Sheppard

Healthy AI

2025 Year In Review

Introduction

Artificial intelligence (AI) in healthcare may have reached an inflection point in 2025. After being defined by proofs of concept and isolated pilots for more than a decade, recent advances signal that AI is moving decisively into the operational core of health systems, life sciences organizations, and payer workflows. What once lived primarily in innovation labs and limited production environments is now increasingly embedded in day-to-day infrastructure across the industry.

In clinical decision support, AI has evolved from experimental triage tools to systems capable of continuously learning and integrating with electronic health records. Many organizations now leverage these technologies to help clinicians interpret images and tests, anticipate patient deterioration, and personalize care pathways at scale. In revenue cycle management and utilization management, automation has begun shifting from back-office augmentation to supporting real-time adjudication and documentation, reducing administrative friction, and allowing staff to focus on higher-value work, though adoption varies regionally and by institution. Meanwhile, AI is becoming foundational in drug development and clinical trials, accelerating target discovery, optimizing trial design, and improving patient recruitment and monitoring with leading life sciences firms actively deploying these capabilities. AI has also been widely deployed across provider organizations to reduce administrative burden and assist with record preparation.

2025 marked not only a period of technological progress, but a transition towards embedding AI into an integral part of healthcare's operational fabric. This 2025 review provides a high-level summary of how this transition is unfolding at the international, federal, and state levels, and what it may signal for the future of healthcare delivery and innovation.

Authors

Carolyn Metnick
Partner • Chicago
+1.312.499.6315
cmetnick@sheppard.com
Bio

Megan Miller
Associate • Dallas
+1.469.391.7427
memiller@sheppard.com
Bio

Arushi Pandya
Associate • Washington DC
+1.202.747.3228
apandya@sheppard.com
Bio

Alexandria Foster
Associate • Washington DC
+1.202.747.3267
afoster@sheppard.com
Bio

Alyanna Gallo
Associate • New York
+1.212.653.8709
agallo@sheppard.com
Bio

Daniel Shemano
Associate • Century City
+1.424.288.5371
dshemano@sheppard.com
Bio

Tina Watson
Associate • New York
+1.212.896.0678
twatson@sheppard.com
Bio

Lena Zinner
Associate • New York
+1.212. 896.0674
lzinner@sheppard.com
Bio

Michael Sutton
Associate • Dallas
+1.469.391.7455
msutton@sheppard.com
Bio

Michael Orlando
Partner • San Diego (Del Mar)
+1.858.720.8932
morlando@sheppard.com
Bio

Madeline Cunnion
Associate • Washington DC
+1.202.747.2672
mcunnion@sheppard.com
Bio

Stephanie Awanyai-Ufondu
Associate • Century City
+1.310.228.6168
sawanyai-ufondu@sheppard.com
Bio

Esperance Becton
Associate • Washington DC
+1.202.747.3230
ebecton@sheppard.com
Bio

Christina Nguyen
Associate • New York
+1.212.896.0676
cmnguyen@sheppard.com
Bio

Timothy Rozier-Byrd
Associate • Washington DC
+1.202.747.2671
trozier-byrd@sheppard.com
Bio

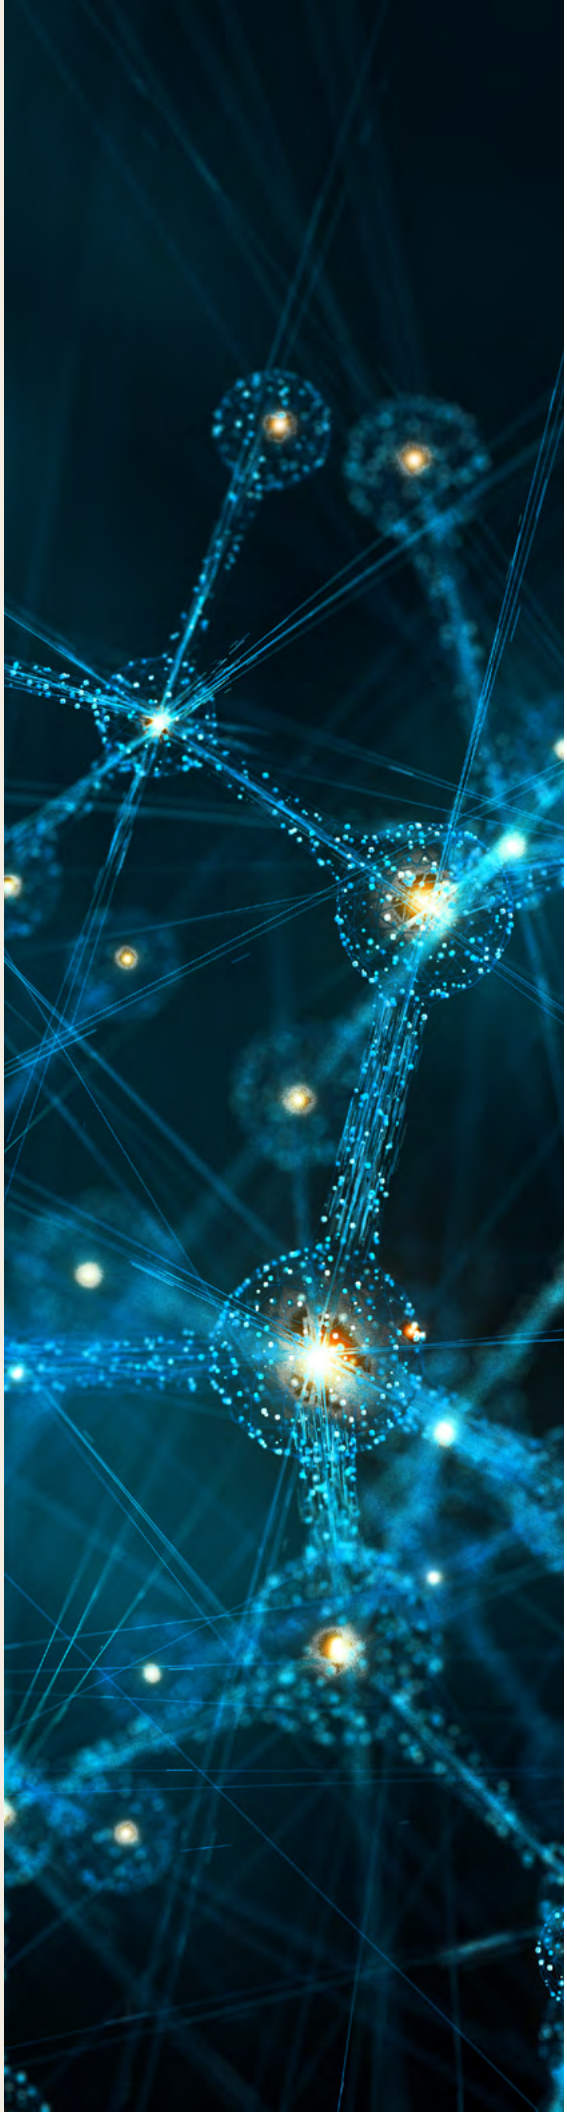

Contents

Legal and Policy Developments

International Level

Status of the EU AI Act in 2025

The European Union (EU) AI Act (EUAI Act) was formally adopted as a regulation on August 1, 2024. It applies directly across all EU member states without the need for national implementing legislation. The legislation, at a high level, aims to provide a harmonized legal framework "for the development, the placing on the market, the putting into service and the use of artificial intelligence systems" in the EU. The EUAI Act adopts a phased implementation in which different categories of obligations are poised to take effect at various times over a multiyear period.

On February 2, 2025, the first set of obligations under the EUAI Act took effect, including prohibitions on "unacceptable risk" AI systems. This category covers a defined set of banned uses, including systems that manipulate behavior through subliminal techniques, exploit vulnerable populations, engage in social scoring, or attempt to predict criminality based solely on profiling. This first set of obligations also bars untargeted scraping for facial-recognition databases, emotion recognition in workplaces and schools, and certain public-space biometric identification uses outside narrow law-enforcement exceptions.

On August 2, 2025, additional requirements entered into force, including governance rules that require EU member states to designate competent national authorities, establish notifying bodies, and ensure that the EU AI Office is operational. In the same phase, new obligations applied to providers of general-purpose AI models, meaning broadly capable foundation models (such as large language or image models) that can be used for many different downstream applications. These obligations include preparing technical documentation demonstrating how the model was developed, trained, and intended to be used, publishing a summary of the training data used, supporting downstream deployers, and, for models designated as posing systemic risk, implementing risk-mitigation measures, and reporting serious incidents.

Remaining obligations continue to phase in, with high-risk AI system requirements generally scheduled to apply starting August 2, 2026, and later compliance deadlines extending to August 2027.

Global Initiative on AI for Health

The Global Initiative on Artificial Intelligence for Health (GI-AI4H) is a joint program of the World Health Organization, the International Telecommunication Union, and the World Intellectual Property Organization, and is established to coordinate international standards, policy development, and regulatory capacity relating to AI in health systems. In 2025, the initiative continued its focus on creating interoperable governance frameworks, technical benchmarks, and evaluation methodologies to support the safe and equitable adoption of AI in healthcare, with particular emphasis on expanding regulatory readiness in low- and middle-income settings. Workstreams include guidance on ethical deployment, evidence generation, data quality, intellectual property considerations, and system-level integration, alongside multilateral collaboration and knowledge-sharing mechanisms intended to support member states in designing oversight structures.

GI-AI4H operates as both a normative body and a convening platform, coordinating multilateral collaboration across governments, industry, and research institutions. The initiative also facilitates knowledge-sharing, model benchmarking, and dissemination of best practices to enable member states to develop oversight structures, support responsible innovation, and implement AI systems within clinical care and public-health service delivery.

