

Augmented Intelligence Development, Deployment, and Use in Health Care

BACKGROUND

The issue of AI first presented itself as an area of potential interest to AMA physicians and medical students that necessitated creation of AMA policy in 2018. At that time, physicians and medical students primarily considered AI-enabled technologies within the context of medical device and clinical decision support, although administrative applications of AI began to grow exponentially and started to gain traction in the hospital, health system, and insurer space. Since the development of the AMA's foundational AI policy in 2018 and subsequent policy on coverage and payment for AI in 2019, the number of AI-enabled medical devices approved by the U.S. Food and Drug Administration (FDA) has grown to over 800. In 2022, the concept of "generative AI" and what it can do became better understood to the public. Generative AI is a broad term used to describe any type of artificial intelligence that can be used to create new text, images, video, audio, code, or synthetic data. Generative AI and LLMs have rapidly transformed the use cases and policy considerations for AI within health care, necessitating updated AMA policy that reflects the rapidly evolving state of the technologies.

AMA policy adopted in [2018](#) and [2019](#) enabled the AMA to be a strong advocate on behalf of patients and physicians and has been the bedrock of AMA's advocacy on AI in the form of lobbying key congressional committees, participating in expert panel discussions, creating educational resources, and working with our Federation colleagues at the federal and state levels. However, as AI has rapidly developed beyond AI-enabled medical devices and into LLMs/generative AI, new policy and guidance are needed to ensure that they are designed, developed, and deployed in a manner that is ethical, equitable, responsible, accurate, and transparent.

As an initial step, in November 2023, the AMA Board of Trustees approved a set of advocacy principles developed by the Council on Legislation (COL) that serve as the framework of this Board report. The main topics addressed in the principles include AI oversight, disclosure requirements, liability, data privacy and security, and payor use of AI. In addition to the COL, these principles have been vetted among multiple AMA business units, and AMA staff has worked with several medical specialty societies that have an expertise in AI and has received additional guidance and input from outside experts that have further refined these principles, which serve as the foundation for the included policy. The resulting policy builds upon and are supplemental to the AMA's existing AI policy, especially Policy [H-480.940](#), "Augmented Intelligence in Health Care," Policy [H-480.939](#), "Augmented Intelligence in Health Care," and Policy [D-480-956](#), "Use of Augmented Intelligence for Prior Authorization," as well as the [AMA's Privacy Principles](#).

This report highlights the AMA's recognition of the issues raised at the A-23 and I-23 HOD meetings, as well as the comments heard during the A-24 HOD meeting regarding BOT Report 15-A-24. It also introduces and explains major themes of the report's recommendations and provides background information on the evolution of AI policy in health care and the direction that policy appears to be headed.

Current Status of Oversight of Augmented Intelligence-Enabled Technologies

There is currently no whole-of-government strategy for oversight and regulation of AI. The U.S. Department of Health and Human Services (HHS) did establish an AI Office in March 2021 and developed a general strategy to promote the use of trustworthy AI but has not produced a department-wide plan for the oversight of AI. While many other federal departments and agencies also have some authority to regulate health care AI, many regulatory gaps exist. The Assistant Secretary for Technology Policy/Office of the National Coordinator for Health Information Technology (ASTP/ONC) recently created a position for a Chief AI Officer. However, the job role is targeted at the internal use of AI within HHS and less about public policy. To address the lack of a national strategy and national governance policies directing the development and deployment of AI, the federal government has largely defaulted to public “agreements” representing promises by large AI developers and technology companies to be good actors in their development of AI-enabled technologies.

In December 2023, the Biden Administration released a reasonably comprehensive [executive order](#) on the “Safe, Secure, and Trustworthy Development and Use of Artificial Intelligence.” While the executive order does not create new statutory or regulatory requirements, it does serve to direct federal departments and agencies to take action to provide guidance, complete studies, identify opportunities, etc. on AI across several sectors, including HHS. The AMA was pleased to see close alignment between the executive order’s direction and AMA principles. However, executive orders do not represent binding policy, so the regulatory status quo remains unchanged at present.

The Biden Administration had also previously released a “[Blueprint for an AI Bill of Rights](#),” setting forth five principles that should guide the design, use, and deployment of AI. Those include recommendations for creating safe and effective systems; algorithmic discrimination protections; data privacy; notice and explanation; and human alternatives, considerations, and fallback. Like executive orders, this blueprint does not create new or binding policy with the force of law.

There have been few, but notable, additional actions by federal agencies that may serve to impact patient and physician interaction with AI-enabled technologies. In 2022, the Centers for Medicare & Medicaid Services (CMS) and HHS Office for Civil Rights (OCR) introduced a sweeping liability proposal within its Section 1557 Non-Discrimination in Health Programs and Activities proposed rule. The AMA submitted detailed comments opposing this section of the proposed rule. OCR ultimately finalized the rule, including the new section prohibiting discrimination by clinical algorithms. The final rule requires physicians to make “reasonable efforts” at identifying and mitigating discriminatory harms from algorithms, including AI.

In addition, the ASTP/ONC* proposed and finalized, with some modifications, policies that will require electronic health record (EHR) technology developers to make certain information about AI used in EHRs available to physicians and other users. ASTP/ONC refers to these AI tools as Predictive Decision Support Interventions (Predictive DSIs). Starting in 2025, EHR developers that supply Predictive DSIs as part of the developer’s EHR offering must disclose specific attributes and inform users if patient demographic, social determinants of health, or health assessment data are used in the Predictive DSI. EHRs will be subject to regulatory requirements regarding the design, development, training, and evaluation of Predictive DSIs along with mandated risk management practices. ASTP/ONC’s stated goal is to ensure that physicians understand how these tools work, how data are used, the potential for bias, and any known limitations.

FDA Approved AI-Enabled Medical Devices

The FDA continues to rapidly approve AI-enabled medical devices. While FDA approval and clearance of algorithmic-based devices date back to 1995, clearance and approval of these devices has rapidly accelerated in the last several years. As of May 2024, 882 devices that FDA classifies as Artificial Intelligence/Machine Learning (AI/ML) devices have been approved for marketing. The overwhelming number of these devices are classified as radiology devices and this category of devices has seen the steadiest increases in the number of applications for FDA approval. However, the number of applications is increasing in several specialties, including cardiology, neurology, hematology, gastroenterology, urology, anesthesiology, otolaryngology, ophthalmology, and pathology. A significant number of cleared or approved devices are considered diagnostic in nature and many currently support screening or triage functions.

In 2017, the FDA announced that it was evaluating a potentially new regulatory approach towards Software as a Medical Device, which would include AI/ML technologies. The so-called Pre-Certification program, or “Pre-Cert,” progressed to an initial pilot program involving nine manufacturer applicants. The program proposed to pre-certify manufacturers of software-based medical devices. Devices developed by pre-certified manufacturers would be subject to varying levels of FDA review based on risk to patients, including potentially being exempt from review if the risk is low. However, the Pre-Cert program has been tabled and the pilot dismantled for the time being, leaving FDA to utilize traditional review pathways for AI-enabled medical devices. In the absence of new regulatory strategies tailored to Software as a Medical Device (SaMD) and AI/ML, FDA has issued some proposed guidance for developers of these devices but has not yet moved forward with additional guidance for important, physician-facing topics, such as transparency and labeling requirements. In June 2024, the FDA released a set of “guiding principles” for AI transparency in conjunction with Health Canada and the Medicines and Healthcare Products Regulatory Agency of the United Kingdom. However, these guiding principles do not represent official FDA guidance nor are they mandatory requirements of applicants for FDA review. The continued lack of transparency mandates leaves a critical gap in the oversight of AI-enabled medical devices.

