



**RECENT TRENDS AND DEVELOPMENTS IN
ARTIFICIAL INTELLIGENCE
AND DIGITAL TECHNOLOGIES**

**Presidential Task Force on Artificial Intelligence
and Digital Technologies**

APRIL 2026

THE ASSOCIATION OF THE BAR OF THE CITY OF NEW YORK
42 West 44th Street, New York, NY 10036
212.382.6600 | www.nycbar.org

ask for warranties that go far beyond accuracy in outputs and that would trigger a warranty breach for facts and circumstances outside the control of a software provider.¹⁶⁰

Finally, a growing trend in enterprise negotiations is the requirement for “kill switches” or the ability to disable specific AI functionalities. This allows organizations to adopt the core software while opting out of AI features that may not yet meet their internal risk tolerance or security standards. Procurers of enterprise software that contain discrete AI functionality may push for this approach, thereby allowing for the larger system to be procured and implemented while the AI tools are evaluated. Doing so, however, could surface an internal tension at organizations between the procurement team and the business unit that champions the product. On the one hand, a risk-management-first conservative approach by the procurement team will want optionality to navigate risk, while the business units will call for full functionality of the product. A procurer that licenses a product that includes various AI functionality and then turns off all of that functionality may find that it has paid for the equivalent of an empty box. Working with organizational leaders to navigate these dynamics before the contract negotiation stage will reap long term value from AI related processes.

G. Subcommittee on Artificial Intelligence in Healthcare¹⁶¹

The AI Healthcare Subcommittee focuses on the widespread use of AI in healthcare, its potential and risks, and the changing regulatory landscape. AI is used in many aspects of healthcare, for example, radiology, clinical decision support, administrative tasks, research, and drug development. The use of AI in healthcare carries risks including bias, hallucination, error, privacy-related violations, overreliance on automation outputs (known as automation bias), insufficient transparency, skill degradation, and reduced interaction between patients and doctors. As new AI tools may combine data from clinical trials, wearables, and clinical care with environmental, diet and lifestyle data, the cybersecurity and privacy risks are widespread with many points at which a breach can occur. There is an ongoing debate about whether existing laws are applicable and sufficient or whether new laws are needed. This report describes uses of AI in healthcare and then focuses on the legal and governance landscape of 2024 and 2025.

Uses of AI in Healthcare

i. Radiology

Technology has been embedded in radiology since its inception. Professional societies, including the American College of Radiology, emphasize that AI should assist radiologists, not supplant their educated decision making.¹⁶² In the same vein, the American Medical Association refers to augmented, rather than artificial, intelligence.¹⁶³ AI in radiology calls for scrutiny to ensure that AI-generated or AI-enhanced images are an accurate representation. Radiologists must minimize harm through responsible use.¹⁶⁴ Imaging AI often relies on “shortcuts” whereby the AI detects a proxy for rather than evidence of a disease or condition. For example, in the case of pneumonia, an AI tool that detects surrogate markers rather than pathological evidence associated the presence of chest tubes with pneumonia.¹⁶⁵

In addition to the way in which AI employs shortcuts, some radiology tools pose special risks. Deep learning image reconstruction (DLR) relies on Generative Adversarial Networks

(GANs) to allow CT scans to deliver lower radiation.¹⁶⁶ DLR is used to reduce the radiation exposure from CT scans by using AI to reconstruct images, filling in gaps in lower resolution, lower radiation scans. This means the AI is not just analyzing images, it is contributing to the image itself. “Additionally, the potential for DLR methods to generate false structures or ‘hallucinations,’ especially at low radiation doses, emphasizes the need for effective monitoring and mitigation strategies from both technical and clinical perspectives.”¹⁶⁷ DLR enhances image quality and can reduce noise. Some tools decrease contrast-to-noise ratio.¹⁶⁸ Training data is especially important when AI is used this way. The AI could obscure a lesion or bring in an anomaly that does not exist in the patient. Because of its risks and because it may not compensate for the image quality reduction when radiation is decreased, experts warn to use caution.