Principles for the Secure Integration of Artificial Intelligence in Operational Technology

International cybersecurity and intelligence agencies have articulated a common set of expectations for safe AI deployment that provide an important reference point for health sector stakeholders, including the American Hospital Association. In December 2025, the U.S. Cybersecurity & Infrastructure Security Agency, the Australian Signals Directorate's Australian Cyber Security Centre, the National Security Agency's Artificial Intelligence Security Center, the Federal Bureau of Investigation, the Canadian Centre for Cyber Security, the German Federal Office for Information Security, the Netherlands National Cyber Security Centre, the New Zealand National Cyber Security Centre, and the United Kingdom National Cyber Security Centre jointly issued the "Principles for the Secure Integration of Artificial Intelligence in Operational Technology." This guidance is directed to critical infrastructure owners and operators that rely on operational technology to provide essential public services and identifies AI integration as a source of both efficiency gains and significant new safety and security risks.

The authoring agencies organized their recommendations around four core principles: (i) understanding AI and its unique risks in Operational Technology (OT) environments; (ii) evaluating whether AI is appropriate for the specific OT use case and managing associated data security risks; (iii) establishing governance and assurance frameworks for AI in OT, including clear roles, responsibilities, and integration with existing cybersecurity controls; and (iv) embedding monitoring, human oversight, and failsafe mechanisms into AI and AI enabled OT systems. The guidance further emphasizes secure system development lifecycles, protection of OT data (including data assurance and sovereignty), transparency and security obligations for OT vendors that embed AI into their products, thorough testing and evaluation of AI systems in non-production environments before deployment, and explicit incorporation of AI related failure modes and attack vectors into incident response and functional safety planning.

U.S. Federal Level – Executive Orders and Other Guidance

Executive Order 14179 on Removing Barriers to American Leadership in Artificial Intelligence (EO 14179)

On January 23, 2025, President Trump signed EO 14179. This was the Trump administration's first executive order addressing AI, but later executive orders were issued in support of the goals outlined in EO 14179.

The Trump administration's stated goal is to ensure that the U.S. continues to be at the forefront of the development of AI. The administration believes that AI should be free of any ideological biases or social agendas. The prior administration under President Joe Biden established a policy framework in 2023 through Executive Order 14410 (EO 14410), that was designed to establish the first AI Governance Model. Several objectives of the Biden order included ensuring safety and security, advancing equity, and protecting citizens' privacy. The Trump administration believed that EO 14410 stifled the ability of AI developers to innovate, and repealed EO 14410 in EO 14179, outlining a new policy to promote AI development. All policies, directives, regulations, orders, and other actions taken pursuant to the revoked Biden administration EO 14410 on AI were directed to be reviewed, and any actions taken pursuant to EO 14410 that are or may be inconsistent with or present obstacles to the stated goals of EO 14179 were to be suspended, revised, or rescinded.

EO 14179 outlined an Artificial Intelligence Action Plan that must be established within 180 days of the signing of EO 14179. The national policy memorandum, Winning the Race: America's AI Action Plan, was published on July 10, 2025, and outlines the policy priorities for the administration on AI. The AI Action Plan identifies healthcare as "critical infrastructure" and contains several healthcare priorities, focusing on (i) accelerating healthcare AI innovation, (ii) reducing regulatory barriers, (iii) developing national healthcare AI standards, (iv) improving AI cybersecurity for healthcare systems, and (v) establishing AI testing sandboxes and Centers of Excellence.

EO 14179 and the AI Action Plan both signify the Trump administration's goal to ensure the competitiveness of the U.S. in AI innovation, as a leader that is in the global forefront of AI development. The administration believes that the best way to maintain American competitiveness is to remove barriers and regulations that may put restrictions on the efforts of AI developers. This is a definitive break from the AI vision of the Biden administration's efforts to take broad precautions and promulgate regulations in the AI field, in favor of a regulatory scheme that focuses on specific risks and eliminates broad based restrictions that are believed to stifle competition. The administration believes that the country will be able to modernize the AI infrastructure more effectively by removing barriers on AI developers in both the private and public sector.

For healthcare companies and providers, EO 14179 provides insights into the Trump administration's vision on the future of AI development. The removal of the prior emphasis on regulations on AI could also mean decreased compliance costs for healthcare companies that are developing AI tools. The potential removal of regulations could also lead to the acceleration of AI-assisted medical research efforts.

OMB Memoranda: M-25-21 and M-25-22

On April 7, 2025, the White House Office of Management and Budget (OMB) issued two memoranda, M-25-21 and M-25-22, to implement EO 14179 and to provide guidance to federal agencies on the use and acquisition of AI. These memoranda established the administration's policy framework for federal AI deployment and procurement and formally rescinded prior AI guidance issued under the Biden administration.

Memorandum M-25-21

OMB Memorandum M-25-21, "Accelerating Federal Use of AI through Innovation, Governance, and Public Trust," addresses how federal agencies may develop, deploy, and oversee AI systems. The memorandum reflects the administration's emphasis on accelerating AI adoption across the federal government by removing unnecessary procedural and regulatory barriers that could limit innovation. It encourages agencies to expand their use of AI to improve mission delivery, operational efficiency, and public services, while requiring agencies to establish internal governance structures to oversee AI use. Rather than imposing broad, precautionary restrictions, the memorandum adopts a risk-based approach that focuses oversight on higher-impact AI applications. Although considerations related to privacy, civil rights, and security remain part of agency governance obligations, these considerations are framed to support responsible innovation rather than to constrain AI development. M-25-21 formally rescinds the Biden administration's prior OMB guidance governing federal agency use of AI.

Memorandum M-25-22

OMB Memorandum M-25-22, "Driving Efficient Acquisition of Artificial Intelligence in Government," focuses on the procurement and acquisition of AI technologies by federal agencies. This memorandum is intended to streamline how agencies purchase and integrate AI tools and services by reducing complexity, delays, and compliance burdens associated with AI acquisition. It emphasizes that agencies should acquire AI systems that are mission-effective, secure, and reliable, while avoiding procurement requirements that could discourage private-sector innovation or slow adoption. The memorandum aligns federal acquisition practices with the administration's broader goal of removing barriers to AI development and encourages engagement with the private sector to ensure access to cutting-edge AI capabilities. M-25-22 rescinds and replaces prior Biden-era OMB acquisition guidance related to AI.

The issuance of OMB Memoranda M-25-21 and M-25-22 has particular significance for healthcare entities that develop, deploy, or rely on AI in connection with federal healthcare programs and agencies. For health information technology vendors that sell AI-enabled products or services to federal healthcare agencies, including agencies such as the U.S. Department of Health and Human Services (HHS), the U.S. Department of Veterans Affairs, and the Centers for Medicare & Medicaid Services (CMS), the memoranda signal a more favorable procurement environment. By streamlining acquisition processes and reducing broad compliance requirements, the guidance lowers barriers to entry for vendors offering innovative AI solutions, particularly in areas such as clinical decision support, operational analytics, and population health management. At the same time, vendors should expect continued scrutiny of AI systems used in high-impact or mission-critical contexts, requiring careful attention to security, reliability, and performance expectations established by federal customers.

Hospitals and health systems participating in federally sponsored pilot programs or demonstration projects may also see expanded opportunities to deploy AI tools under the new framework. The administration's emphasis on accelerating AI adoption and reducing regulatory friction may enable federal agencies to launch or scale AI-enabled healthcare pilots more quickly, including programs focused on care delivery, research, and administrative efficiency. Hospitals engaged in such pilots should be prepared for increased flexibility in how AI technologies are tested and evaluated, while recognizing that federal agencies will continue to apply risk-based oversight to ensure patient safety, data protection, and operational integrity.

Although the memoranda reflect a shift away from broad, AI-specific regulatory requirements, they do not eliminate existing healthcare and information security obligations. Healthcare companies and providers using AI systems in federal contexts must continue to align their practices with applicable requirements under the HIPAA Security Rule, particularly with respect to safeguarding electronic protected health information. In addition, AI systems that are hosted in cloud environments or otherwise subject to federal information security standards may still be required to comply with the Federal Risk and Authorization Management Program (FedRAMP), depending on the nature of the deployment and the agency involved. As a result, healthcare entities should view the OMB memoranda as reducing AI-specific compliance burdens while reinforcing the importance of integrating AI governance with established healthcare privacy and security frameworks.

Executive Order 14363 on Genesis Mission (EO 14363)

On November 24, 2025, President Trump signed EO 14363, establishing the "Genesis Mission," a landmark national initiative designed to accelerate scientific discovery and bolster U.S. leadership in AI-driven research and innovation. The order charges the U.S. Department of Energy (DOE) with leading this mission and creating an American Science and Security Platform, an integrated AI research infrastructure that unites high-performance computing, federal scientific datasets, and advanced modeling tools into a secure, collaborative system for hypothesis testing, experiment automation, and accelerated scientific workflows.

The scope of the Genesis Mission extends across the foundational pillars of national competitiveness, including energy dominance, national security, and transformational scientific breakthroughs. EO 14363 directs the Secretary of Energy, in coordination with the Assistant to the President for Science and Technology, to harness the U.S.'s national laboratory network, supercomputing resources, and federal research assets to support large-scale AI model training, simulation, and discovery efforts. Priority scientific challenges identified for this mission include biotechnology, quantum information science, advanced materials, semiconductors, and critical energy technologies.

In healthcare and biomedical research, the Genesis Mission's unified AI infrastructure could lower the barrier to participation for clinical and biomedical researchers by providing access to advanced computing platforms and large, integrated datasets that were previously siloed across agencies and institutions. By enabling sophisticated AI-assisted hypothesis generation and automated experiment design, the initiative aims to shorten research timelines for drug discovery, genomic analysis, and disease modeling. Such capabilities have the potential to accelerate the pace of innovation, strengthen public-private research partnerships, and improve the coordination of federal health research and development efforts, particularly as agencies leverage shared platforms and tools.