Data Privacy and Cybersecurity Considerations in Health Care AI

The integration of AI into health care signifies a transformative era, with potential to greatly enhance patient care and operational efficiency. However, this advancement also introduces considerable challenges, particularly in data privacy and cybersecurity. As health care facilities, technology vendors, clinicians, and users increasingly adopt AI, it is vital to focus on protecting patient and user data and securing AI systems against cyber threats. Handling vast amounts of sensitive data raises critical questions about privacy and security. Survey data has shown that nine out of 10 patients believe privacy is a right and nearly 75 percent of people are concerned about protecting the privacy of their health data.¹ Addressing these concerns necessitates a multifaceted approach that includes advanced data privacy techniques, data use transparency, robust cybersecurity strategies, and compliance with regulatory standards.

Ensuring the protection of patient data in the context of AI requires sophisticated privacy techniques. Key methods such as anonymization and pseudonymization can remove or replace personal identifiers in data sets and significantly reduce the risk of re-identification. Additionally, implementing a robust data management system empowers patients by providing clear ways to grant, deny, or revoke consent for the use of their data, enhancing patient trust and ensuring compliance with global data protection regulations such as the General Data Protection Regulation and the Health Insurance Portability and Accountability Act (HIPAA). Moreover, the collection of data should be kept to a minimum. By collecting only the data necessary for the intended purpose, AI systems can mitigate the risks associated with data breaches and misuse.



Cybersecurity plays a crucial role in health care, especially in the context of the increasing digitalization of medical records, patient data, and health care services. The health care sector is a prime target for cyber-attacks due to the sensitivity and value of the data it handles, including personal health information (PHI), financial data, and intellectual property related to medical research. The integration of technology in health care has undoubtedly brought significant benefits such as improved patient care, streamlined operations, and enhanced data analytics. However, it also introduces vulnerabilities. These include potential unauthorized access, data breaches, and disruptions to health care services, which can have dire consequences for patient privacy and safety. In 2017, 83 percent of surveyed physicians had already experienced a cyberattack and 85 percent stated that they want to share electronic PHI but were concerned about the data security necessary to protect it.² This risk is amplified by the recent increased use of interconnected devices and systems, such as EHRs, telemedicine platforms, and mobile health applications.

The attack on Change Healthcare in February 2024 is a stark reminder of the critical importance of cybersecurity in health care. Change Healthcare, a division of UnitedHealth Group, was struck by a ransomware attack that significantly disrupted the largest health care payment and operations system in the United States. This incident led to widespread disruptions, affecting thousands of medical practices, hospitals, pharmacies, and others. The attack was attributed to ransomware. Despite efforts to recover from this attack, the impact on health care operations was profound, including the disruption of claims processing, payments, and electronic prescriptions leading to financial strain on physicians and delays in patient care. The health care sector's reliance on interconnected digital systems for patient records, billing, and payments, means that the impact of a cyberattack can be both immediate and widespread, affecting patient care and operational continuity.

The implications of cybersecurity in health care AI are multifaceted. AI in health care, encompassing machine learning algorithms, predictive analytics, and robotic process automation, holds immense potential for diagnostic accuracy, personalized medicine, and operational efficiency. However, the deployment of AI in health care settings creates unique cybersecurity challenges. AI systems require large datasets to train and operate effectively, increasing the risk of large-scale data breaches. Additionally, the complexity of AI algorithms can make them opaque and vulnerable to manipulation, such as adversarial attacks that can lead to misdiagnoses or inappropriate treatment recommendations. AI-driven health care solutions often rely on continuous data exchange across networks, escalating the risk of cyber-attacks that can compromise both the integrity and availability of critical health care services.

A model stealing attack represents a significant cybersecurity threat in the realm of AI, where a malicious actor systematically queries an AI system to understand its behavior and subsequently replicates its functionality. This form of intellectual property theft is particularly alarming due to the substantial resources and time required to develop sophisticated AI models. An example of this issue involves a health care organization that has invested heavily in an AI model designed to predict patient health outcomes based on a wide range of variables. If a malicious entity were to engage in model stealing by extensively querying this predictive model, it could essentially duplicate the original model's predictive capabilities along with capitalizing on sensitive health care information and physicians, users, or the entity's intellectual property. Absent strong protections against input manipulation and malicious attacks, AI can become a new conduit for bad actors to compromise health care organizations and harm patients. This not only undermines the original investment but also poses a direct threat to the competitive advantage of the innovating organization.

* On July 25, 2024, HHS announced that ONC will be renamed the Assistant Secretary for Technology Policy and Office of the National Coordinator for Health Information Technology (ASTP/ONC).



Moreover, the risk extends beyond intellectual property theft to encompass serious privacy concerns. This is exemplified by incidents where generative AI models, trained on vast datasets, inadvertently reveal sensitive information contained within their training data in response to certain prompts. In the health care sector, where models are often trained on highly sensitive patient data, including personally identifiable information, the unauthorized extraction of this data can lead to significant breaches of patient confidentiality. The dual threat of intellectual property theft and data privacy breaches underscores the critical need for robust cybersecurity measures in safeguarding AI models, particularly those developed and utilized within the health care industry, to maintain the integrity of both their intellectual property and the confidentiality of the sensitive data they handle.

While there are new federal policies to increase data transparency when AI is used in conjunction with health information technology, such as those issued by ASTP/ONC, these new policies only cover the certified EHR developer and stop short of holding AI developers accountable for robust data governance or data security and privacy practices.³

Generative AI

The broad introduction of generative AI into the public sphere in 2022 saw a paradigm shift in how physicians contemplated AI. Open-source LLM Chat GPT presented a new, easily accessible AI-enabled technology with significant capabilities to generate new content and provide readily available access to information from a huge number of sources. Generative AI tools have significant potential to relieve physician administrative burdens by helping to address actions such as in-box management, patient messages, and prior authorization requests. They also show promise in providing clinical decision support and highly personalized treatment recommendations.

However, these generative AI tools can also pose significant risk, particularly for clinical applications. As these LLMs are constantly evolving, they run the risk of providing inconsistent responses on the same fact pattern on potentially a daily, weekly, monthly, or yearly basis. The risks of these tools fabricating content are well known and could serve to propagate the spread of medical misinformation as content fabricated by the AI technologies is more broadly disseminated. They also pose potentially significant data privacy concerns.

At the present time, these technologies are largely unregulated, as there is no current regulatory structure for generative AI clinical decision support tools unless they meet the definition of a medical device regulated by the FDA. The U.S. Federal Trade Commission (FTC) has limited authority to regulate data privacy issues that may be associated with generative AI. The FTC does have some authority to regulate activities considered to be an unfair, deceptive, or abusive business practice and can enforce laws for consumer protection. However, these authorities are not specific to AI and the agency is generally under-resourced in this area. CMS has some authority to regulate use of AI by entities receiving funds from Medicare and Medicaid, including use by Medicare Advantage plans. OCR has some additional authorities to regulate data privacy and nondiscrimination.

While some federal agencies may have oversight and authorities to regulate some aspects of AI, there are many regulatory gaps. These regulatory gaps are particularly significant when considering generative AI, as tools like ChatGPT and others currently fall well outside the definition of a regulated medical device. While generative AI use for clinical applications is relatively limited currently, it is expected to grow and patients and physicians will need assurances that it is providing safe, accurate, non-discriminatory answers to the full extent possible, whether through regulation or generally accepted standards for design, development, and deployment.

Physician Liability for Use of AI

One of the most significant concerns raised by physicians regarding the use of AI in clinical practice is concern



over potential liability for use of AI that ultimately performs poorly. The question of liability for the use of AI is novel and complex given that the use of AI for activities, such as clinical decision making and treatment recommendations, introduces an element of shared decision making between the patient, physician, and now the machine. While it is likely that liability will mostly be determined by the legal system through decisions in courts of law, some federal agencies have considered the idea of physician liability in these instances. Notably, the HHS Office of Civil Rights has finalized a rule creating new liability for physicians utilizing AI that results in discriminatory harms to patients. This could include, for example AI that utilizes algorithms with race adjustments or returns otherwise biased results to physicians and patients. The final rule prohibits discrimination by clinical algorithms and requires physicians, hospitals, health systems, and others to use “reasonable efforts” to both identify algorithmic discrimination and to mitigate resulting harms. While the AMA supports a prohibition on discrimination by clinical algorithms, the AMA strongly opposed efforts to create new physician liability for the use of AI.

