what is needed to accomplish a specific clinical task. A defensible compliance structure clearly defines the treatment or operations function for which the tool is used, limits data input/output accordingly, and documents business associate agreements for any vendor-hosted tool.

Diagnostic Imaging presents a heightened risk because imaging data often contains embedded metadata or identifiers that survive de-identification efforts, leaving a reasonable basis for re-identification (*Steinberg v. CVS Caremark Corp.*, 2012). Even "de-identified" imaging data should continue to be treated as PHI unless de-identification can be confirmed with high confidence. Compliance therefore requires particular attention to whether de-identification is adequate, whether secondary uses exceed the tool's original scope, and whether storage and transmission controls are sufficient.

Operations functions such as scheduling, staffing, resource allocation, fraud monitoring, and utilization analysis generally fall within HIPAA's routine use exception, the broadest carve-out under the law. The principal compliance risk here is scope creep: organizations must ensure AI-assisted operational tools remain confined to recognized operational categories and do not drift into unintended secondary uses that would trigger the Minimum Necessary Rule.

Across all three domains, an effective compliance program applies the minimum-necessary standard by default, maintains legally compliant business associate agreements with any third-party vendor, and documents the specific operational or clinical purpose for which each AI tool is deployed.

5.2 Mapping Common Risks

5.2.1 Re-identification of Data

As discussed above, there are several benefits to de-identified data, including being able to train an LLM. While the law sets forth specific standards for how a covered entity and its business associates

handle ePHI (45 CFR Subtit. A, Subch. C, Pt. 164, Subpt. C), de-identified ePHI data falls outside the scope of the law because it, in theory, should lack the ability to identify a patient (42 U.S.C. § 17935(b)(4); 45 C.F.R. §164.502(d)). However, even with the best efforts, there is always a risk that the data is re-identifiable. If data is susceptible to re-identification and found to be re-identifiable, then HIPAA's Privacy Rule, Minimum Necessary Rule, and Business Associate Rule would apply to said data. Also, HIPAA and HITECH breach notifications would be triggered if any information were breached.

Complex AI systems pose high linkage risk when they are asked, for example, to analyze large data sets, free text, or multimedia data—such as imaging data—which increases re-identification risks that need to be carefully considered when an organization depends on the de-identification of data to reduce its security protocol. A best practice is to treat all data as identifiable when it is used in conjunction with an AI system, particularly a complex one.

5.2.2 Secondary Uses of Health Information

There are infinite secondary uses for PHI, including training AI models, research, validation, workflow analytics, and more. In some instances, healthcare organizations arguably have an ethical obligation—such as when training AI models—to use secondary data to drive quality improvement and better patient outcomes. The law restricts the use of PHI without an individual's consent unless an exception applies (45 C.F.R. §164.502). And, as discussed above, the law indeed establishes exceptions under which PHI and ePHI may be used without individual consent and sets forth parameters for such use (45 C.F.R. §164.506).

Secondary uses of data are beneficial but also carry inherent risks that must be accounted for within the larger AI management framework. For example, if an organization relies on its secondary use of data being exempt from the law under a "routine use" exemption, the organization must ensure that the data remains within that use zone and put clear measures in place to monitor its use. Under the laws already discussed, there are additional concerns about the re-identification of data used for a secondary purpose and the permissibility of its use.

5.2.3 Algorithmic Bias

While algorithmic bias is not yet expressly covered under federal law, there is a deep ethical obligation for healthcare organizations to engage methodologies and processes that specifically work to reduce or eliminate algorithmic bias.

5.2.4 *Reliance on External Vendors*

As discussed in Section 3.4, reliance on external AI vendors triggers HIPAA's Business Associate Rule (45 C.F.R. §164.502(e)) and requires a written agreement meeting specific statutory terms (45 C.F.R. §164.502(e)(2)). Because many healthcare AI tools are vendor-hosted, organizations must ensure these agreements are legally sound and reviewed regularly as vendor relationships and the underlying technology evolve.