Overall, the Genesis Mission reflects a strategic push to harness AI as a transformative scientific tool, not only to maintain American dominance in global AI and science, but also to expand the nation's capacity to tackle complex scientific and health-related challenges more efficiently and collaboratively than ever before.

Executive Order 14365 on Ensuring a National Policy Framework for Artificial Intelligence (EO 14365)

On December 11, 2025, President Trump signed EO 14365 signaling a major federal push to centralize AI governance and reduce regulatory fragmentation across the U.S. EO 14365 underscores the Trump administration's policy that U.S. leadership in AI is critical to national and economic security, and that a "minimally burdensome national policy framework" for AI is needed to maintain competitiveness and avoid the complications of 50 divergent state regulatory regimes.

EO 14365 directed several federal actions aimed to preempt conflicting state AI laws and advance national consistency. It creates an AI Litigation Task Force within the U.S. Department of Justice (DOJ) tasked with identifying and legally challenging state laws that are inconsistent with federal policy, including those argued to interfere with interstate commerce or be preempted by federal standards. The Secretary of Commerce, in coordination with senior AI advisors, is also required to evaluate state AI laws and recommend those that may be referred to the Task Force. Federal agencies such as the Federal Communications Commission (FCC) and Federal Trade Commission (FTC) are instructed to explore national reporting and disclosure standards for AI that could preempt conflicting state regulations. Additionally, EO 14365 calls for legislative recommendations to Congress for a uniform federal AI policy framework that, where appropriate, would supersede state AI laws.

Healthcare implications of EO 14365 flow from its drive toward regulatory uniformity. For healthcare organizations and AI developers, EO 14365 aims to reduce the compliance burden created by disparate state laws governing AI applications in areas such as clinical decision support, patient data use, diagnostics, and patient-facing algorithms. A consolidated federal policy could streamline multi-state deployment of AI-enabled tools, eliminate conflicting legal requirements, and provide clearer guidance for innovation in health tech. At the same time, stakeholders should be aware that states may pursue litigation or challenge the federal preemption strategy, and federal standards are still emerging. As a result, healthcare clients are advised to continue to comply with current state AI laws while closely monitoring developments in federal policy and potential legal challenges to EO 14365's implementation.

Overall, EO 14365 reflects the Administration's effort to balance national AI leadership and innovation with the complexities of the U.S. federal system, creating both opportunities and legal uncertainties for sectors like healthcare that are integrating AI technologies at scale.

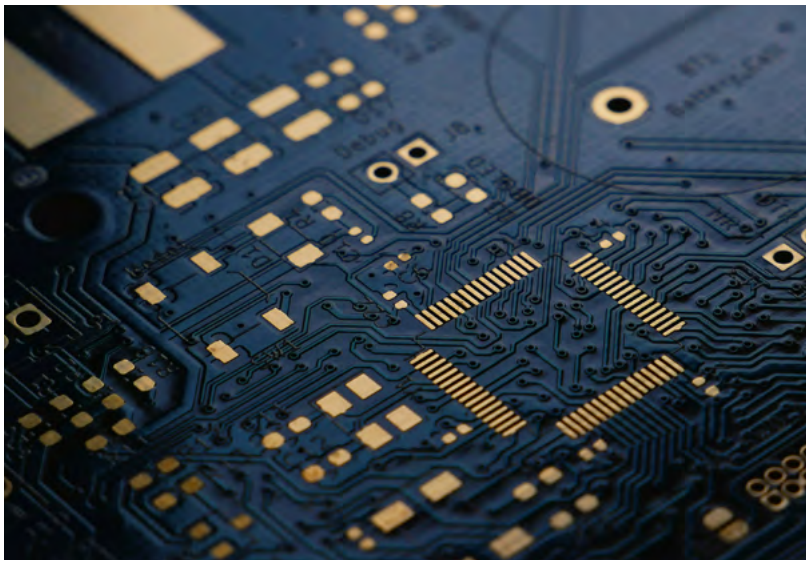

U.S. Federal Level – Agency and Other Guidance on AI in Healthcare

01 FDA's Guidance on AI in Healthcare

At the beginning of 2025, the Food and Drug Administration (FDA) released a few anticipated drug and device guidance documents related to AI. The FDA's January 7, 2025 guidance, "Considerations for the Use of Artificial Intelligence (AI) to Support Regulatory Decision-Making for Drug and Biological Products," marks a pivotal step toward integrating AI into the decision-making process for drugs and biologics. It outlines how AI output data can be incorporated into regulatory submissions for drugs and biologics on safety, effectiveness, or quality. For AI to influence drug approvals or safety decisions, it must be credible, explainable, and reproducible, ensuring that algorithms do not introduce bias or compromise patient safety. As the AI output may be difficult to interpret and explain, the FDA places emphasis on transparency and recommends a risk-based approach that includes human oversight and intervention. For healthcare executives, this means there will be a growing need to allocate resources for technology assessment of AI tools and outputs, while also preparing for potential shifts in the regulatory inquiry during product review.

A second January 2025 guidance document from the FDA, "AI-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations" (AI-DSF Guidance), addresses the unique total product lifecycle challenges of AI-enabled devices. This is only the second guidance issued by the FDA specific to AI-enabled device software functions (AI-DSF). The scope not only includes recommendations on initial submissions, but also design, development, deployment, post-market modifications, performance monitoring, and cybersecurity. The impact is important for both device manufacturers and healthcare users, as the FDA sets expectations for pre-market planning, post-market surveillance, and robust change management systems for a market device segment that is experiencing rapid growth. Of note, the AI-DSF guidance places particular emphasis on maintaining transparency across the entire life span of AI-enabled devices, with a user-centered focus that supports a "human-in-the-loop" approach for the safe and effective use of AI-enabled devices beginning with the initial design.

While the AI-DSF Guidance does address cybersecurity, it mostly refers to the June 2023 guidance document, which has since been superseded by the June 2025 guidance document, "Cybersecurity in Medical Devices: Quality System Considerations and Content of Premarket Submissions." The updated 2025 guidance is generally consistent with the total product lifecycle and transparency approach the FDA takes in the AI-DSF Guidance and now includes specific detailed recommendations for "cyber devices" as defined in Section 524B of the Food, Drug and Cosmetic Act. Specifically, the FDA clarifies that devices that "are or contain software, including software that is firmware or programmable logic," or that are able to connect to the internet, "whether intentionally or unintentionally, through any means" are included in the definition of cyber devices. The FDA recommendations align with the increased cybersecurity risks related to cyber devices, which device manufacturers and healthcare users should also consider with any AI-DSFs.

Finally, in August 2025 the FDA updated its final guidance, "Marketing Submission Recommendations for a Predetermined Change Control Plan for AI-Enabled Device Software Functions," which was originally issued at the end of 2024. The final guidance recognizes the need for an innovative approach for managing device software updates that leverage AI and supersedes and expands upon the 2024 version and provides more details on Predetermined Change Control Plan (PCCP) updated regulatory expectations around the FDA's total product lifecycle approach. The FDA understands that AI-DSFs involve iterative improvements and has adopted an approach that enables more rapid deployment of device improvements while simultaneously ensuring devices are safe and effective. Device manufacturers can submit a PCCP to the FDA for devices that include an AI-DSF, including where the modifications to the AI model are implemented automatically. If the FDA authorizes the PCCP as part of a marketing submission, the device company will be able to modify its AI-DSF in the future without obtaining separate FDA authorization for each subsequent change.

02 HHS AI Strategic Plan

HHS's "Strategic Plan for the Use of Artificial Intelligence in Health, Human Services, and Public Health" was published on January 10, 2025, and describes "HHS's vision and goals for AI in health, human services, and public health." This lengthy and detailed report reinforces HHS's healthcare-specific priorities across four strategic goals: (i) catalyzing health AI innovation and adoption to unlock new ways to improve people's lives; (ii) promoting trustworthy AI development and ethical and responsible use to avoid potential harm; (iii) democratizing AI technologies and resources to promote access; and (iv) cultivating AI-empowered workforces and organization cultures to effectively and safely use AI. The Strategic Plan provides HHS's strategy and detailed action plan for achieving these goals across several domains: (i) medical research and discovery; (ii) medical product development, safety and effectiveness; (iii) healthcare delivery; (iv) human services delivery; (v) public health; (vi) cybersecurity and critical infrastructure protection; and (vii) internal operations. These actions are expected to evolve into federal regulations as the underlying technologies advance and as additional real-world evidence is generated. Future new laws are likely to focus on stronger requirements for data quality, bias mitigation, transparency, explainability, privacy protection, and cybersecurity, as well as oversight and remediation of AI use in the healthcare sector.

03 HHS AI Strategy

Looking ahead, HHS has stated intentions to harmonize its strategic initiatives with broader congressional and FDA actions, setting the stage for a more unified national approach to AI regulation in healthcare. On December 4, 2025, HHS unveiled its "Artificial Intelligence Strategy", with the complex goal to use AI to transform agency operations. Of note is the establishment of a "OneHHS" approach where all divisions of the HHS, including the Centers for Disease Control and Prevention, CMS, FDA, the National Institutes of Health, and others, will streamline AI procurement, development, data sharing, workflows, and cybersecurity enhancements. Some of the main goals that HHS focuses on in the Artificial Intelligence Strategy are to improve research with gold standard science and to modernize healthcare delivery for better outcomes.

04 FTC

In 2025, the FTC continued to sharpen its focus on AI-related enforcement and oversight, applying longstanding consumer protection laws to modern AI technologies rather than seeking new AI-specific statutes. The agency's public guidance and enforcement actions make clear that existing rules against unfair or deceptive practices under Section 5 of the FTC Act apply fully to AI products and services, meaning companies must ensure their AI-related claims are truthful, substantiated, and transparent. Self-reported data use and model capabilities must be accurately described to consumers, or companies risk enforcement action for misleading representations.