Use of AI By Payors

There have been numerous reports recently regarding the use of what has been termed “automated decision-making tools” by payors to process claims. However, numerous reports regarding the use of these tools show a growing tendency toward inappropriate denials of care or other limitations on coverage. Reporting by ProPublica claims that tools used by Cigna denied 300,000 claims in two months, with claims receiving an average of 1.2 seconds of review.⁴ Two class action lawsuits were filed during 2023, charging both United Health Care and Humana with inappropriate claims denials resulting from use of the nHPredict AI model, a product of United Health Care subsidiary NaviHealth. Plaintiffs in those suits claim the AI model wrongfully denied care to elderly and disabled patients enrolled in Medicare Advantage (MA) plans with both companies. Plaintiffs also claim that payors used the model despite knowing that 90 percent of the tool’s denials were faulty.

There is growing concern among patients and physicians about what they perceive as increasing and inappropriate denials of care resulting from the use of these automated decision-making tools. In his recent Executive Order on AI, President Biden addressed this issue as an area of concern, directing HHS to identify guidance and resources for the use of predictive and generative AI in many areas, including benefits administration, stating that it must take into account considerations such as appropriate human oversight of the application of the output from AI.

There are currently no statutory and only limited regulatory requirements addressing the use of AI and other automated decision-making tools by payors. States are beginning to look more closely at this issue given the significant negative reporting in recent months and are a likely place for near-term action on this issue. Congress has also shown increasing concern and has convened hearings for testimony on the issue; however, there has been no further Congressional action or legislation to pursue further limitations on use of these algorithms. Additionally, CMS has not taken broad regulatory action to limit the use of these algorithms by entities administering Medicare and Medicaid benefits.

AMA POLICY

The AMA has existing policies, [H-480.940](#) and [H-480.939](#) both titled “Augmented Intelligence in Health Care,” which stem from a 2018 and 2019 Board report and cover an array of areas related to the consequences and benefits of AI use in the physician’s practice. In pertinent part to this discussion, AMA Policy H-480.940 seeks to “promote development of thoughtfully designed, high-quality, clinically validated health care AI, encourage education for patients, physicians, medical students, other health care professionals, and health administrators to promote greater understanding of the promise and limitations of health care AI, and explore the legal



implications of health care AI, such as issues of liability or intellectual property, and advocate for appropriate professional and governmental oversight for safe, effective, and equitable use of and access to health care AI." This policy reflects not only the significance of attribution on the part of the developer, but furthermore emphasizes that physicians and other end users also play a role in understanding the technology and the risks involved with its use.

AMA Policy H.480.939 also addresses key aspects of accountability and liability by stating that "oversight and regulation of health care AI systems must be based on risk of harm and benefit accounting for a host of factors, including but not limited to: intended and reasonably expected use(s); evidence of safety, efficacy, and equity including addressing bias; AI system methods; level of automation; transparency; and, conditions of deployment." Furthermore, this policy asserts that "liability and incentives should be aligned so that the individual(s) or entity(ies) best positioned to know the AI system risks and best positioned to avert or mitigate harm do so through design, development, validation, and implementation. Specifically, developers of autonomous AI systems with clinical applications (screening, diagnosis, treatment) are in the best position to manage issues of liability arising directly from system failure or misdiagnosis and must accept this liability with measures such as maintaining appropriate medical liability insurance and in their agreements with users."

AMA Policy [D-480.956](#) supports "greater regulatory oversight of the use of augmented intelligence for review of patient claims and prior authorization requests, including whether insurers are using a thorough and fair process that: (1) is based on accurate and up-to-date clinical criteria derived from national medical specialty society guidelines and peer reviewed clinical literature; (2) includes reviews by doctors and other health care professionals who are not incentivized to deny care and with expertise for the service under review; and (3) requires such reviews include human examination of patient records prior to a care denial."

AMA Policy [H-480.935](#) directs our AMA to study and develop recommendations on the benefits and unforeseen consequences to the medical profession of LLMs such as generative pretrained transformers (GPTs), and other augmented intelligence-generated medical advice or content. In addition to a report back to the HOD, this policy directs AMA to work with the federal government and other appropriate organizations to protect patients from false or misleading AI-generated medical advice; encourage physicians to educate patients about the benefits and risks of consumers facing LLMs including GPTs; and support publishing groups and scientific journals in efforts to ensure transparency and accountability of authors in the use and validation of text generated by augmented intelligence.

DISCUSSION

As the number of AI-enabled health care tools and systems continues to grow, these technologies must be designed, developed, and deployed in a manner that is ethical, equitable, responsible, accurate, and transparent. With a lagging effort towards adoption of national governance policies or oversight of AI, it is critical that the physician community engage in development of policies to help drive advocacy, inform patient and physician education, and guide engagement with these new technologies. It is also important that the physician community help guide development of these tools in a way that best meets both patient and physician needs, and help define their own organization's risk tolerance, particularly where AI impacts direct patient care. AI has significant potential to advance clinical care, reduce administrative burdens, and improve clinician well-being. This may only be accomplished by ensuring that physicians engage only with AI that satisfies rigorous, clearly defined standards to meet the goals of the quadruple aim,⁵ advance health equity, prioritize patient safety, and limit risks to both patients and physicians.

Oversight of Health Care Augmented Intelligence

There is currently no national policy or governance structure in place to guide the development and adoption of non-medical device AI. As discussed above, the FDA regulates AI-enabled medical devices, but many types of AI-enabled technologies fall outside the scope of FDA oversight.⁶ This potentially includes AI that may have clinical applications, such as some generative AI technologies serving clinical decision support functions. While the FTC and OCR have oversight over some aspects of AI, their authorities are limited and not adequate to ensure appropriate development and deployment of AI generally, and specifically in the health care space. Likewise, ASTP/ONC's enforcement is limited and focused on EHR developers' use and integration of AI within their federally certified EHRs. While this is a major first step in requiring AI transparency, it is still the EHR developer that is regulated with few requirements on the AI developer itself. Encouragement of a whole-of-government approach to implement governance policies will help to ensure that risks to consumers and patients arising from AI are mitigated to the greatest extent possible.

In addition to the government, health care institutions, practices, and professional societies share some responsibility for appropriate oversight and governance of AI-enabled systems and technologies. Beyond government oversight or regulation, purchasers and users of these technologies should have appropriate and sufficient policies in place to ensure they are acting in accordance with the current standard of care. Similarly, clinical experts are best positioned to determine whether AI applications are high quality, appropriate, and whether the AI tools are valid from a clinical perspective. Clinical experts can best validate the clinical knowledge, clinical pathways, and standards of care used in the design of AI-enabled tools and can monitor the technology for clinical validity as it evolves over time.

Transparency in Use of Augmented Intelligence-Enabled Systems and Technologies

As implementation of AI-enabled tools and systems increases, it is essential that use of AI in health care be transparent to both patients and physicians. Transparency requirements should be tailored in a way that best suits the needs of the end users. Care must be taken to preserve the integrity of data sets used in health care such that individual choice and data privacy are balanced with preserving algorithms that remain as pristine as possible to avoid exacerbating health care inequities. Disclosure should contribute to patient and physician knowledge without increasing administrative burden. When AI is utilized in health care decision-making at the point of care, that use should be disclosed and documented to limit risks to, and mitigate inequities for, both patients and physicians, and to allow each to understand how decisions impacting patient care or access to care are made. While transparency does not necessarily ensure AI-enabled tools are accurate, secure, or fair, it is difficult to establish trust if certain characteristics are hidden.

Heightened attention to transparency and additional transparency requirements serve several purposes. They help to ensure that the best possible decisions are made about a patient's health care and help patients and physicians identify critical decision points and possible points of error. They can also serve as mechanisms to help shield physicians from liability so that potential issues related to use of AI-enabled technologies can be isolated and accountability apportioned appropriately.