6. MODEL LEGAL AND ETHICAL COMPLIANCE FRAMEWORK FOR HEALTHCARE ORGANIZATIONS TO GOVERN AI DEPLOYMENT AND USE

This framework is intended to be adaptable and is focused on a risk-based approach to managing AI use in a healthcare organization. A risk-based approach will examine the risks posed by the use and adoption of AI within the organization and provide solutions to mitigate them. A risk-based framework can be implemented enterprise-wide, governing at the larger organizational level, or deployed at the divisional level within the organization. Depending upon the tailored needs of an institution, organizations may implement these policies through existing governance channels that focus on micro-areas—such as a patient-safety framework that would ask a narrower question of whether this AI will harm a patient—rather than creating new structures, provided accountability, review, documentation, and escalation remain clear within each of those narrower areas.

More complex or higher-risk organizations may require deeper review, more formal approvals, and stronger monitoring; smaller organizations may use lighter processes while preserving the same core

controls across different areas using the technology. Going beyond the legal minimum allows an organization to incorporate ethical considerations and its own values. An organization may look to its institutional or divisional mission(s) to better understand how ethics and ethical considerations can be incorporated into its compliance framework.

6.1 Purpose and Scope

Healthcare organizations should expressly establish and articulate the purpose of their framework and set forth the scope it intends to govern. As discussed above, an organization should consider the approach it will take to accomplish its goals and articulate its purpose and scope accordingly. Stating a purpose and scope is critical to laying the foundation for the framework. Language in this section will address the why (purpose) and what (scope) of the model. Why the model is being implemented—to address the deployment, implementation, and management of AI in the organization- and what it is intended to do—direct and support the legal and ethical use of AI across the organization.

6.2 Governance and Accountability

Clear lines of accountability for the governance of artificial intelligence (AI) systems, encompassing responsibility for policy development, intake procedures, review processes, approval determinations, ongoing monitoring, and incident escalation protocols, are essential. Governance structures should incorporate the involvement of relevant institutional stakeholders. Depending on organizational context, such stakeholders may include representatives from compliance, legal affairs, privacy, information security, clinical leadership, procurement, information technology, and relevant operational units.

AI applications identified as high-risk should be subject to formal escalation procedures, with approval authority vested in senior leadership. Applications carrying material implications for patient safety, legal or regulatory compliance, ethical integrity, or broader enterprise risk warrant an elevated level of oversight, proportionate to the magnitude and nature of the risks involved. For each AI use case, the organization should identify: a business or operational owner, a technical or implementation owner, a privacy and security review path, a clinical owner for patient-care uses, and an escalation path for exceptions.

6.3 Risk Classification and Approval

Risk classification should account for several interconnected factors. Chief among them is intended use and user profile: clinical decision support, administrative automation, patient-facing interaction, and research applications each carry a different risk calculus, particularly with respect to potential impact on patient care, clinical safety, individual rights, and equitable access. The nature, volume, and sensitivity of the underlying data matter as well, including whether it constitutes PHI or another regulated category and the practical likelihood of re-identification.

Additional considerations include the degree of automation and human reliance on system outputs (a fully autonomous clinical tool poses different risks than one offering suggestions subject to clinician review), the involvement of third-party vendors or processors, the depth of integration with clinical or operational infrastructure, and the applicable legal and regulatory landscape. Organizations should also assess the potential consequences of error, bias, model drift, misuse, or system outages. No AI tool should advance to pilot testing or production until all required reviews and approvals are complete.

6.4 Data, Privacy, and Security Controls

Effective healthcare AI requires diverse, high-volume medical datasets that a single provider typically cannot supply under today's siloed data infrastructure (Barbaria, et al., 2025; Feldman et al., 2018). The industry's continued reliance on a small number of centralized, privately owned Electronic Health Record systems—built primarily for record-keeping, billing, and baseline regulatory compliance—makes balancing innovation with privacy protection an ongoing challenge (Holmgren et al., 2025).

Federated learning offers one path forward: healthcare providers share only de-identified AI model updates generated from locally maintained records, rather than sharing the underlying PHI itself. This limits direct PHI exposure, though inference or reconstruction attacks on shared model updates remain a residual risk (Pati et al., 2024). A complementary approach pairs federated learning with blockchain technology, using smart contracts to embed HIPAA and GDPR compliance directly into data-sharing transactions; because blockchain transactions are immutable and automatically vetted, this model can make compliance and secure data-sharing inherent to the infrastructure itself rather than dependent on manual enforcement (Barbaria et al., 2025). More broadly, shifting from traditional "data-to-model" architectures toward decentralized "model-to-data" architectures—where PHI remains in place, and AI models access it only to gather updates—represents a promising direction for reconciling AI's need for diverse training data with patients' privacy interests (Pati et al., 2024).