A cornerstone of the FTC's more recent approach to AI enforcement is "Operation AI Comply", a law enforcement sweep launched in 2024 to target companies that allegedly use AI "hype" to mislead consumers or engage in unfair schemes. The initiative resulted in complaints and enforcement actions against businesses that advertised exaggerated AI benefits, from purported legal services to tools promising unrealistic income-generation, underscoring the agency's view that deceptive AI marketing can constitute unlawful conduct. The emphasis on truthful communication around AI claims continues to be reiterated by FTC officials as AI technologies become more integrated into consumer products.

Beyond advertising, the FTC's oversight encompasses data practices tied to AI deployment, especially where consumer privacy is concerned. According to FTC guidance, companies that fail to honor commitments about how they collect or use user data, including using it to train or improve AI models without clear and conspicuous disclosure, may be liable under applicable consumer protection laws. Previous FTC actions have even required companies to delete models or algorithms developed using unlawfully obtained data.

In late 2025, the FTC also initiated a broader information gathering inquiry into AI chatbots, requesting data from major platforms about how they test and monitor consumer-facing AI tools for safety risks and negative impacts. This move suggests the agency is expanding its scrutiny beyond claims and data practices to consider how companies measure and mitigate harms associated with advanced AI systems.

05 OCR

In 2025, the HHS's Office for Civil Rights (OCR) continued to intensify enforcement activity related to the privacy and security of health data, including where AI-assisted technologies are used to process protected health information (PHI). While OCR has not issued AI-specific regulations, it has made clear through guidance and enforcement that existing HIPAA Privacy and Security Rules fully apply to emerging technologies. Covered entities and business associates are expected to conduct thorough risk analyses and implement appropriate safeguards before deploying AI tools that access or analyze PHI.

OCR's recent enforcement actions reflect a strong focus on risk analysis, cybersecurity controls, and vendor oversight, all of which are particularly relevant when AI systems are involved in data handling. In 2024 and 2025, OCR finalized numerous settlements and civil monetary penalties tied to failures to safeguard electronic PHI, emphasizing that inadequate security practices, regardless of whether advanced technologies are used, can lead to significant liability. These trends signal heightened expectations for organizations leveraging AI to demonstrate that privacy and security risks are proactively identified and managed.

06 DOJ

In 2025, the DOJ maintained robust enforcement of the False Claims Act (FCA), with continued emphasis on healthcare fraud and emerging scrutiny of how AI-assisted coding and billing systems may contribute to improper claims. DOJ's collaboration with HHS through the DOJ-HHS False Claims Act Working Group underscores the government's intention to leverage data analytics to identify and address potential fraud, waste, and abuse in federal healthcare programs, including issues like manipulation of electronic health records (EHRs) and billing practices, which may be amplified by automated systems.

A notable trend has been the expansion of FCA enforcement into areas where automated or AI-driven tools influence billing accuracy. In one high-profile example, a healthcare provider agreed to a $23 million settlement after its automated coding system generated billing that allegedly failed to align with official coding requirements, leading to overpayments from Medicare and Medicaid. This case illustrates how errors emerging from automated or algorithmic processes can trigger FCA liability when claims are submitted to federal programs, even if the underlying intent was not fraudulent.

Law firms and compliance experts warn that AI systems can systemize improper coding at scale when models are trained on biased or outdated data or when human oversight is insufficient, compounding risk across thousands of claims. Because FCA liability can be based on "reckless disregard" for accuracy, a relatively low intent threshold, hospitals and billing entities must ensure that any AI-assisted tools reflect current coding guidelines and robust compliance checks.

Looking forward, DOJ's enforcement landscape suggests that AI-related billing and coding errors will remain a focal point for FCA investigations. Healthcare organizations should prioritize proactive compliance by integrating strong oversight, documentation, and validation frameworks for any AI tools used in claims generation or processing. Continuous monitoring and human review will be critical to mitigate FCA exposure as federal agencies expand their use of data analytics and predictive tools to detect irregularities before they evolve into costly enforcement actions.

07 CMS

In 2025, federal oversight of utilization management and prior authorization at CMS increasingly intersected with AI and advanced technology, marking a notable shift in how Medicare evaluates coverage decisions. On January 1, 2026, CMS launched a landmark pilot under its Wasteful and Inappropriate Service Reduction (WISeR) Model, in six states to test AI-enabled review of prior authorization requests for select Medicare services that CMS identifies as vulnerable to fraud, waste, and abuse. While licensed clinicians will ultimately make final determinations, AI and machine learning tools will support the review process to help ensure timelier and more consistent coverage assessments.

The WISeR pilot reflects CMS's broader effort to modernize utilization management by incorporating technology that can sift through clinical information and flag requests that merit closer scrutiny. The model will apply to services such as skin and tissue substitutes, nerve stimulation devices, and certain orthopedic procedures, all subject to enhanced prior authorization criteria under the demonstration. Participating technology vendors are expected to leverage AI to accelerate decision workflows while maintaining clinical standards for medical necessity.

These developments have generated a mix of responses from the health sector. Some providers and clinicians express optimism that AI-driven tools could reduce administrative burdens and streamline workflows, while others raise concerns about potential delays in care and the transparency of automated components in authorization decisions. A national survey conducted in 2025 found widespread provider confidence in AI's ability to support prior authorization processes, even as formal regulatory guardrails continue to evolve.

CMS's direction on AI in prior authorization is part of a broader emphasis on utilization management reform and oversight, as the agency also works on improving interoperability and data exchange standards that support electronic prior authorization across payers. As the WISeR model rolls out and additional guidance emerges, stakeholders should monitor how AI tools are deployed, the results of early implementations, and how CMS balances efficiency with patient access and clinical integrity.

08 House Task Force on AI – Guiding Principles for Congress

The "Bipartisan House Task Force on Artificial Intelligence" was released on December 17, 2024, by the House Task Force on Artificial Intelligence, and is important to understand the congressional position on the deployment of AI in healthcare in 2025. Per Chairman Jay Obernolte (R-Ca.), "[t]he report details a roadmap for Congress to follow to both safeguard consumers and foster continued U.S. investment and innovation in AI." The report includes a chapter dedicated to healthcare and has specific recommendations for policymakers. Congress recommends laws and regulations concerning AI technologies in healthcare, including to ensure that they are continually and sufficiently monitored for safety and efficacy. The report also expresses concern that equitable access is provided to AI-assisted technologies. Looking towards the future, healthcare businesses should anticipate further legislative initiatives and regulations, which could include mandatory disclosure requirements for algorithmic decision-making in patient care, and new liability frameworks for adverse events caused by AI, with a focus on transparency and accountability in the use of AI.

State Level

Texas Responsible AI Governance Act

Texas enacted the Texas Responsible Artificial Intelligence Act (TRAIGA) in 2025, which became effective on January 1, 2026. TRAIGA applies to any party who (i) promotes, advertises, or conducts business in Texas; (ii) produces a product or service used by residents of Texas; or (iii) develops or deploys an AI system (an AI System) in Texas. Significantly, a "consumer" for purposes of the Act is limited to residents of Texas acting only in an individual or household context (i.e., excluding commercial activities).

Among regulated parties, TRAIGA assigns duties based on the party's role, which includes both deployers (i.e., a party who deploys an AI System for use in Texas), as well as developers (i.e., a party who develops an AI System that is offered, sold, leased, given, or otherwise provided in Texas). TRAIGA imposes a number of duties and prohibitions on regulated parties, including at a high level:

1. Developers and deployers are prohibited from intentionally using AI Systems to promote or encourage self-harm, violence, or criminal acts.
2. AI Systems must not be developed or used with the intent to infringe, limit, or otherwise impair consumer rights protected by the U.S. Constitution or with the intent to unlawfully discriminate against consumers of a protected class in violation of federal or state law. The existence of disparate impact alone does not demonstrate intent to discriminate.
3. Regulated parties may not develop or deploy AI Systems with the sole intent of producing or distributing AI-generated child pornography or "deepfake" sexual content depicting minors.

Specific to the healthcare setting, TRAIGA requires that providers of healthcare services disclose to patients that an AI System is being used to provide such service. In particular, TRAIGA requires that such notice be furnished no later than the date the service is first provided, except in the case of an emergency.

TRAIGA also imposes a number of obligations on governmental entities leveraging AI Systems, such as prohibiting use of AI Systems to categorize individuals based on their social behavior or characteristics for the purpose of assigning "social scores" that could result in detrimental or unfavorable treatment or to uniquely identify a consumer using biometric data sourced from publicly available information without that consumer's consent, among other activities.

Apart from limiting use of AI Systems, TRAIGA also strives to promote innovation by establishing a "sandbox" program permitting qualifying participants to test AI Systems without obtaining the licenses or permits that would otherwise be required. Aspiring participants must seek approval from the Texas Department of Information Resources and any applicable state agency before testing an AI System within the sandbox program.

TRAIGA also establishes the Texas Artificial Intelligence Council, which is charged with (i) overseeing that AI Systems are used in an ethical manner and developed in the public's interest; (ii) ensuring that AI Systems used in Texas do not harm public safety or undermine individual freedoms by making recommendations to the legislature; and (iii) making recommendations to the applicable state agencies addressing use of AI Systems with an aim to improve efficiency and effectiveness, among other activities. The Texas Artificial Intelligence Council will be a part of the Texas Department of Information Resources.

In terms of enforcement, TRAIGA empowers the Texas Attorney General with the authority to enforce TRAIGA and expressly nullifies any city or county ordinance regulating AI. Where there is a determination that a party has violated TRAIGA, the Texas Attorney General must notify the pertinent party in writing of such determination, must identify the specific provisions which appear to have been violated, and must provide the party with an opportunity to cure the identified violation within sixty days following the notice. Notably, TRAIGA expressly states that it does not create a private right of action for individuals. Further, TRAIGA arms the Texas Attorney General with the power to assess civil penalties, to pursue injunctive relief, as well as to seek recovery of attorney's fees, court costs, and investigative expenses. TRAIGA authorizes civil penalties ranging from $10,000 to $12,000 for curable violations, $80,000 to $200,000 for non-curable violations, and $2,000 to $40,000 per day for continued violations. Separately, TRAIGA also permits state agencies to impose sanctions against parties licensed, registered, or certified by such agency. It is important to note that TRAIGA prohibits the Texas Attorney General from pursuing "charges against a participant within the sandbox program for violation that occurs during the testing period." Similarly, TRAIGA also prohibits state agencies from filing or pursuing punitive action against participants.