There are currently few federal requirements for transparency regarding AI. The FDA requires product labeling to provide certain information to physicians and other users, but requirements for device labeling are generally considered to be less stringent and have more leeway than drug product labeling. While FDA has stated that transparency is a key priority for the agency to address, they have not taken any additional action to update the labeling requirements for AI-enabled medical devices or put into place additional transparency requirements for AI-enabled devices. As discussed above, ASTP/ONC also has new transparency requirements applicable to the



use of AI within EHRs; however, again, those requirements are limited to AI within an EHR or other applications integrated and made available through the EHR. They will not apply to AI-enabled tools accessible through the Internet, cellular phones, etc. There is an urgent need for additional federal action to ensure AI transparency.

Transparency: Attributes and the Importance of Disclosure

During consideration of an earlier version of this report at the 2024 Annual Meeting, comments were heard during the online forum and Reference Committee B hearing regarding the recommendations on disclosure of use of AI to physicians and, ultimately, to patients. Commentors raised concerns that transparency regarding the use of AI would be overly burdensome to health systems and hospitals deploying AI and that transparency would entail disclosure of use of algorithms in any instance, including those used in EHRs, those for administrative purposes, and others that do not directly impact physician and patient decision-making. There were also concerns that the recommendations around transparency were akin to calling for burdensome informed consent for the use of AI and that disclosure of the use of AI to patients risks damaging the patient-physician relationship.

For the purposes of this report and its recommendations, “disclosure” should be understood to mean communicating to physicians or patients about the use of AI-enabled systems or technologies that directly impact medical decision making and treatment recommendations at the point of care.

Documentation involves recording of an AI system’s design, development, and decision-making processes. This is primarily intended for internal teams, regulators, and researchers, and to enhance understanding, maintenance, and improvement of AI systems. Disclosure, on the other hand, refers to communicating essential information about AI systems to external stakeholders, e.g., end users. Disclosure focuses on essential aspects and, in this context, denotes the “when” and not the “what” to disclose. Concise and targeted disclosure is easier to disseminate and understand than comprehensive and nuanced details. It is important to note that disclosure should not be confused with informed consent. Informed consent is multifaceted, including benefits and drawbacks depending on its implementation and context of use. It can introduce burdens such as time-consuming paperwork, complex legal language, and potential delays in receiving care or participating in research. These burdens can deter individuals from providing their medical information or utilizing AI. Disclosure, on the other hand, is a form of transparency that builds trust, ensures accountability, supports risk management efforts, and informs users about the AI system’s behavior without adding undue burden. Together, documentation and disclosure foster a comprehensive approach to AI transparency, addressing both internal and external needs.

The National Institute of Standards and Technology (NIST) frames AI risk management as a path to minimize potential negative impacts of AI systems, such as threats to civil liberties and rights, while also providing opportunities to maximize positive impacts. NIST adopted the International Organization for Standardization’s (ISO) position that transparency and ethical behavior are a social responsibility when decisions and activities impact society and the environment (ISO 26000:2010).⁷ NIST further states that addressing, documenting, disclosing, and managing AI risks and potential negative impacts effectively can lead to more trustworthy AI systems.⁸ Moreover, multiple medical specialty organizations, including the American College of Radiology (ACR) and the American College of Physicians (ACP) support disclosure.

ACR’s *Ethics of AI in Radiology* states that, for a model to be transparent, it must be both visible and understandable to outsiders, including patients. A practical approach to achieving transparency is through clear disclosure. Further, when AI is the main point of contact in health care, it is ACR’s position that patients



should be clearly informed that they are interacting with an AI tool. In its 2024 position paper *AI in the Provision of Health Care*, ACP emphasizes that AI transparency is important for patients as well as physicians and other clinicians. Even if patients are not, at present, explicitly informed of all the ways technology is involved in their care—for example, they may or may not be told about computer-assisted electrocardiogram or mammography interpretation—ACP asserts that, due to the novelty of AI and its potential for significant clinical impacts, honesty and transparency about its use are crucial.^{9,10}

Given that transparency and disclosure are not static, their practicality or applicability are dependent on the situation and environment. ACP, for example, recognizes that transparency with patients about the integration of AI into certain devices may be reasonably feasible. In these cases, disclosure is more attuned to AI used in medical treatment and decision making and not the underlying algorithm, which could be overly burdensome. Algorithms are not new in health care; they are widely used, and many have become the standard of care. On the other hand, transparency with patients about AI integration into EHR systems and other common sources of information may be less feasible, especially given that physicians are often not made aware of the integration.

Nevertheless, as NIST notes, meaningful transparency should provide access to appropriate levels of information based on the stage of the AI lifecycle and tailored to the role or knowledge of individuals interacting with or using the AI system.

Ethical Considerations for Disclosure of the Use of AI that Impacts Clinical Decision Making

The AMA was founded in part to establish the world's first national code of medical ethics. Opinions included in the AMA Code of Medical Ethics aim to address issues and challenges confronting the medical profession and represent AMA policy. Promoting adherence to the professional standards promulgated in the Code is essential to preserving patient trust and public confidence in the medical profession.

Included as part of the Code are the ethical responsibilities of physicians as they relate to transparency in health care.¹¹ The Code states that “[p]atients must rely on their physicians to provide information that patients reasonably would want to know to make informed, well-considered decisions about their health care,” and that “physicians have an obligation to inform patients about...tools that influence treatment recommendations and care.” The Code additionally states that, where treatment recommendations are concerned, “[p]atients have the right to receive information and ask questions about recommended treatments so that they can make well-considered decisions about care. Successful communication in the patient-physician relationship fosters trust and supports shared decision-making.”¹²

Physician use of AI is not an exception to the Code, nor is there separate ethical guidance for the use of AI at this time. The Code suggests that communication to physicians and patients about the use of AI that may directly impact medical decision making and treatment recommendations is in line with prevailing ethical principles. It may be particularly important seeing that, at this time, patients are expressing broad discomfort with the notion of their physicians relying on AI in their own health care.¹³ To best foster trust, both between physicians and developers/deployers, and between physicians and patients, use of AI that may directly impact medical decision making should be communicated to parties involved in that decision making.

Intersections between Physician Liability and Disclosure of the Use of AI in Clinical Practice

AI transparency, both in disclosing use to physicians and to patients as well as disclosure of key information to physicians regarding the tools by AI developers and deployers, is an essential component to managing risk and



potentially reducing physician liability resulting from the use of AI. As with hardware devices and other medical products, physicians are ultimately responsible for the appropriate selection and use of devices, diagnostics, and other products in clinical practice. Claims of lack of knowledge or understanding of the system in question will likely weaken a defense in any medical liability case involving AI-enabled technology. Therefore, it is essential that both physicians and patients are aware when AI impacts clinical decision-making and understand how it factors into the process. This ensures that accountability and liability can be appropriately assigned when poor AI performance leads to poor patient outcomes, or where the AI-technology is itself defective (similar to when a device or diagnostic product is defective).

Required Disclosures by Health Care Augmented Intelligence-Enabled Systems and Technologies

Along with significant opportunity to improve patient care, all new technologies in health care will likely present certain risks and limitations that physicians must carefully navigate during the early stages of clinical implementation of these new systems and tools. AI-enabled tools are no different and are perhaps more challenging than other advances as they present novel and complex questions and risks. To best mitigate these risks, it is critical that physicians understand AI-driven technologies and have access to certain information about the AI tool or system being considered, including how it was trained and validated, so that they can assess the quality, performance, equity, and utility of the tool to the best of their ability. This information may also establish a set of baseline metrics for comparing AI tools. Transparency and explainability regarding the design, development, and deployment processes should be mandated by law where feasible, including potential sources of inequity in problem formulation, inputs, and implementation. Additionally, sufficient detail should be disclosed to allow physicians to determine whether a given AI-enabled tool would reasonably apply to the individual patient they are treating.