ii. *Decision support*

Large language models for clinical decisions. There is an increasing number of large language models (LLMs) designed specifically for clinical care and research. For example, OpenEvidence includes only peer reviewed articles from leading journals and has relationships with the New England Journal of Medicine and JAMA. Some products note each journal’s impact factor or allow searches to include only specified journals or those with a minimum impact factor. Institutions can build AI tools using platforms like Vertex AI to specifically meet their needs, including the need to prioritize quality of sources. These specific LLMs tailored to medicine provide clinicians with a shortcut when they evaluate patients. Nonetheless, clinicians increasingly use general purpose AI (foundation models) like GPT-3 and GPT-4, and platforms including ChatGPT for information about drug interactions, diagnosis support, treatment planning, literature searches, analysis, and educating themselves and patients.¹⁶⁹ AI tools trained to deliver in medical contexts may have lower hallucination rates and less likelihood of error than general purpose (foundation) models, including ChatGPT, and DeepSeek, a multimodal model. More experienced clinicians may be more likely to notice an error, misinterpretation, or hallucination. Many argue that without a need to read each long article completely and extract important information, clinicians may be free to turn their attention to things AI cannot do, like build personal patient relationships.¹⁷⁰

Machine learning tools for clinical decision support. Non-device machine learning tools for clinical decision support make recommendations and provide options rather than singular diagnoses or answers. Examples of clinical decision support tools include the Statin Choice Decision Aid and the Kidney Failure Risk Equation.¹⁷¹ Clinical decision support tools help emergency room clinicians with triage and estimates of risk of sepsis, infection, and need for ICU. Clinical decision support tools evaluate drug options in oncology. For example, CancerLinQ provides data from 9 million patients and is designed to provide evidence-based recommendations.¹⁷² Clinical decision support tools are often built into electronic health records systems and may make recommendations in the form of pop-up alerts and notifications.

Digital Twins. Digital twin technology can replicate cells, systems, patients, and populations using numerical representations and computations. Digital twin clinical tools are multimodal and used for clinical decision making and interventions. For example, digital twins allow oncologists to do virtual testing for safety and efficacy before delivery of a drug to the patient. Similarly, replicas of a patient’s organ allow surgeons to prepare and test surgical procedures. Digital twins have a significant amount of patient data in one product, which must be

kept secure. Digital twins are best approached as a more complex AI tool rather than a new categorically different product. It remains unclear whether institutions will take additional precautions when approaching governance of these AI tools.

iii. Administration

AI used in healthcare administration can ease coordination of care, communications, and scheduling. AI is built into electronic health records systems (EHR), improving document and calendar management and allowing clinicians to focus on the delivery of medical care. The American Medical Informatics Association Task Force to Reduce Documentation Burden and accompanying toolkit¹⁷³ highlights the many sources of healthcare documentation and their role in clinician burnout.¹⁷⁴ It provides strategies for governance, metrics to measure time spent on documentation, and ways to reduce duplication or redundancy. LLMs are a new method of facilitating patient communication. One study found patients were skeptical of AI-generated emails when they were aware that the emails were generated by AI. The same study found approval of the content and quality when patients did not know the communications were AI generated.¹⁷⁵ Importantly, a patient-centered approach to governance calls for informing patients of the use of AI tools when they could be mistaken for communications written and sent by humans. The use of ambient AI scribes can reduce documentation burden and clinician burnout.¹⁷⁶ Yet there are ethical risks and legal uncertainty about transparency, explainability, accuracy,¹⁷⁷ and governance.¹⁷⁸ Institutional governance policies should do more to ensure transparency and patient consent as ambient scribes are widely used.

iv. Research and drug discovery

GANs are two-pronged systems in which generated data is compared to and distinguished from real data. GANs help improve medical imaging,¹⁷⁹ enhance data quality, improve disease detection, and improve research efficiency by creating datapoints and replicas. GANs also allow testing variations of drugs for efficacy and side effects.¹⁸⁰ They are used in radiology, digital twins, and drug discovery, among other categories.

New molecules. GANs enable scientists to identify new molecules as there would be a nearly infinite number of possible molecules in a “chemical space.”¹⁸¹ GANs can be used for protein folding. (Protein folding is twisting of an amino acid chain to form a three-dimensional structure). Protein folding occurs naturally; faulty protein folding is linked to diseases including Alzheimer’s Disease and Parkinson’s disease. Using GANs, scientists apply protein folding to create de novo proteins¹⁸² and proteins with specific properties.¹⁸³ A noteworthy risk of lab-made molecules is that pathogens can be used for biowarfare. Researchers and clinicians may overlook the global risks in their effort to treat and cure disease. New research suggests that a loophole may exist in biosecurity screening software in that screening software may not flag dangerous proteins not found in nature.¹⁸⁴ There are laws, policies, and guidance covering this arena, yet some open questions about oversight and how to ensure safety in lab-based medical research remain.¹⁸⁵