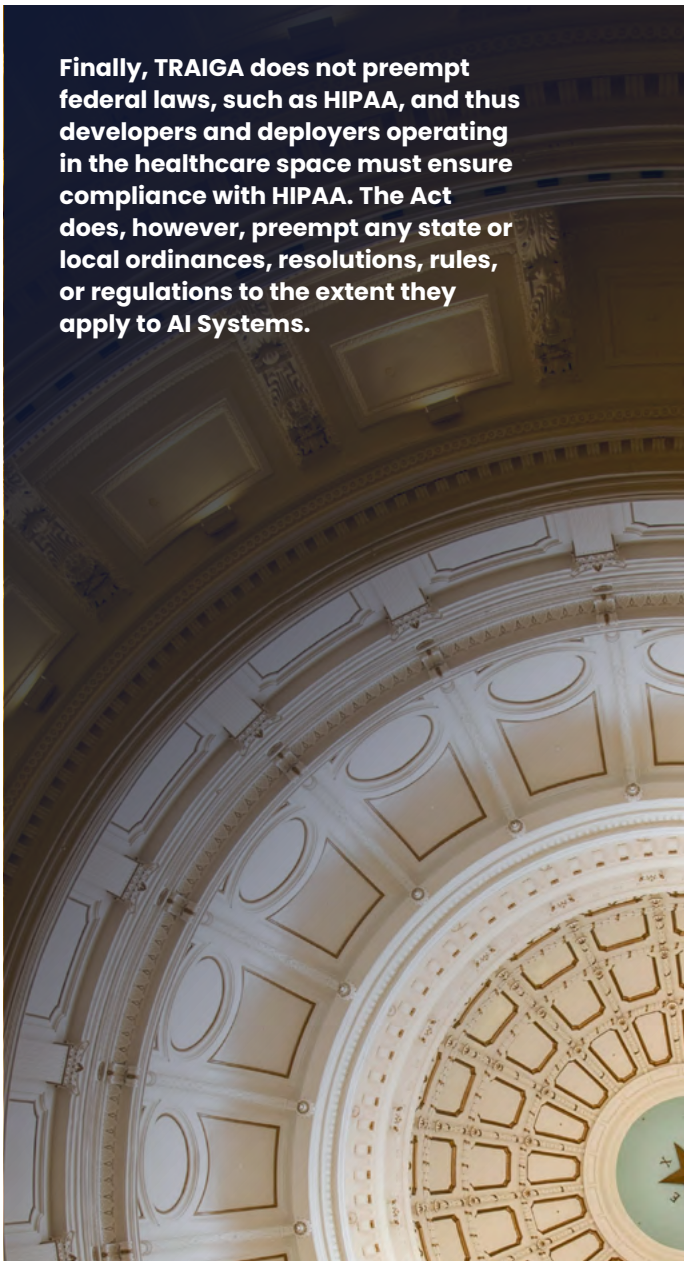

Finally, TRAIGA does not preempt federal laws, such as HIPAA, and thus developers and deployers operating in the healthcare space must ensure compliance with HIPAA. The Act does, however, preempt any state or local ordinances, resolutions, rules, or regulations to the extent they apply to AI Systems.

Utah Artificial Intelligence Act and AI Amendments

In 2025, Utah continued to be a leader in AI regulation, enacting a series of AI laws including HB 452, SB 226, and SB 332. These three laws amend Utah's Artificial Intelligence Policy Act (UAIPA), a first-of-its-kind consumer protection law, enacted in 2024. Notably, HB 452, SB 226, and SB 332 include key provisions impacting the healthcare industry.

HB 452 creates disclosure requirements, advertising restrictions, and privacy protections for the use of mental health chatbots. In particular, the law protects users' individually identifiable health information from being sold to or shared with third parties without authorization. It also creates an affirmative defense for suppliers of mental health chatbots who implement and maintain policies for these chatbots that, among other requirements, clearly state their purposes, abilities, and limitations and ensure the chatbots are used in a manner consistent with clinical best practices.

SB 226 limits certain UAIPA requirements to disclose that an interaction is with AI. Sellers using generative AI in a consumer transaction must now disclose such use in response to "clear and unambiguous requests" to determine whether an interaction is with a human or AI made during a "high-risk" interaction, rather than any request in any generative AI interaction. In addition, providers of regulated services must "prominently disclose" when an interaction is with AI. The amendment creates a safe harbor for clear and conspicuous disclosures, establishing incentives for suppliers to be transparent about their use of AI tools.

Finally, SB 332 extends UAIPA's repeal date from May 1, 2025 to July 1, 2027. These laws emphasize Utah's continued focus on transparency in the use of generative AI, the importance of consumer choice, and the state's attempt to strike a balance between protecting consumers and avoiding placing undue burdens on the use of AI in commercial settings.

California AB 3030

California's AB 3030 took effect on January 1, 2025, requiring health facilities, clinics, physicians' offices, and group practice offices to provide disclaimers when generative AI is used to generate patient communications. The law also requires these entities to provide clear instructions to patients on how to reach a human provider or employee. AB 3030 extends not only to the first AI-generated communication with a patient regarding clinical information, but also to all AI-generated communications. An exception applies if the communication has been read and reviewed by a licensed or certified healthcare provider, and the disclosure requirements do not apply to non-clinical information, such as scheduling or billing matters.

Beyond AB 3030 taking effect in 2025, California also adopted other AI laws, a handful of which impact the healthcare industry, including SB 243, Companion Chatbots Safety Act, and AB 489, Health Advice from AI Act. AB 489, which was sponsored by the California Medical Association, prohibits developers or deployers of AI tools from falsely indicating that it possesses a healthcare license or that the advice is being provided by someone with a license. SB234 requires protocols to limit chatbots from engaging in conversations about suicidal ideation, self-harm or sexually explicit content, which would have implications for any healthcare related chatbots or chatbots used in the behavioral health space.

Notably, California also adopted SB53, the Transparency in Frontier Artificial Intelligence Act (TFAIA), which applies to limited developers who either fall into the "frontier developer" or "large frontier developer" category and took effect on January 1, 2026. Frontier developers are those that have "trained, or initiated the training of, a frontier model," which is a large foundation model trained using a high level of computing power. A large frontier developer is a developer who together with affiliates collectively has annual gross revenues in excess of $500 million in the prior year. While very few in developers will fall into this category due to the revenue and computing power thresholds, TFAIA still merits mention given that it requires much transparency and reporting of these developers and could serve as important precedent for transparency steps in other states as they develop their own AI legislation.

Broader Trends in the State Legislative Landscape

2025 saw a notable increase in state legislative activity regulating the use of AI in the healthcare industry. Four areas with implications for the healthcare industry were at the forefront:

1. **Algorithmic Impact Assessments:** Though the use of AI has become routine in many industries, concerns arise when AI tools fail to evolve in tandem with their users' goals. Algorithmic impact assessments (AIAs), structured evaluations designed to identify and mitigate risks from using automated systems, help allay these concerns. These risks include security, privacy, and non-discrimination concerns. Several state legislatures introduced bills this year requiring developers of "high-risk" AI systems, including those substantially contributing to decisions having a material effect on access to healthcare, to implement AIAs.
2. **Automated Decision-Making Restrictions:** In addition to requiring AIAs, a number of states introduced bills restricting the degree to which AI could be used to make automated decisions without human intervention. Many of these bills, such as Florida's HB 527 and Pennsylvania's HB 1925, mandate that all health insurance coverage decisions by AI resulting in claim denials be reviewed by a human being.
3. **Consumer Transparency Rights:** Several states also introduced laws governing businesses' obligations to disclose when a consumer is interacting with AI rather than with a human being. Most of these proposed bills bolster consumer protections, including through mandatory opt-out provisions; however, Utah's SB 226 pared back some of the AI disclosure requirements the state previously implemented in 2024.
4. **Mental Health Chatbots:** Six states enacted legislation regulating the use of AI chatbots in 2025, including Utah's HB 452, discussed above. Among other provisions, these laws include disclosure requirements, enhanced protections for minors, and requirements for prompt referral to crisis support providers. California's SB 243, for example, aims to protect minors from emotional manipulation and unsafe interaction with AI chatbots through notification requirements and prevention protocols. Notably, Governor Newsom of California vetoed AB 1064, which would have prohibited most uses of AI chatbots by minors, including for mental health therapy.

States to Monitor

While many states introduced or enacted regulations on the use of AI in 2025, California, Colorado, and New York stand out as leaders in the space. These three states have had elevated levels of legislative activity, and several laws from their legislatures will be coming into effect in 2026, including the Colorado Artificial Intelligence Act, the first comprehensive U.S. law regulating high-risk AI intelligence systems. Providers interested in harnessing AI tools in these states should carefully monitor new and evolving obligations around the use of AI.

Fragmentation Risk for Multi-State Providers

2025 presented robust activity with some states introducing a variety of laws regulating AI, while others remained quiet with effectively no or limited legislative activity. In the absence of a governing AI federal law to date, this patchwork of state laws complicates providers' legal obligations surrounding the use of AI, particularly for those with patients or consumers in multiple states. These laws are often sector-specific and vary widely in what they require. Accordingly, until a federal regulatory framework is in place, providers operating in multiple states must carefully navigate this fragmented and expanding legal landscape.