Physicians should be aware and understand that, where they utilize AI-enabled tools and systems without transparency provided by the AI developer, their risks of liability for reliance on that AI will likely increase. The need for full transparency is greatest where AI-enabled systems have greater impact on direct patient care, such as by AI-enabled medical devices, clinical decision support, and interaction with AI-driven chatbots. Transparency needs may be somewhat lower where AI is utilized for primarily administrative, practice-management functions.

While some of this information may be provided in labeling for FDA cleared and approved medical devices, the labeling requirements for such devices have not been specifically tailored to clearly convey information about these new types of devices. Updated guidance for FDA-regulated medical devices is needed to provide this critical information. Congress should consider actions to ensure appropriate authorities exist to require appropriate information to be provided to users of AI so that they can best evaluate the technology to determine reported performance, intended use, intended population, and appropriateness for the task. Developers and vendors should provide this information about their products, and physicians and other purchasers should consider this information when selecting the AI tools they use.

Generative AI

Generative AI is a type of AI that can recognize, summarize, translate, predict, and generate text and other content based on knowledge gained from large datasets. Generative AI tools are finding an increasing number of uses in health care, including assistance with administrative functions, such as generating office notes, responding to documentation requests, and generating patient messages. Additionally, there has been increasing discussion about clinical applications of generative AI, including use as clinical decision support to provide differential diagnoses, early detection and intervention, and to assist in treatment planning. While



generative AI tools show tremendous promise to make a significant contribution to health care, there are a number of risks and limitations to consider when using these tools in a clinical setting or for direct patient care. These risks are especially important to consider for clinical applications that may impact clinical decision-making and treatment planning where risks to patients are higher.

Given that there are no regulations or generally accepted standards or frameworks to govern the design, development, and deployment of generative AI, consideration and mitigation of the significant risks are paramount. To manage risk, health care organizations should develop and adopt appropriate policies that anticipate and minimize negative impacts. Physicians who consider utilizing a generative AI-based tool in their practice should ensure that all practice staff are educated on the risks and limitations, including patient privacy concerns, and should have appropriate governance policies in place for its use prior to adoption. Also, as raised in Resolution 206-I-23, physicians should be encouraged to educate their patients about the benefits and risks of using AI-based tools, such as LLMs, for information about health care conditions, treatment options, or the type of health care professionals who have the education, training, and qualifications to treat a particular condition. Patients and physicians should be aware that chatbots powered by LLMs/generative AI could provide inaccurate, misleading, or unreliable information and recommendations. This principle is incorporated in the recommendations in this report and current AMA Policy [H-480.940](#), "Augmented Intelligence in Health Care."

Liability

The question of physician liability for use of AI-enabled technologies presents novel and complex legal questions and poses risks to the successful clinical integration of AI-enabled technologies. It is also one of the most serious concerns for physicians when considering integration of AI into their practice. Concerns also arise for employed physicians who feel they may have no choice but to utilize the AI, should hospitals or health systems mandate its use or utilize an EHR system that incorporates AI-based applications as standard.

The challenge for physicians regarding questions of liability for use of AI is that there is not yet any clear legal standard for determining liability. While there are clear standards for physician liability generally and for medical device liability, AI presents novel and potentially complex legal questions. When AI has suggested a diagnosis, the question of how appropriate it is for a physician to rely on that result is yet to be determined and will likely continue to evolve as AI improves. Ultimately the "standard of care" will help guide physician liability. It is expected that, as it improves over time, AI will be incorporated into what is likely to be specialty-specific standards of care. However, until that occurs, AI-transparency is of critical importance and physicians will need to be diligent in ensuring that they engage with AI tools where performance has been validated in their practice setting.

As AI continues to evolve, there may ultimately be questions regarding liability when physicians fail to use AI and rely only on their professional judgment. Again, this question may ultimately turn on what evolves to be considered the standard of care.

It should be noted that, when using AI, physicians will still be subject to general legal theories regarding medical liability. Negligent selection of an AI tool, including using tools outside their intended use or intended population, or choosing a tool where there is no evidence of clinical validation, could be decisions that expose a physician to a liability claim.

Data Privacy and Augmented Intelligence

Data privacy is highly relevant to AI development, implementation, and use. The AMA is deeply invested in ensuring individual patient rights and protections from discrimination remain intact, that these assurances are guaranteed, and that the responsibility rests with the data holders. AI development, training, and use requires assembling large collections of health data. AI machine learning is data hungry; it requires massive amounts of data to function properly. Increasingly, more electronic health records are interoperable across the health care system and, therefore, are accessible by AI trained or deployed in medical settings. AI developers may enter into legal arrangements (e.g., business associate agreements) that bring them under the HIPAA Privacy and Security Rules. However, physicians and medical providers are often seen as the sole responsible parties, expected to bear the burden of data protection. This position is not sustainable. Given the newness of AI and its potential for clinically significant effects on care, equitable accountability must be established. While some uses of AI in health care, such as research, are not allowed by HIPAA absent patient authorization, the applicability of other HIPAA privacy protections to AI use is not as clear and HIPAA cannot protect patients from the “black box” nature of AI which makes the use of data opaque. AI system outputs may also include inferences that reveal personal data or previously confidential details about individuals. This can result in a lack of accountability and trust and exacerbate data privacy concerns. Often, AI developers and implementers are themselves unaware of exactly how their products use information to make recommendations.

It is unlikely that physicians or patients will have any clear insight into a generative AI tool’s conformance to state or federal data privacy laws. LLMs are trained on data scraped from the web and other digital sources, including one well-documented instance where HIPAA privacy protections were violated.¹⁴ Few, if any, controls are available to help users protect the data they voluntarily enter in a chatbot query. For instance, there are often no mechanisms in place for users to request data deletion or ensure that their inputs are not stored or used for future model training. While tools designed for medical use should align with HIPAA, many “HIPAA-compliant” generative tools rely on antiquated notions of deidentification, i.e., stripping data of personal information. With today’s advances in computing power, data can easily be reidentified. Rather than aiming to make LLMs compliant with HIPAA, all health care AI-powered generative tools should be designed from the ground up with data privacy in mind. Additionally, some companies have intentionally misled the public and end-users by labeling their software tools as “HIPAA compliant”, when the entity itself was not a covered entity or business associate and therefore not subject to HIPAA Privacy Rules.

[The AMA’s Privacy Principles](#) were designed to provide individuals with rights and protections and shift the responsibility for privacy to third-party data holders. While the Principles are broadly applicable to all AI developers, e.g., entities should only collect the minimum amount of information needed for a particular purpose, the unique nature of LLMs and generative AI warrant special emphasis on entity responsibility and user education.

Augmented Intelligence Cybersecurity

Data privacy relies on strong data security measures. There is growing concern that cyber criminals will use AI to attack health care organizations. AI poses new threats to health IT operations. AI-operated ransomware and AI-operated malware can be targeted to infiltrate health IT systems and automatically exploit vulnerabilities. Attackers using ChatGPT can craft convincing or authentic emails and use phishing techniques that entice people to click on links—giving them access to the entire electronic health record system.



AI is particularly sensitive to the quality of data. Data poisoning is the introduction of “bad” data into an AI training set, affecting the model’s output. AI requires large sets of data to build logic and patterns used in clinical decision-making. Protecting this source data is critical. Threat actors could also introduce input data that compromises the overall function of the AI tool. Failure to secure and validate these inputs, and corresponding data, can contaminate AI models—resulting in patient harm.

Because stringent privacy protections and higher data quality standards might slow model development, there could be a tendency to forgo essential data privacy and security precautions. However, strengthening AI systems against cybersecurity threats is crucial to their reliability, resiliency, and safety.

Mis- and Disinformation Propagated by AI

Health mis- and disinformation poses a serious threat to public health. It can cause significant confusion among patients, increase patient mistrust in science and in physicians, result in patients making decisions that cause themselves harm, and undermine the ability to manage public health threats. The dissemination of mis- and disinformation in health care significantly increased during the COVID-19 pandemic and shows no signs of abating. Whether intentionally or unintentionally, AI, in particular generative AI, runs the risk of contributing to the creation and dissemination of scientific and medical mis- and disinformation. Physicians, staff, and patients must all be aware of the risks of mis- and disinformation when engaging with generative and other forms of AI.