Synthetic data. LLMs and GANs are used to create synthetic datasets. Synthetic data has many forms used in healthcare including text, images, cells, and biological data, etc. Although some synthetic data is transformed from underlying personal data,¹⁸⁶ it is not governed by data collection and use laws for privacy and consent.¹⁸⁷ Some call for more regulation, noting risks,¹⁸⁸

while others emphasize the benefits, including the absence of traditional privacy concerns that accompany data collection and the promise of less biased AI tools. If the potential of synthetic data is realized, the ability to carry out meaningful medical research without as many research subjects stands to save time and effort as well as present a way to experiment safely. To protect end users, patients, and the public, synthetic data governance could focus on the type of inference made from the data,¹⁸⁹ the impact on research, and eliminating bias¹⁹⁰ through oversight of how the synthetic data correlates to real-world data.

Policy, soft law, and standards

Many sets of standards and principles (collectively, soft law) point to how to ethically deploy AI in healthcare settings. Institutions should rely on them when formulating institutional governance policies. As healthcare providers, payors and consumers anticipate the benefits of AI, they should address the risks and challenges through compliance with soft law, federal law, state law, and good governance. With the increasing use of AI in healthcare, government, non-governmental organizations, standards associations, professional associations, and nonprofits have developed soft law, frameworks, and principles for the ethical use of AI as applied to healthcare.

i. World Health Organization

In 2024, the World Health Organization (WHO) published new guidance, Ethics and governance of artificial intelligence for health: Guidance on large multi-modal models (LMMs), (the Guidance).¹⁹¹ LMMs can analyze different types of data, for example text and images, and are thus broader than LLMs, which rely on natural language processing to generate text. For example, DeepSeek Janus Pro is an LMM while ChatGPT is an LLM. LMMs are a type of generative AI – the results they produce are not limited to the data entered. The Guidance builds on the six principles the WHO set forth in 2021 for the use of AI in healthcare:

- protect autonomy;
- promote human well-being, human safety and the public interest;
- ensure transparency, “explainability” and intelligibility;
- foster responsibility and accountability;
- ensure inclusiveness and equity; and
- promote AI that is responsive and sustainable

WHO addresses ethics and governance at three stages: development, provision, and deployment. Provision includes direct use (i.e., a user prompting an LMM to provide a healthcare result) or integrating an LMM into an app or product. The recommended role of government is strong throughout the three stages.

The Guidance points out that “regurgitating medical knowledge” is not the same as providing adequate clinical care.¹⁹² It also recommends keeping pediatric data separate from adults and not basing pediatric generalizations on adult data. Additionally, LMMs allow patients to

evaluate their own health conditions, highlighting risks beyond healthcare systems and practitioners. Risks of individuals using LMMs yet being underrepresented in training data or lacking the medical knowledge to properly interpret results and recommendations can negatively impact outcomes – the Guidance refers to this as well as the risk that clinicians will rely on inappropriately trained LMMs as “epistemic injustice.” The Guidance distinguishes risk-based and human rights-based frameworks, finding human rights and ethical principles must be upheld regardless of risk. This stance is stronger than the EU’s risk-based framework and notes the importance of scrutinizing even low-risk AI tools to ensure dignity, and autonomy.

ii. *NIST Clinical Internet of Things (CIoT) Standards*

In 2024, the Institute of Electrical and Electronics Engineers (IEEE), a global standards and professional organization, adopted new standards for healthcare tools including wearables. Trust, Identity, Privacy, Protection, Safety, and Security (TIPSS) principles govern the new standards.¹⁹³ The Clinical Internet of Things (CIoT) includes wearables, implantables, sensors, in-hospital devices, devices that communicate with each other, with electronic health records systems, and with healthcare systems. The goals of the framework include promoting data sharing, improving outcomes, and protecting patients, including protecting privacy. Interoperability was a priority. The IEEE plans to socialize and establish the standard by engaging healthcare industry and IT partners and then “Engage industry partners, device manufacturers and providers to enable adoption of the TIPSS standard and policies in their clinical IoT devices.”¹⁹⁴

iii. *Joint Commission and Coalition for Health AI*

In 2025, the Joint Commission¹⁹⁵ and Coalition for Health (CHAI)¹⁹⁶ jointly published Elements of Responsible Use of AI in Healthcare (RUAIH™), a guidance framework for implementation and use of AI in healthcare¹⁹⁷ so organizations can “fully harness the transformative potential of AI.” The guidance provides recommendations to healthcare professionals – primarily hospitals and health systems – regarding use and deployment of AI capabilities in a responsible manner. The elements of responsible use set forth are the following:

- AI Policies and Governance Structures
- Patient Privacy and Transparency
- Data Security and Data Use Protections
- Ongoing Quality Monitoring
- Voluntary, Blinded Reporting of AI Safety-Related Events
- Risk and Bias Assessment
- Education and Training

RUAIH is a broad framework that will “feed into the development and release of a series of community - and resource informed Responsible Use of AI Playbooks.” The Joint Commission’s

accreditation of healthcare institutions and quality and patient safety make RUAIH potentially more effective than other soft law instruments.

iv. *American Medical Association*

In 2025, the AMA published practical guidance regarding AI governance, including a toolkit for Governance for Augmented Intelligence.¹⁹⁸ The toolkit provides a high-level playbook emphasizing clear accountability and decision rights (e.g., executive sponsorship and defined roles/committees); standardized intake and risk-triage processes with a systemwide AI inventory; vendor evaluation and change-control requirements; training and education for AI tool users; and post-deployment monitoring (e.g., for safety, performance, bias, and user impact); and includes an adaptable model AI policy template. This operational guidance generally aligns with AMA's November 2024 Principles for Augmented Intelligence Development, Deployment, and Use in Health Care,¹⁹⁹ which emphasize physician-led oversight, transparency to patients, rigorous clinical validation, bias mitigation, and governance that is embedded across the AI lifecycle.

Federal Law

i. *FDA*

As of 2025, the FDA has authorized approximately 1,250 AI-enabled devices for use in clinical settings, approximately 65 percent of which are used for radiology. The FDA regulates these technologies under its authority in the Federal Food, Drug, and Cosmetic Act, which treats software that meets the definition of a medical device – including “Software as a Medical Device” (SaMD) and AI/ML-enabled functions - as subject to FDA clearance, approval and post-market oversight. In 2025, the FDA published draft guidance titled *Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations* (the AI-enabled Device Software Guidance).²⁰⁰ Together with the 2024 and 2025 guidance on Predetermined Change Control Plans²⁰¹ for AI-enabled device software, the FDA is refining the framework that allows adaptive algorithms to update while preserving ongoing safety and effectiveness. The AI-enabled Device Software Guidance provides nonbinding recommendations with a focus on the total product lifecycle and performance over time.²⁰² The FDA clarified its thinking about transparency and bias. The FDA views transparency holistically and as context dependent suggesting the “right information at the right time” is crucial. The guidance describes transparency in labeling, manuals, user interfaces, and marketing materials. In this guidance, transparency refers to clearly communicating the contextually relevant performance and design information of a device to the appropriate stakeholders in a manner that they can understand and act on. Transparency involves ensuring that important information is both accessible and functionally comprehensible and is connected both to the sharing of information, and to the usability of a device. As such, a user-centered approach to transparency design helps support the safe and effective use of AI-enabled devices. Including appropriate transparency information has also been shown to more than double willingness to use a device.²⁰³ The FDA's draft guidance suggests that attention to safety, efficacy, transparency, and risk of bias may lead clinicians to feel more confident with AI-enabled medical devices.

ii. *Health and Human Services Office of Civil Rights*

When AI is trained on unrelated tasks or biased training data, the risk of bias is high, and may lead to a higher rate of misdiagnosis, or impact safety and effectiveness in certain groups by age, sex, and race.²⁰⁴ The use of biased AI may violate antidiscrimination laws.²⁰⁵ The Health and Human Services Office of Civil Rights (OCR) promulgated a final rule in May 2024 to clarify that covered entities are prohibited from discriminating through the use of clinical algorithms on the basis of race, color, national origin, sex, age, or disability under Section 1557 of the Affordable Care Act.²⁰⁶ Covered entities must identify algorithms posing risk of bias and mitigate the risk. Portions of the 2024 Section 1557 final rule have been subject to multiple judicial challenges.