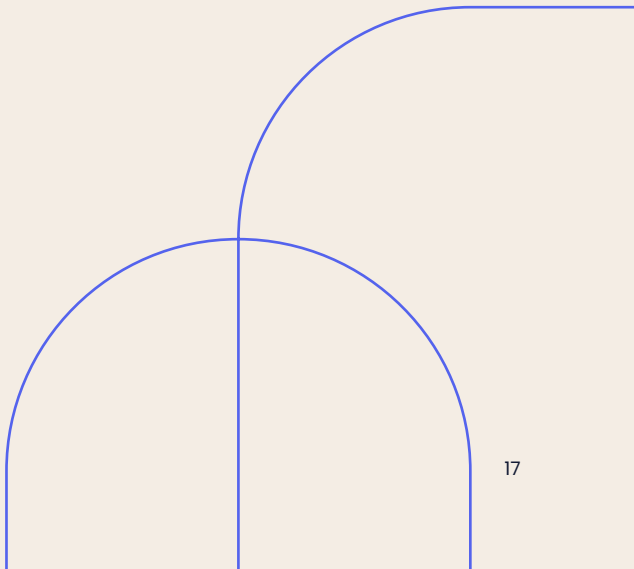

Other Industry Considerations

Coalition for Health AI (CHAI)

On December 24, 2025, CHAI released its new best practice guides and testing/evaluation frameworks for four additional use cases, "General Health Advice Chatbot, Prior Authorization, Electronic Health Record Information Retrieval, and Clinical Decision Support." These resources offer actionable recommendations for developers and implementers to evaluate, deploy, and continually monitor AI solutions in certain areas of healthcare. Specifically, the case-specific best practice guides offer developers and implementers practical insight into the current challenges associated with applying particular AI tools in healthcare, emphasizing the importance of integrating responsible AI principles throughout development and deployment. The corresponding testing/evaluation frameworks provide structured options for organizations to operationalize the evaluation and ongoing monitoring of responsible AI practices as they pertain to each specific use case. Recognizing the dynamic nature of health AI, CHAI maintains these guides and frameworks as living documents to help healthcare organizations safely scale responsible AI. These efforts underscore CHAI's ongoing mission, as described in its 2024 assurance standards framework, to equip organizations to address evolving AI challenges while promoting ethical, effective, and patient-centered care.

Other Industry-Led Risk Frameworks

In May 2025, the National Academy of Medicine (NAM) released the special publication, "An Artificial Intelligence Code of Conduct for Health and Medicine: Essential Guidance for Aligned Action," which provides an AI Code of Conduct (the Code) framework for responsible and coordinated development and use of AI across technologies, systems, and practices. The Code serves as a reference for organizations and developers as they design, deploy, and oversee AI technologies in health and medicine. The Code includes six central commitments: advancing health equity, maintaining human oversight and interests, protecting privacy and security, ensuring safety and effectiveness, supporting transparency and accountability, and promoting responsible AI governance. According to NAM, the Code provides a common foundation to guide policy development, inform best practices, and support collaboration among stakeholders, facilitating consistent and coordinated integration of AI technologies throughout the healthcare sector.

On September 17, 2025, the Joint Commission in partnership with CHAI released comprehensive, high-level guidance, "Responsible Use of AI in Healthcare," to support responsible adoption of AI use in healthcare organizations.

The guidance is designed to facilitate internal governance and the safe deployment of AI tools within healthcare organizations, focusing on key areas including patient safety, clinical outcomes, data protection, patient privacy, quality monitoring, reporting, and operational efficiency.

The guide provides foundational recommendations for AI policies and governance, ongoing monitoring and reporting of AI safety events, and the protection of patient data, clarifying that its priority is safe operational use of AI tools rather than their technical development.

On November 11, 2025, Health Sector Coordinating Council (HSCC) Cybersecurity Working Group (CWG) published a preview of upcoming AI cybersecurity guidance, providing healthcare organizations insight into recommended practices for securing AI technologies and protecting sensitive information as digital threats evolve. As described by HSCC, this preview outlines key areas of vulnerability and suggests practical steps to manage AI-related cyber risks, enhance organizational readiness, and respond to emerging challenges. The guidance documents will be the first product of the HSCC CWG's AI Cybersecurity Task Group, which was formed in October 2024 and comprises 115 healthcare organizations. The final version of the guidance is expected to be released in 2026, with more detailed resources to support healthcare entities in safeguarding their operations with respect to cyber safety.

Core Healthcare AI Risk Domains That Matured in 2025

In 2025, the use of AI in healthcare continued to expand across both clinical and payer environments. As adoption increased, legal and regulatory analysis increasingly focused on how existing liability, reimbursement, and oversight frameworks applied to AI-enabled tools in practice. Developments during the year reflected closer examination of accountability, oversight, and risk allocation within established regulatory and legal structures.

Clinical Liability and Medical Malpractice

Throughout 2025, attention remained focused on the distinction between AI tools deployed as clinical decision support and those used in a more autonomous diagnostic or treatment capacity. While many AI systems were described as supportive of clinician decision-making, their integration into clinical workflows often resulted in AI-generated outputs playing a significant role in diagnoses, treatment planning, and patient prioritization.

Regulatory guidance and industry practice continued to emphasize the importance of maintaining human oversight over AI-enabled clinical functions. In evaluating risk, the manner in which AI tools were operationalized, such as the degree of clinician review, the availability of override mechanisms, and the documentation of human involvement, became increasingly relevant.

Additionally, the growing presence of AI in clinical settings prompted renewed focus on how the standard of care applies when AI tools are used as part of healthcare delivery.

In 2025, scrutiny centered on the role of AI-generated information within clinical decision-making rather than on the technology itself. Providers remained responsible for exercising independent professional judgment, and AI outputs were generally treated as one of several inputs informing care decisions. Documentation of clinical reasoning, including how AI recommendations were considered or disregarded, remained an important component of risk management.

Relatedly, increased awareness of AI limitations heightened the expectation that clinicians remain attentive to inconsistencies between AI outputs and patient-specific information. The use of AI did not alter existing professional duties but informed how those duties were carried out in AI-enabled environments.

Product Liability Exposure

The broader deployment of AI in clinical care also highlighted product liability considerations involving both developers and healthcare organizations.

With respect to AI vendors, developers of clinical AI systems continued to face scrutiny related to model design, training data, performance representations, and disclosures regarding limitations and appropriate use. FDA guidance addressing AI-enabled device software functions, including predetermined change control plans, provided additional clarity regarding regulatory expectations for lifecycle management and documentation.

With respect to integrated health systems, health systems implementing AI tools also encountered liability considerations, particularly where systems customized, configured, or expanded the use of AI beyond its original design.

Internal governance practices, validation efforts, and oversight mechanisms were relevant in assessing responsibility for AI-enabled outcomes within clinical operations.

Reimbursement and Payer Use of AI

In 2025, payers continued to rely on AI across claims adjudication, prior authorization, and fraud, waste, and abuse detection. These systems were used to automate administrative processes, manage utilization, and analyze large volumes of data. As AI tools became embedded in reimbursement workflows, their role in influencing coverage determinations became more visible to regulators, providers, and patients.

Regulatory attention during 2025 focused on the use of AI in coverage determinations, particularly in connection with automated denials and potential discriminatory impact. Federal and state authorities reiterated that AI tools may be used to support administrative efficiency but may not replace individualized assessment where medical necessity determinations are required.

Concerns regarding transparency and bias also remained prominent. Regulators emphasized the importance of human review, documentation of decision-making processes, and monitoring of AI performance over time. Guidance and oversight activity reinforced existing expectations that payers deploying AI systems remain accountable for ensuring compliance with consumer protection, nondiscrimination, and due process requirements.

Litigation, Investigations, and Case Law Watch

Emerging theories of liability included algorithmic negligence, failure to warn, and inadequate human oversight.

Under theories of algorithmic negligence, courts and regulators are beginning to articulate theories addressing liability for AI-driven harm. Traditional frameworks, rooted in human fault, often fail to properly account for the distributed and opaque nature of "black box" algorithms, where errors may originate from developers, clinicians, or deploying institutions at varying points in the system lifecycle. The emerging doctrine of algorithmic negligence posits that when anyone responsible for developing or deploying an AI system fails to exercise reasonable care to prevent foreseeable risks, they may be deemed legally negligent, akin to the liability incurred from manufacturing defective products.

Defining what constitutes "reasonable" development of, or supervision over, autonomous systems present unique challenges, especially when these systems act at speeds and scales beyond human oversight. Courts are increasingly focused on duties relating to data quality, testing and validation, ongoing monitoring, and transparency about known limitations, particularly the need for robust human oversight when failures could result in significant harm. While algorithmic negligence is not yet an accepted cause of action in courts and harm is often still viewed through the lens of conventional malpractice, recent cases signal an evolving willingness to adapt traditional liability concepts to encompass AI-specific risks.

One such adaptation is the failure to warn, treating AI tools as products subject to product liability principles. Under this theory, developers and deployers may be held accountable when users are not adequately warned of non-obvious risks such as bias, limitations in model accuracy, over-reliance on automated outputs, or privacy concerns. Liability may arise if foreseeable risks are not disclosed, leading to user harm, such as the failure to advise that an AI model's recommendations may be unreliable for certain patient populations.

Compounding these risks is the issue of inadequate human oversight, often manifested as "automation bias." Clinicians may defer to algorithmic outputs, particularly when marketed as evidence-based and clinically validated, without independent means to verify those outputs due to the opacity of many AI systems. Despite such advancements, the practice of medicine remains the province of licensed professionals, who must retain ultimate authority over patient care decisions. As standards evolve, industry efforts to establish evaluation tools and guidelines will play a critical role in mitigating liability and shaping accountability for emerging AI applications in healthcare.

Whistleblower risk is increasingly tied to AI-assisted billing and automated upcoding.