Generative AI can propagate mis- and disinformation in several ways. It can engage in the unintentional or intentional creation of incorrect information on its own. The risk of generative AI “hallucinating,” “confabulating,” or otherwise fabricating information in response to a user-generated query has been well documented.^{15,16} Notably, tools such as ChatGPT have shown a not-uncommon tendency to falsify references cited in response to these queries. Generative AI tools have demonstrated the ability to generate fraudulent scientific/medical literature.¹⁷ They are also capable of plagiarizing, falsifying, or misrepresenting data in ways that could compromise research integrity. Additionally, retracted papers may have the ability to continue to impact the content generated by LLM-based tools, potentially leading to dissemination or inaccurate or otherwise discredited information.

AI can also be responsible for intentionally or unintentionally disseminating false information or intentional misinformation, which can happen when that information is used as part of the training data set for the model, used as a reference in a response to a query, or otherwise presented to a user in a query response. Information presented to users by generative AI models can be extremely convincing, with the users potentially having little reason to doubt what is presented.

There is little opportunity currently to regulate AI’s role in propagation of health mis- and disinformation under current oversight structures. The FTC is the most likely agency to take action against mis- and disinformation, as it has broad authorities to regulate unfair and deceptive business practices. However, as discussed above, the FTC will require additional resources to appropriately regulate the role of AI in propagating mis- and disinformation. Regulation of mis- and disinformation is further complicated by the intersection of false and misleading information with free speech rights guaranteed by the First Amendment.

It is critical that the health care industry and health care stakeholders broadly take action to limit AI’s ability to create or disseminate mis- or disinformation. Developers of AI should be accountable for their product creating or disseminating false information and should have mechanisms in place to allow for reporting of mis- and disinformation. Federal regulations should seek to eliminate the propagation of mis- and disinformation by



AI-enabled tools. Ethical principles for use of AI in medical and scientific research should be in place to ensure continued research integrity. Journals should ensure that they have clear guidelines in place to regulate the use of AI in scientific publications that include documenting and detailing the use of AI in research and to exclude the use of AI systems as authors. Policies should also detail the responsibility of authors to validate the veracity of any text generated by AI. (See Policy [H-480.935](#), Assessing the Potentially Dangerous Intersection Between AI and Misinformation).

Payor Use of Augmented Intelligence in Automated Decision-Making

Payors and health plans are increasingly using AI and algorithm-based decision-making in an automated fashion to determine coverage limits, make claim determinations, and engage in benefit design. Payors should leverage automated decision-making systems that improve or enhance efficiencies in coverage and payment automation, facilitate administrative simplification, and reduce workflow burdens. While the use of these systems can create efficiencies such as speeding up prior authorization and cutting down on paperwork, there is concern these systems are not being designed or supervised effectively creating access barriers for patients and limiting essential benefits.

Increasingly, evidence indicates that payors are using automated decision-making systems to deny care more rapidly, often with little or no human review. This manifests in the form of increased denials, stricter coverage limitations, and constrained benefit offerings. For example, a payor allowed an automated system to cut off insurance payments for Medicare Advantage patients struggling to recover from severe diseases, forcing them to forgo care or pay out of pocket. In some instances, payors instantly reject claims on medical grounds without opening or reviewing the patient's medical record. There is also a lack of transparency in the development of automated decision-making systems. Rather than payors making determinations based on individualized patient care needs, reports show that decisions are based on algorithms developed using average or "similar patients" pulled from a database. Models that rely on generalized, historical data can also perpetuate biases leading to discriminatory practices or less inclusive coverage.^{18,19,20,21}

While AI can be used inappropriately by payors with severe detrimental outcomes to patients, it can also serve to reduce administrative burdens on physicians, providing the ability to more easily submit prior authorization and documentation requests in standardized forms that require less physician and staff time. Given the significant burden placed on physicians and administrative staff by prior authorization requests, AI could provide much needed relief and help to increase professional satisfaction among health care professionals. With clear guidelines, AI-enabled decision-making systems may also be appropriate for use in some lower-risk, less complex care decisions.

While payor use of AI in well-defined situations with clear guidelines has the potential to reduce burdens and benefit physician practices, new regulatory or legislative action is necessary to ensure that automated decision-making systems do not reduce needed care, nor systematically withhold care from specific groups. Steps should be taken to ensure that these systems do not override clinical judgment. Patients and physicians should be informed and empowered to question a payor's automated decision-making. There should be stronger regulatory oversight, transparency, and audits when payors use these systems for coverage, claim determinations, and benefit design. [See Policy [D-480.956](#), "Use of Augmented Intelligence for Prior Authorization;" and Policy [H-320.939](#), "Prior Authorization and Utilization Management Reform"]

CONCLUSION

As the number of AI-enabled health care tools and systems continue to grow, these technologies must be designed, developed, and deployed in a manner that is ethical, equitable, responsible, accurate, and transparent. In line with AMA Policy [H-480-935](#) and Resolution 206-I-23, this report highlights some of the potential benefits and risks to the medical profession and patients of LLMs (e.g., GPTs) and other AI-generated medical decision-making tools, and recommends adoption of policy to help inform patient and physician education and guide engagement with this new technology, as well as position the AMA to advocate for governance policies that help to ensure that risks arising from AI are mitigated to the greatest extent possible.

RECOMMENDATION

The AMA House of Delegates adopted the following policy on November 12, 2024:

Augmented intelligence development, deployment, and use in health care

- 1) General Governance
 - a) Health care AI must be designed, developed, and deployed in a manner which is ethical, equitable, responsible, accurate, transparent, and evidence-based.
 - b) Use of AI in health care delivery requires clear national governance policies to regulate its adoption and utilization, ensuring patient safety, and mitigating inequities. Development of national governance policies should include interdepartmental and interagency collaboration.
 - c) Compliance with national governance policies is necessary to develop AI in an ethical and responsible manner to ensure patient safety, quality, and continued access to care. Voluntary agreements or voluntary compliance is not sufficient.
 - d) AI systems should be developed and evaluated with a specific focus on mitigating bias and promoting health equity, ensuring that the deployment of these technologies does not exacerbate existing disparities in health care access, treatment, or outcomes.
 - e) Health care AI requires a risk-based approach where the level of scrutiny, validation, and oversight should be proportionate to the overall potential of disparate harm and consequences the AI system might introduce. [See also Augmented Intelligence in Health Care [H-480.939](#) at (1)]
 - f) AI risk management should minimize potential negative impacts of health care AI systems while providing opportunities to maximize positive impacts.
 - g) Clinical decisions influenced by AI must be made with specified qualified human intervention points during the decision-making process. A qualified human is defined as a licensed physician with the necessary qualifications and training to independently provide the same medical service without the aid of AI. As the potential for patient harm increases, the point in time when a physician should utilize their clinical judgment to interpret or act on an AI recommendation should occur earlier in the care plan. With few exceptions, there generally should be a human in the loop when it comes to medical decision making capable of intervening or overriding the output of an AI model.
 - h) Health care practices and institutions should not utilize AI systems or technologies that introduce overall or disparate risk that is beyond their capabilities to mitigate. Implementation and utilization of AI should avoid exacerbating clinician burden and should be designed and deployed in harmony with the clinical workflow and, in institutional settings, consistent with AMA Policy H-225.940 - Augmented Intelligence and Organized Medical Staff.
 - i) Medical specialty societies, clinical experts, and informaticists are best positioned and should identify the most appropriate uses of AI-enabled technologies relevant to their clinical expertise and set the standards for AI use in their specific domain. [See Augmented Intelligence in Health Care [H-480.940](#) at (2)]