These interoperability gains matter beyond compliance: without broad, representative datasets, AI models risk underestimating the health needs of communities with less frequent healthcare access, compounding existing disparities in care (Mittermaier et al., 2023). Models trained predominantly on data from insured, highly connected patient populations risk becoming biased against and less accurate for underserved communities. Emerging federated and decentralized data architectures, if adopted broadly, could meaningfully reduce this risk while strengthening privacy protections simultaneously.

6.5 Third-party and Vendor Controls

Beyond the Business Associate Agreement requirements discussed above, third-party AI providers warrant risk-based due diligence given their access to sensitive data and influence over regulated workflows. At minimum, this review should assess the product's purpose and functionality, data flows, hosting architecture, processing locations, and the vendor's scope of access.

Vendor contracts should address incident and change management, governance responsibilities, and a clear exit plan. By default, organizations should prohibit vendors from using organizational data for unrelated model training, commercialization, or generalized product improvement absent separate authorization.

6.6 Validation, Human Oversight, and Monitoring

It is important to manage and monitor automation maturity, use AI to enhance safety capabilities, highlight potential hazards, and align AI system responses with the perceived level of hazard. Autonomous functionality should be integrated with other safety systems to maintain redundancy, and the scale and speed of AI development should align with the capacity to analyze and learn from these activities. Systematic efforts should also be made to gather data from people interacting with AI systems (Macrae, 2025).

6.7 Incident Management and Reporting

Organizations should formally define processes for identifying, escalating, investigating, and resolving AI-related incidents. These incidents may encompass a broad range of categories, including unauthorized access and loss of sensitive data. Incidents involving sensitive data breaches, threats to patient safety, and regulated activities must be escalated promptly to the appropriate security team. Additionally, organizations should systematically evaluate whether such incidents trigger obligations for formal notification, regulatory reporting, and suspension of the AI tool, in accordance with policy requirements.

6.8 End-User Training, Documentation, and Review

Role-appropriate training programs for all personnel engaged with artificial intelligence systems are essential. Such training should address the scope of approved and prohibited applications of AI, requirements governing data handling, protocols for incident reporting, and the consequences associated with unapproved or improper use. Organizations should concurrently maintain governance records demonstrating the application has been appropriately reviewed. These records should

document the intended use and designated owner, the applicable risk classification, outcomes of data and privacy assessments, security reviews, and formal monitoring outcomes.

Organizations should conduct periodic reviews of their overall AI portfolio, governance framework, and the effectiveness of associated controls. Reviews should assess whether all AI tools have been comprehensively inventoried and monitoring activities have been performed promptly. Additionally, the review process should evaluate whether incidents and associated lessons learned have been adequately addressed, and whether existing policies and controls remain aligned with applicable legal requirements, industry standards, and the organization's evolving risk profile. This layered approach supports a comprehensive and adaptive framework for the responsible oversight of AI deployment within the organization.

7. CONCLUSION

The integration of artificial intelligence into healthcare delivery and information systems offers substantial gains in diagnostic accuracy, clinical efficiency, and operational performance while exposing organizations to complex legal, ethical, and regulatory risks. AI deployment across clinical decision support, diagnostic imaging, and administrative operations implicates specific HIPAA and HITECH provisions while introducing risks such as re-identification, algorithmic bias, staff over-reliance, and vendor dependency that existing compliance frameworks were not designed to anticipate.

The fragmented nature of U.S. privacy law, combined with data analytics entities operating outside HIPAA's regulatory perimeter, creates accountability gaps not resolved by current statutory and common law doctrine, suggesting litigation will likely serve as an important, if imperfect, mechanism for establishing accountability norms in the interim.

This paper proposes a risk-based governance framework to provide healthcare organizations with a practical, adaptable structure for managing the legal and ethical dimensions of AI deployment. By emphasizing clear governance and accountability structures, rigorous risk classification, robust data and vendor controls, sustained human oversight, and systematic incident management and training protocols, the framework translates abstract regulatory and ethical obligations into operational practice. It adjusts to organizational size, risk tolerance, and mission, allowing institutions to calibrate governance intensity to the materiality of the risks presented by a given AI use case.