As hospitals and physician groups increasingly deploy AI-assisted billing tools to streamline coding and claims submission, they are also exposing themselves to heightened FCA risk. These systems can perpetuate systematic errors at scale if not carefully monitored, and internal staff are often the first to identify patterns of improper billing driven by algorithmic logic rather than clinical documentation. Federal enforcement trends continue to show that DOJ and relators' counsel are attuned to automated billing processes, making robust oversight, audit trails, and human review essential to mitigate whistleblower exposure.

AI-enabled documentation tools designed to improve specificity can also create pressure toward higher-acuity coding. If models are trained on historical data skewed toward revenue maximization, the result may be subtle but pervasive upcoding disconnected from actual clinical complexity. Regulators are scrutinizing the role of AI recommendations in coding decisions, and 2025 saw more focus on guardrails, such as transparency requirements, clinician attestation, and periodic model validation, to reduce the risk that automation translates into fraudulent billing practices.

Going forward, a wave of class actions and regulatory test cases is expected.

The rapid adoption of clinical and administrative AI has triggered a new wave of class action litigation, particularly where automated tools influence reimbursement, utilization management, or patient access to care. Plaintiffs' theories in 2025 increasingly focused on systemic harms caused by model-driven decisions, arguing that uniform, algorithmic workflows create common injury ripe for class treatment. Courts have shown growing willingness to allow discovery into how models are trained, validated, and governed, making transparency around AI development and deployment a critical defense consideration for health systems, payers, and vendors.

In the past year, regulators have used targeted enforcement and administrative actions to build precedent. Investigations into algorithmic bias, documentation integrity, and payment accuracy serve as early "test cases" to clarify how existing fraud, consumer protection, and health privacy rules apply to machine-generated outputs. Even in the absence of comprehensive AI-specific statutes, agencies emphasized that long-standing standards, including accuracy, fairness, and truthful claims, remain fully enforceable against automated systems, signaling that compliance expectations evolve faster than legislation.

Strategic Governance Lessons from 2025

Throughout 2025, advancements in AI technologies continued. AI's role in healthcare rapidly expanded, including, but not limited to, a wider range of applications, transformation of healthcare operations, and enhanced patient experiences. This development in AI prompted a fundamental shift in strategic governance models. The following lessons emerged as critical for effective governance:

AI Governance Shift from IT to Enterprise Risk Issue

Since the first iterations of regulation around AI emerged, the regulatory landscape has been a moving target. AI oversight has evolved beyond a purely technical or IT concern. AI adoption is accelerating across enterprises. Enterprise level AI risk is the potential of AI to cause losses for an organization and can originate from various sources. In 2025, leading organizations recognized AI as an enterprise-wide risk management issue, integrating AI considerations into broader risk frameworks. This shift reflected the growing realization that AI impacts compliance, reputation, operational resilience, and business continuity, and thus must be addressed at the highest levels of corporate governance.

Importance of Board Oversight, Cross-Functional AI Committees and Model Inventory and Lifecycle Tracking

As healthcare AI solutions rapidly evolve, organizations must navigate a complex landscape filled with promise and risk. To address the surge of AI solutions and models, it is imperative for all healthcare organizations to build a resilient AI governance program. There was an increased demand for meaningful board engagement in AI governance. Boards transitioned from passive oversight and assumed a more active role in setting AI risk appetites, approving policies, and monitoring organizational AI deployments. Successful AI implementation and integration require collaboration among compliance, IT, legal, and clinical teams. Part of this collaboration includes the creation of multidisciplinary cross-functional AI governance committees to oversee implementation, approval processes, and ongoing oversight. AI committees became best practice, ensuring diverse perspectives and alignment with overall corporate strategy. Collaboration must be proactive, continuous, and integrated into the AI lifecycle. Additionally, organizations prioritized the creation and maintenance of robust AI model inventories and lifecycle tracking processes, enabling transparency, accountability, and ready responses to regulatory inquiries or incidents.

Vendor Contract Evolution

Vendor relationships were also re-examined through a strategic governance lens. With such a wide array of applications, it is crucial to understand the specific AI service being offered when negotiating a vendor contract. AI arrangements are complex, and it is important to consider pre-negotiation strategies before entering into a vendor contract. Certain considerations to contemplate include: (i) evaluating the contract within an AI governance scope; (ii) engaging stakeholders; (iii) considering AI-specific contracts; (iv) assessing upstream contract requirements; and (v) performing vendor due diligence. After a deep dive into these considerations surrounding AI use, there are key contracting provisions to consider. Vendor contracts increasingly include specific provisions addressing ownership and licensing, indemnification, AI risk, accountability, transparency, and ongoing monitoring. Other requirements may include disclosure of AI models used, maintaining detailed model documentation, and agreeing to audit and compliance requirements throughout the AI model lifecycle. This evolution reflects the importance of third-party risk management with respect to AI. Finally, data rights remain a top priority for vendors who are increasingly experiencing pushback from customers who are developing a more sophisticated understanding of how data may be used in compliance with law and the significant limitations on the use and disclosure of PHI subject to HIPAA.

The key lessons of 2025 highlight the urgent need for organizations to treat AI governance as a core element of strategic risk management, driven by engaged boards, empowered cross-functional teams, and robust contractual frameworks. Positioning AI oversight at the center of enterprise governance will remain essential for responsible, resilient growth in the years ahead.

Forecast for 2026

Regulatory Outlook

In light of the administration and industry's push for AI innovation and adoption, we anticipate a wave of regulatory developments in AI at all levels. Agencies are expected to continue utilizing AI tools in their own operations and encouraging the adoption of those tools in the entities, companies, and industries they regulate. Given the administration's deregulatory agenda, regulations that limit or unduly obstruct AI use in healthcare may be limited while regulatory developments that enhance adoption or provide agencies with enforcement discretion for AI-related governance may be on the rise. The regulatory landscape will continue to rapidly evolve, driven by continual change.

Operational Reality for Providers and Payers

In 2026, we expect that AI in healthcare will likely no longer be treated as an emerging innovation layered onto existing systems. It is becoming operationally indistinguishable from other forms of mission-critical infrastructure, bringing with it heightened compliance exposure, financial risk, and governance obligations for both providers and payers.

AI as Mission-Critical Infrastructure and Core Compliance Risk

Clinical decision support, utilization management, claims processing, fraud detection, and population health analytics are now deeply reliant on AI-enabled systems. As these technologies become embedded in core clinical and financial operations, failures such as system outages, performance degradation, or vendor disruption can produce impacts comparable to EHR downtime or revenue cycle interruptions. In 2026, we expect AI resilience, encompassing redundancy, availability, continuous monitoring, and incident response, to likely be treated as baseline operational infrastructure, not a discretionary technical feature.

In parallel, AI is emerging as a primary compliance risk surface. We expect providers and payers to continue to face overlapping and sometimes conflicting obligations across fraud and abuse enforcement, nondiscrimination, medical necessity, privacy, and consumer protection frameworks. Risks stemming from opaque algorithms, model drift, biased outputs, or undocumented model changes cannot be effectively managed through traditional compliance controls alone. As a result, we expect AI governance failures to increasingly be understood not as IT shortcomings, but as compliance failures carrying material legal, regulatory, and reputational consequences.

Rising Cost of AI Risk Management

As AI becomes embedded in core clinical and administrative operations, the cost of managing AI risk is rising, particularly through increased cyber insurance scrutiny and expanding regulatory reporting obligations.

Insurers are recalibrating underwriting to account for AI-specific threats such as model manipulation, data poisoning, and downstream harms from automated decision-making. In 2026, we expect organizations deploying AI at scale, especially in high-risk use cases, to likely face higher premiums, narrower coverage, and enhanced disclosure requirements tied to AI governance maturity, vendor dependencies, and incident response readiness. As noted in *Forbes*, *How Agentic AI Is Reshaping Cyber Risk And Challenging The Insurance Model*, inadequate AI governance can result in "higher premium[s], restricted coverage, sublimit[s] on autonomous action losses, or even declined renewals," while organizations demonstrating visibility, auditability, and human override may be viewed as more insurable. Industry analyses, including the Institute for Law and AI, further observe that insurers struggle to quantify AI risk due to limited loss data and immature risk metrics, driving more conservative underwriting and broader exclusions.

At the same time, regulatory expectations around AI transparency and accountability are translating into concrete documentation and reporting burdens. Although the federal government continues to assess how best to regulate AI, and the ultimate treatment or preemption of state AI laws remains uncertain, state-level requirements are presently in effect and actively governing regulated entities. As a result, we expect providers and payers to incur increasing costs to inventory AI systems, document model purpose and logic, maintain audit trails, and respond to regulator or payer inquiries regarding AI-driven decisions. These obligations increasingly resemble established compliance regimes for clinical quality or financial reporting. State laws illustrate this increasing reporting obligation: Colorado requires documentation and risk management practices for high-risk AI systems to prevent algorithmic discrimination; Maryland mandates quarterly reporting on AI use in adverse coverage or payment determinations; and California requires public disclosures and incident reporting for certain large-scale or high-impact AI systems.