2) When to Disclose: Transparency in Use of Augmented Intelligence-Enabled Systems and Technologies That Impact Medical Decision Making at the Point of Care

- a) Decisions regarding transparency and disclosure of the use of AI should be based upon a risk- and impact-based approach that considers the unique circumstance of AI and its use case. The need for transparency and disclosure is greater where the performance of an AI-enabled technology has a greater risk of causing harm to a patient.
 - i) AI disclosure should align and meet ethical standards or norms.
 - ii) Transparency requirements should be designed to meet the needs of the end users. Documentation and disclosure should enhance patient and physician knowledge without increasing administrative burden.
 - iii) When AI is used in a manner which impacts access to care or impacts medical decision making at the point of care, that use of AI should be disclosed and documented to both physicians and/or patients in a culturally and linguistically appropriate manner. The opportunity for a patient or their caregiver to request additional review from a licensed clinician should be made available upon request.
 - iv) When AI is used in a manner which directly impacts patient care, access to care, medical decision making, or the medical record, that use of AI should be documented in the medical record.
- b) AI tools or systems cannot augment, create, or otherwise generate records, communications, or other content on behalf of a physician without that physician's consent and final review.
- c) When AI or other algorithmic-based systems or programs are utilized in ways that impact patient access to care, such as by payors to make claims determinations or set coverage limitations, use of those systems or programs must be disclosed to impacted parties.
- d) The use of AI-enabled technologies by hospitals, health systems, physician practices, or other entities, where patients engage directly with AI, should be clearly disclosed to patients at the beginning of the encounter or interaction with the AI-enabled technology. Where patient-facing content is generated by AI, the use of AI in generating that content should be disclosed or otherwise noted within the content.

3) What to Disclose: Required Disclosures by Health Care Augmented Intelligence-Enabled Systems and Technologies

- a) When AI-enabled systems and technologies are utilized in health care, the following information should be disclosed by the AI developer to allow the purchaser and/or user (physician) to appropriately evaluate the system or technology prior to purchase or utilization:
 - i) Regulatory approval status.
 - ii) Applicable consensus standards and clinical guidelines utilized in design, development, deployment, and continued use of the technology.
 - iii) Clear description of problem formulation and intended use accompanied by clear and detailed instructions for use.
 - iv) Intended population and intended practice setting.
 - v) Clear description of any limitations or risks for use, including possible disparate impact.
 - vi) Description of how impacted populations were engaged during the AI lifecycle.
 - vii) Detailed information regarding data used to train the model:
 - (1) Data provenance.
 - (2) Data size and completeness.
 - (3) Data timeframes.

- (4) Data diversity.
- (5) Data labeling accuracy.
- viii) Validation Data/Information and evidence of:
 - (1) Clinical expert validation in intended population and practice setting and intended clinical outcomes.
 - (2) Constraint to evidence-based outcomes and mitigation of "hallucination"/"confabulation" or other output error.
 - (3) Algorithmic validation.
 - (4) External validation processes for ongoing evaluation of the model performance, e.g., accounting for AI model drift and degradation.
 - (5) Comprehensiveness of data and steps taken to mitigate biased outcomes.
 - (6) Other relevant performance characteristics, including but not limited to performance characteristics at peer institutions/similar practice settings.
 - (7) Post-market surveillance activities aimed at ensuring continued safety, performance, and equity.
- ix) Data Use Policy:
 - (1) Privacy.
 - (2) Security.
 - (3) Special considerations for protected populations or groups put at increased risk.
- x) Information regarding maintenance of the algorithm, including any use of active patient data for ongoing training.
- xi) Disclosures regarding the composition of design and development team, including diversity and conflicts of interest, and points of physician involvement and review.

b) Purchasers and/or users (physicians) should carefully consider whether or not to engage with AI-enabled health care technologies if this information is not disclosed by the developer. As the risk of AI being incorrect increases risks to patients (such as with clinical applications of AI that impact medical decision making), disclosure of this information becomes increasingly important. [See also Augmented Intelligence in Health Care [H-480.939](#)]

4) Generative Augmented Intelligence

- a) Generative AI should: (a) only be used where appropriate policies are in place within the practice or other health care organization to govern its use and help mitigate associated risks; and (b) follow applicable state and federal laws and regulations (e.g., HIPAA-compliant Business Associate Agreement).
- b) Appropriate governance policies should be developed by health care organizations and account for and mitigate risks of:
 - i) Incorrect or falsified responses; lack of ability to readily verify the accuracy of responses or the sources used to generate the response.
 - ii) Training data set limitations that could result in responses that are out of date or otherwise incomplete or inaccurate for all patients or specific populations.
 - iii) Lack of regulatory or clinical oversight to ensure performance of the tool.
 - iv) Bias, discrimination, promotion of stereotypes, and disparate impacts on access or outcomes.
 - v) Data privacy.
 - vi) Cybersecurity.
 - vii) Physician liability associated with the use of generative AI tools.
- c) Health care organizations should work with their AI and other health information technology (health IT) system developers to implement rigorous data validation and verification protocols to ensure that only



accurate, comprehensive, and bias managed datasets inform generative AI models, thereby safeguarding equitable patient care and medical outcomes. [See Augmented Intelligence in Health Care [H-480.940](#) at (3)(d)]

- d) Use of generative AI should incorporate physician and staff education about the appropriate use, risks, and benefits of engaging with generative AI. Additionally, physicians and healthcare organizations should engage with generative AI tools only when adequate information regarding the product is provided to physicians and other users by the developers of those tools.
- e) Clinicians should be aware of the risks of patients engaging with generative AI products that produce inaccurate or harmful medical information (e.g., patients asking chatbots about symptoms) and should be prepared to counsel patients on the limitations of AI-driven medical advice.
- f) Data and prompts contributed by users should primarily be used by developers to improve the user experience and AI tool quality and not simply increase the AI tool's market value or revenue generating potential.

5) Physician Liability for Use of Augmented Intelligence-Enabled Technologies

- a) Current AMA policy states that liability and incentives should be aligned so that the individual(s) or entity(ies) best positioned to know the AI system risks and best positioned to avert or mitigate harm do so through design, development, validation, and implementation. [See Augmented Intelligence in Health Care [H-480.939](#)]
 - i) Where a mandated use of AI systems prevents mitigation of risk and harm, the individual or entity issuing the mandate must be assigned all applicable liability.
 - ii) Developers of autonomous AI systems with clinical applications (screening, diagnosis, treatment) are in the best position to manage issues of liability arising directly from system failure or misdiagnosis and must accept this liability with measures such as maintaining appropriate medical liability insurance and in their agreements with users.
 - iii) Health care AI systems that are subject to non-disclosure agreements concerning flaws, malfunctions, or patient harm (referred to as gag clauses) must not be covered or paid and the party initiating or enforcing the gag clause assumes liability for any harm.
- b) When physicians do not know or have reason to know that there are concerns about the quality and safety of an AI-enabled technology, they should not be held liable for the performance of the technology in question.
- c) Liability protections for physicians using AI-enabled technologies should align with both current and future AMA medical liability reform policies.

6) Data Privacy and Augmented Intelligence

- a) Entity Responsibility:
 - i) Entities, e.g., AI developers, should make information available about the intended use of generative AI in health care and identify the purpose of its use. Individuals should know how their data will be used or reused, and the potential risks and benefits.
 - ii) Individuals should have the right to opt-out, update, or request deletion of their data from generative AI tools. These rights should encompass AI training data and disclosure to other users of the tool.
 - iii) Generative AI tools should not reverse engineer, reconstruct, or reidentify an individual's originally identifiable data or use identifiable data for nonpermitted uses, e.g., when data are permitted to conduct quality and safety evaluations. Preventive measures should include both legal frameworks and data model protections, e.g., secure enclaves, federated learning, and differential privacy.



b) User Education:

- i) Users should be provided with training specifically on generative AI. Education should address:
 - (1) Legal, ethical, and equity considerations.
 - (2) Risks such as data breaches and re-identification.
 - (3) Potential pitfalls of inputting sensitive and personal data.
 - (4) The importance of transparency with patients regarding the use of generative AI and their data.