5. REFERENCES

Abbas, Q., Jeong, W., & Lee, S. W. (2025). Explainable AI in Clinical Decision Support Systems: A Meta-Analysis of Methods, Applications, and Usability Challenges. *Healthcare*. https://doi.org/10.3390/healthcare13172154

Amount of a civil money penalty, 45 C.F.R. § 160.404 (2025).

Arzikulov, F., & Komiljonov, A. (2025, October 20). AI-powered diagnostic systems in radiology: enhancing precision, speed, and clinical decision-making. *Academic Journal of Science, Technology and Education, 1*(6), 16-23.

Barbaria, S., Jemai, A., Ceylan, H. İ., Muntean, R. I., Dergaa, I., & Rahmouni, H. B. (2025). Advancing Compliance with HIPAA and GDPR in Healthcare: A Blockchain-Based Strategy for Secure Data Exchange in Clinical Research Involving Private Health Information. *Healthcare, 13*(20), 2594. https://doi.org/10.3390/healthcare13202594

Basis for a civil money penalty, 45 C.F.R. § 160.402 (2026).

Citizens for Health v. Leavitt, 428 F.3d 167 (3d Cir. 2005)

Conduah, A. K., Ofoe, S., & Siaw-Marfo, D. (2025). Data privacy in healthcare: Global challenges and solutions. DIGITAL HEALTH, 11, 1-19. https://doi.org/10.1177/20552076251343959

Dennstädt, F., Hastings, J., Putora, P. M., Schmerder, M., & Cihoric, N. (2025, March 6). Implementing large language models in healthcare while balancing control, collaboration, costs and security. npj digital medicine. https://doi.org/10.1038/s41746-025-01476-7

Feldman, K., Faust, L., Wu, X., Huang, C., & Chawla, N. V. (2017). Beyond Volume: The Impact of Complex Healthcare Data on the Machine Learning Pipeline. *arXiv, 10344*, 150-169. https://doi.org/10.48550/arXiv.1706.01513

Gundlack, J., Negash, S., Thiel, C., Buch, C., Schildmann, J., Unverzagt, S., Mikolajczyk, R., & Frese, T. (2025). Artificial Intelligence in Medical Care – Patients' Perceptions on Caregiving Relationships and Ethics: A Qualitative Study. Health Expectations. https://doi.org/10.1111/hex.70216

Health Information Technology for Economic and Clinical Health Act, 42 U.S.C. § 17931 (2010).

Health Information Technology for Economic and Clinical Health Act, Notification in the Case of Breach, 42 U.S.C. § 17932 (2010).

Health Information Technology for Economic and Clinical Health Act, 42 U.S.C § 17935(b)(4) (2023).

HIPAA Covered Entity Definition, 45 C.F.R. § 160.103 (2026).

Holmgren, A. J., Apathy , N. C., & Kanter , G. P. (2025, August 18). Electronic health record market consolidation and implications for cybersecurity. Health Affairs Scholar, 3(8). https://doi.org/10.1093/haschl/qxaf164

IBM Corporation. (2025). Cost of a Data Breach Report 2025 The AI Oversight Gap. https://www-api.ibm.com/adobe/assets/urn:aaid:aem:607b9590-38e0-4c91-b433-aa8a17f5b5e8/original/as/cost-of-a-data-breach-2025-full-report.pdf

Karunanayake, N. (2025). Next-generation agentic AI for transforming healthcare. *Informatics and Health*, 73-83. https://www.keaipublishing.com/en/journals/informatics-and-health/

Macrae, C. (2025). Managing risk and resilience in autonomous and intelligent systems: Exploring safety in the development, deployment, and use of artificial intelligence in healthcare. *Society for Risk Analysis, 45*, 910–927. https://doi.org/10.1111/risa.14273

Magrabi, F., Ammenwerth, E., McNair, J. B., De Keizer, N. F., Hyppönen, H., Nykänen, P., Rigby, M., Scott, P. J., Vehko, T., Wong, Z.-Y., & Georgiou, A. (2019). Artificial Intelligence in Clinical Decision Support: Challenges for Evaluating AI and Practical Implications. *IMIA Yearbook of Medical Informatics*. https://doi.org/10.1055/s-0039-1677903

Miyoshi, N. (2025, January 15). Use of AI in Diagnostic Imaging and Future Prospects. *Japan Medical Association Journal, 8*(1), 198-203. https://doi.org/10.31662/jmaj.2024-0169