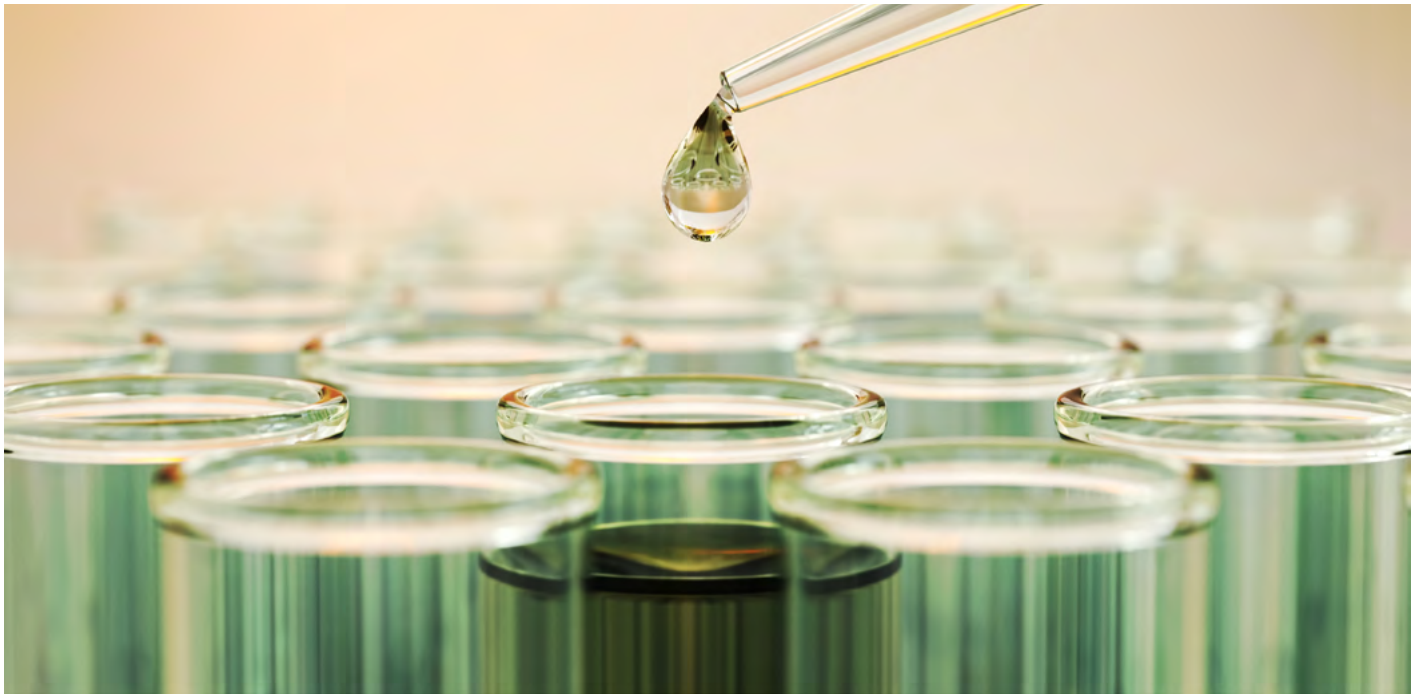

Practical Takeaways for Providers and Payers in 2026

As AI becomes operationally indispensable, healthcare organizations should recalibrate governance, budgeting, and risk management strategies accordingly. Key takeaways to carry into 2026 include:

- **Treat AI as regulated infrastructure, not a pilot program.** AI systems that influence clinical, coverage, payment, or eligibility decisions should be governed with the same rigor as EHRs, billing platforms, or medical devices. This includes formal ownership, documented escalation pathways, uptime and incident response planning, and board-level visibility into AI risk.
- **Integrate legal counsel early into AI lifecycle decisions.** Legal review should occur at multiple points: AI use-case selection, vendor contracting, data sourcing, deployment, monitoring, and incident response. Counsel can help assess whether an AI system triggers healthcare fraud and abuse risk, nondiscrimination concerns, privacy obligations, or emerging AI-specific regulatory requirements – before those risks crystallize into enforcement actions.
- **Reevaluate cyber insurance coverage in light of AI exposure.** Providers and payers should work with risk advisors and legal counsel to understand whether existing cyber and professional liability policies adequately cover AI-related harms, including algorithmic errors, downstream patient impact, or regulatory investigations tied to automated decision-making.
- **Prepare for increased AI-related regulatory inquiries and audits.** Organizations should maintain a current inventory of AI systems, clearly defined intended uses, and contemporaneous documentation explaining how AI outputs are reviewed, overridden, or validated by humans. Legal counsel can help ensure that documentation is aligned with enforcement expectations and does not inadvertently create new liability.
- **Avoid over-reliance on vendor assurances.** Contractual representations regarding AI performance, compliance, and security are increasingly scrutinized after adverse events. Legal counsel should be involved in negotiating AI-specific provisions addressing audit rights, model updates, validation support, indemnification, and regulatory cooperation.
- **Ensure AI governance is cross-functional and legally informed.** Effective AI oversight in 2026 requires collaboration across compliance, clinical leadership, IT, security, risk management, and legal teams. Isolated governance, particularly governance driven solely by technical stakeholders, will be insufficient as AI risk becomes a core enterprise concern.

Navigating AI, Compliance, and Data Privacy in Healthcare

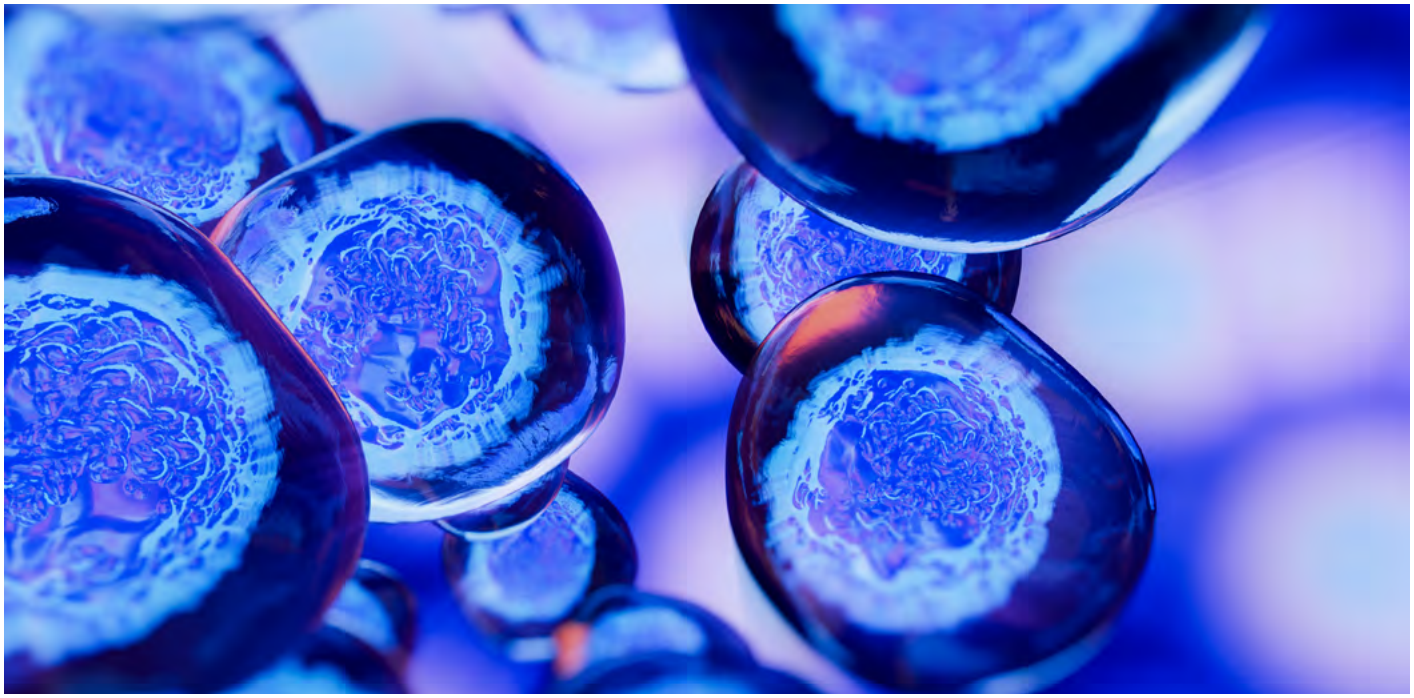

As artificial intelligence transforms healthcare, organizations are racing to integrate new tools—but with innovation comes complexity. In this series by the American Health Information Management Association® (AHIMA), Sheppard outlines how hospitals, health systems, and technology companies can adopt AI responsibly by understanding the legal, ethical, and privacy implications behind every implementation.

Without proper legal review and governance, healthcare organizations risk exposing sensitive patient data or allowing vendors to commercialize trained algorithms using protected health information without the necessary patient authorization.

By reviewing vendor contracts, establishing clear guardrails, and ensuring "humans remain in the loop," healthcare leaders can embrace AI with confidence. The goal isn't to slow innovation, but to make it sustainable — protecting patients, providers, and data integrity while enabling technology to improve outcomes and reduce administrative burden.

View Sheppard's videos, and explore the full AHIMA series, here.

Sheppard

Sheppard Healthy AI

Sheppard Healthy AI recognizes the unique need for healthcare clients to have industry expertise around the developments relating to AI, healthcare and the law. The sensitivity of health information makes this area particularly challenging in its implication of data privacy, HIPAA, FDA issues, and others. Sheppard Healthy AI is dedicated to helping clients understand and anticipate their legal and enforcement risks when incorporating AI and other emerging technologies into their operations.

At a glance:

- We are at the forefront of partnering with clients to ensure the development and deployment of healthcare AI technologies in compliance with law and regulations, industry standards, ethical principles and best practices.
- We offer a full suite of services focused on healthcare clients, including regulatory compliance, risk management, IP protection, and data strategy.
- We advise clients on how HIPAA regulated information and other sensitive information may be used with AI in compliance with law.
- We assist healthcare companies in developing AI governance programs, including developing policies and procedures and providing training.
- We negotiate with AI vendors to manage risk on behalf of companies contracting to use AI.
- We assist providers in developing appropriate consents, notices and other documentation where needed to mitigate risk and comply with law.

For more information about Sheppard Healthy AI, including its thought leadership and round tables, please contact Carolyn Metnick.

Carolyn Metnick
Partner • Chicago
+1.312.499.6315
cmetnick@sheppard.com
Bio

Connect and share ideas with industry leaders by joining Sheppard's Healthy AI LinkedIn Group, a hub for networking and innovation in healthcare AI.

Learn more about Sheppard's nationally ranked Healthcare Industry Team.