Separately, in late 2024, OCR proposed significant updates to the HIPAA Security Rule with direct implications for AI systems that process electronic protected health information.²⁰⁷ Additionally, in December 2025, HHS published a Request for Information on Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care, soliciting feedback on how current HHS regulations impact AI adoption, including reimbursement barriers to AI clinical tools.²⁰⁸

State Laws

As healthcare is significantly regulated at the state level, many state legislatures are addressing AI in healthcare.

i. *Colorado AI Act*

The Colorado AI Act, with certain obligations taking effect on February 1, 2026, applies to developers and deployers. It imposes reporting requirements and impact assessments. Developers and deployers have a duty to exercise reasonable care to protect consumers from known or foreseeable risks of algorithmic discrimination in high-risk AI.²⁰⁹ High risk AI is defined as a system that substantially contributes to or makes a decision with a material impact in eight contexts, including healthcare.²¹⁰ It also requires consumer notifications to ensure that consumers know when they are interacting with an AI system. Violations will be considered unfair trade practices, but the law does not provide a private right of action.²¹¹ The Colorado Attorney General has enforcement authority. Executive Order 14365, issued in December 2025, specifically cited the Colorado AI Act's anti-discrimination provisions as an example of state regulation that the Trump Administration views as potentially burdensome, and directed the creation of an AI Litigation Task Force to challenge certain state AI laws.²¹²

ii. *California*

The California Privacy Protection Agency also finalized regulations in September 2025, to include requirements around automated decision-making technology (ADMT), among other updates to the California Consumer Privacy Act regulations.²¹³ ADMT is defined to include any technology that “processes personal information and uses computation to replace human decisionmaking or substantially replace human decisionmaking.”²¹⁴ The requirements include, among other obligations, providing individuals with notice before an ADMT is used; disclosing the specific right to opt-out of that use; and, in certain contexts, obtaining the individual’s consent prior to deployment. California’s AB 3030, effective January 1, 2025, mandates that healthcare

providers notify patients when generative AI is used to generate written or verbal communications that contain clinical information. California requires healthcare providers to tell patients when generative AI is used in communications and to provide human contact options.²¹⁵ However, the law exempts communications that were read and reviewed by a healthcare provider.²¹⁶

iii. Additional state laws

Similarly, Utah requires disclosure of the use of generative AI.²¹⁷ Hospitals and healthcare systems must “prominently” disclose that they are using generative AI to communicate. They also have responsibility for the content, i.e., the information provided to consumers/patients in a chatbot conversation, AI-generated email or text, etc. Deployers of AI cannot blame the AI or use the AI as a defense against a consumer protection law violation.²¹⁸ Texas’ Responsible AI Governance Act, effective January 1, 2026, obligates providers to disclose AI use to patients before or at the time of service (or as soon as reasonably possible in case of an emergency).²¹⁹ Several broadly applicable state consumer privacy laws effective in 2025, such as those in Delaware and New Hampshire, expand consumer rights over their personal data, including affording consumers the right to opt out of profiling (e.g., AI-powered processing of personal data to predict aspects of an individual) resulting in significant decisions such as those impacting healthcare. It is advisable to track and prepare for state-law duties regarding automated decisions and ensure vendor contracts mirror applicable obligations. For example, contracts can prohibit training on patient data without express, written approval; require disclosure of material AI model changes; include indemnities for deceptive claims; and require strict incident reporting and post-deployment audits (e.g., for bias, safety, and effectiveness). Steps such as these can serve as evidence of reasonable care, if challenged.

Informed Consent

It is crucial that patients understand the risks and benefits of the care they receive. Frameworks for informed consent vary, highlighting the lack of consensus about whether to disclose the use of AI to patients, what to disclose, and how to do so. Notably, no federal statute or regulation currently requires informed consent specifically for the use of AI in clinical-decision making. While new state AI laws impose requirements for the disclosure of generative AI, and in some cases, of high-risk AI, uncertainty remains. One approach set forth by Susannah Rose and Devora Shapiro includes the degree of autonomy the AI tool has, deviation from the standard of care, whether the AI is patient facing, and the clinical risk.²²⁰ Another approach is the FAIR-AI framework, which couches consent in vulnerability, transparency, and equity best practices noting an “ethical obligation to notify patients when AI is being used and, when appropriate, to obtain their consent—particularly in sensitive or high-stakes situations. This obligation is heightened when there is no human oversight, when the technology is experimental, or when the use of AI is not readily apparent.”²²¹ Another framework emphasizes the stronger need for informed consent when the likelihood that a clinician would overly rely on the AI tool is greater. Some frameworks consider inconvenience and administrative burden, while others prioritize transparency, even attaching it to trust in systems and clinicians.