Murdoch, B. (2021). Privacy and artificial intelligence: challenges. *BMC Medical Ethics*. https://doi.org/10.1186/s12910-021-00687-3

Nguyen, T. D., Whaley, , C. M., Simon, K., Mehta, N., Yu, H., McBain, R. K., Mehrotra, A., & Cantor, J. H. (2025). Adoption of Artificial Intelligence in the Health Care Sector. JAMA Health Forum, 6(11). https://doi.org/10.1001/jamahealthforum.2025.5029

Ogut, E. (2025). Artificial Intelligence in Clinical Medicine: Challenges Across Diagnostic Imaging, Clinical Decision Support, Surgery, Pathology, and Drug Discovery. *Clinics and Practice*. https://doi.org/10.3390/clinpract15090169

Other requirements relating to uses and disclosures of protected health information, 45 C.F.R. § 164.514(d) (2026).

Pati, S., Kumar, S., Varma, A., Edwards, B., Lu, C., Qu, L., Wang, J. J., Lakshminarayanan, A., Wang, S.-h., Sheller, M. J., Chang, K., Singh, P., Rubin, D. L., Kalpathy-Cramer, J., & Bakas, S. (2024, July 12). Privacy preservation for federated learning in health care. Patterns, 5(7). https://doi.org/10.1016/j.patter.2024.100974

Saeidnia, H. R., Firuzpour, F., Kozak, M., & Hooman, S. m. (2025, January 25). Advancing cancer diagnosis and treatment: integrating image analysis and AI algorithms for enhanced clinical practice. *Artificial Intelligence Review*. https://doi.org/10.1007/s10462-025-11117-w

Security and Privacy, 45 C.F.R. §164.500-524 (2026).

Security standards for the protection of electronic protected health information, 45 C.F.R. § 164.302–164.318 (2013).

Security standards for the protection of electronic protected health information: Applicability, 45 C.F.R. § 164.302 (2013).

Steinberg v. CVS Caremark Corp., 899 F. Supp. 2d 331 (E.D. Pa. 2012)

Sutton, R. T., Pincock, D., Baumgart , D. C., Sadowski , D. C., Fedorak, R. N., & Kroeker, K. I. (2020). An overview of clinical decision support systems: benefits, risks, and strategies for success. *Digital Medicine*. https://doi.org/10.1038/s41746-020-0221-y

Uses and Disclosures of Protected Health Information, 45 C.F.R. §164.502(b) (2026).

Uses and disclosures of protected health information, 45 C.F.R. §164.502 (2026).

Uses and disclosures of protected health information: Business associate contracts, 45 C.F.R. §164.502(e) (2024).

Uses and Disclosures to Carry Out Treatment, Payment, or Health Care Operations, 45 C.F.R. §164.506 (2013).

APPENDIX A

Core Principles

Core principles should align with an organization's mission and values. Core Principles should establish the standards the framework is meant to meet. An example set of Core Principles in a risk-based framework can be as follows:

1. AI must have a clear and approved purpose before testing or deployment, approved in accordance with the management framework set forth.

2. Data used in AI must be classified before use, including whether it is identifiable, sensitive, regulated, shared, or otherwise high risk.

3. Only the minimum data reasonably necessary for the intended purpose should be used. There will be accountable human oversight validating said data use.

4. Secondary uses of data, including model training, tuning, benchmarking, analytics, or product improvement, require separate review and approval in accordance with the management framework set forth.

5. AI must be subject to risk-based governance, with stronger controls for higher-risk uses as outlined in this management framework.

6. Human accountability must remain in place for decisions that affect patient care, safety, rights, access, or other material outcomes.

7. AI tools and third-party vendors must be included within the organization's privacy, security, and third-party risk management processes.

8. AI systems must be validated before use and monitored during and after operation.

9. AI systems must always maintain accountable human oversight.

10. Material changes to an AI tool, its use case, its data, or its operating environment require reassessment and approval in accordance with the management framework set forth.
11. The organization must maintain sufficient documentation to demonstrate governance, oversight—including human oversight—and control effectiveness.

Of course, organizations need to ensure that their stated principles are legally defensible and tailored to their organization's nuanced needs. Including relevant stakeholders, such as legal counsel and compliance officers, in the process helps ensure that this is indeed the case.