[See [H-480.940](#), Augmented Intelligence in Health Care, at (4) and (5)]

7) Augmented Intelligence Cybersecurity

- a) AI systems must have strong protections against input manipulation and malicious attacks.
- b) Entities developing or deploying health care AI should regularly monitor for anomalies or performance deviations, comparing AI outputs against known and normal behavior.
- c) Independent of an entity's legal responsibility to notify a health care provider or organization of a data breach, that entity should also act diligently in identifying and notifying the individuals themselves of breaches that impact their personal information.
- d) Users should be provided education on AI cybersecurity fundamentals, including specific cybersecurity risks that AI systems can face, evolving tactics of AI cyber attackers, and the user's role in mitigating threats and reporting suspicious AI behavior or outputs.

8) Mitigating Misinformation in AI-Enabled Technologies

- a) AI developers should ensure transparency and accountability by disclosing how their models are trained and the sources of their training data. Clear disclosures are necessary to build trust in the accuracy and reliability of the information produced by AI systems.
- b) Algorithms should be developed to detect and flag potentially false and misleading content before it is widely disseminated.
- c) Developers of AI should have mechanisms in place to allow for reporting of mis- and disinformation generated or propagated by AI-enabled systems.
- d) Developers of AI systems should be guided by policies that emphasize rigorous validation and accountability for the content their tools generate, and, consistent with AMA Policy [H-480.939\(7\)](#), are in the best position to manage issues of liability arising directly from system failure or misdiagnosis and must accept this liability with measures such as maintaining appropriate medical liability insurance and in their agreements with users.
- e) Academic publications and journals should establish clear guidelines to regulate the use of AI in manuscript submissions. These guidelines should include requiring the disclosure that AI was used in research methods and data collection, requiring the exclusion of AI systems as authors, and should outline the responsibility of the authors to validate the veracity of any referenced content generated by AI.
- f) Education programs are needed to enhance digital literacy, helping individuals critically assess the information they encounter online, particularly in the medical field where mis- and disinformation can have severe consequences.

9) Payor Use of Augmented Intelligence and Automated Decision-Making Systems

- a) Use of automated decision-making systems that determine coverage limits, make claim determinations, and engage in benefit design should be publicly reported, based on easily accessible evidence-based clinical guidelines (as opposed to proprietary payor criteria), and disclosed to both patients and their physician in a way that is easy to understand.



- b) Payors should only use automated decision-making systems to improve or enhance efficiencies in coverage and payment automation, facilitate administrative simplification, and reduce workflow burdens. Automated decision-making systems should never create or exacerbate overall or disparate access barriers to needed benefits by increasing denials, coverage limitations, or limiting benefit offerings. Use of automated decision-making systems should not replace the individualized assessment of a patient's specific medical and social circumstances and payors' use of such systems should allow for flexibility to override automated decisions. Payors should always make determinations based on particular patient care needs and not base decisions on algorithms developed on "similar" or "like" patients.
- c) Payors using automated decision-making systems should disclose information about any algorithm training and reference data, including where data were sourced and attributes about individuals contained within the training data set (e.g., age, race, gender). Payors should provide clear evidence that their systems do not discriminate, increase inequities, and that protections are in place to mitigate bias.
- d) Payors using automated decision-making systems should identify and cite peer-reviewed studies assessing the system's accuracy measured against the outcomes of patients and the validity of the system's predictions.
- e) Any automated decision-making system recommendation that indicates limitations or denials of care, at both the initial review and appeal levels, should be automatically referred for review to a physician (a) possessing a current and valid non-restricted license to practice medicine in the state in which the proposed services would be provided if authorized and (b) be of the same specialty as the physician who typically manages the medical condition or disease or provides the health care service involved in the request prior to issuance of any final determination. Prior to issuing an adverse determination, the treating physician must have the opportunity to discuss the medical necessity of the care directly with the physician who will be responsible for determining if the care is authorized.
- f) Individuals impacted by a payor's automated decision-making system, including patients and their physicians, must have access to all relevant information (including the coverage criteria, results that led to the coverage determination, and clinical guidelines used).
- g) Payors using automated decision-making systems should be required to engage in regular system audits to ensure use of the system is not increasing overall or disparate claims denials or coverage limitations, or otherwise decreasing access to care. Payors using automated decision-making systems should make statistics regarding systems' approval, denial, and appeal rates available on their website (or another publicly available website) in a readily accessible format with patient population demographics to report and contextualize equity implications of automated decisions. Insurance regulators should consider requiring reporting of payor use of automated decision-making systems so that they can be monitored for negative and disparate impacts on access to care. Payor use of automated decision-making systems must conform to all relevant state and federal laws.

(New HOD Policy)

Fiscal Note: Less than \$500.

REFERENCES

- ¹ <https://www.ama-assn.org/system/files/ama-patient-data-privacy-survey-results.pdf>.
- ² <https://www.ama-assn.org/sites/ama-assn.org/files/corp/media-browser/public/government/advocacy/medical-cybersecurity-findings.pdf>.
- ³ https://www.healthit.gov/sites/default/files/page/2024-01/DSI_H1%20Final%20Rule%20Presentation_508.pdf.
- ⁴ <https://www.propublica.org/article/cigna-health-insurance-denials-pwdx-congress-investigation#:~:text=The%20letter%20follows%20an%20investigation,PDX%20system%2C%20spending%20an%20average>.
- ⁵ AI systems should enhance the patient experience of care and outcomes, improve population health, reduce overall costs for the health care system while increasing value, and support the professional satisfaction of physicians and the health care team.
- ⁶ For example, the 21st Century Cures Act includes several exemptions to FDA's oversight, such as software intended for administrative support of a health care facility, maintaining or encouraging a healthy lifestyle (and is unrelated to the diagnosis, cure, mitigation, prevention, or treatment of a disease or condition), is intended to be used as electronic patient records, is intended for transferring, storing, converting formats, or displaying data or results, and otherwise does not meet the definition of a medical device under the Federal Food, Drug, and Cosmetic Act.
- ⁷ <https://www.iso.org/obp/ui/en/#iso:std:iso:26000:ed-1:v1:en>.
- ⁸ <https://nvlpubs.nist.gov/nistpubs/ai/NIST.AI.100-1.pdf>.
- ⁹ <https://www.acr.org/-/media/ACR/Files/Informatics/Ethics-of-AI-in-Radiology-European-and-North-American-Multisociety-Statement-6-13-2019.pdf>
- ¹⁰ https://www.acpjournals.org/doi/10.7326/M24-0146?gl=1*e8j406*_gcl_au*MzQ3MzI3OTcuMTcyMTY1OTI1Mw_*_ga*NzM2MjYxNTIxLjE3MjE2NTkyNTM_*_ga_PM4F5HBGFQ*MTcyMTY1OTI1Mi4xLjAuMTcyMTY1OTI1Mi42MC4wljA.&_ga=2.223886232.2016482322.1721659253-736261521.1721659253.
- ¹¹ <https://code-medical-ethics.ama-assn.org/ethics-opinions/transparency-health-care#:~:text=Respect%20for%20patients%20autonomy%20is,influence%20treatment%20recommendations%20and%20care>.
- ¹² <https://code-medical-ethics.ama-assn.org/chapters/consent-communication-decision-making>.
- ¹³ <https://www.pewresearch.org/science/2023/02/22/60-of-americans-would-be-uncomfortable-with-provider-relying-on-ai-in-their-own-health-care/>.
- ¹⁴ Feathers, T, et al. "Facebook is receiving sensitive medical information from hospital websites. The Markup. June 16, 2022." <https://themarkup.org/pixel-hunt/2022/06/16/facebook-is-receiving-sensitive-medical-information-from-hospital-websites>.
- ¹⁵ <https://www.cnet.com/tech/hallucinations-why-ai-makes-stuff-up-and-whats-being-done-about-it/>.
- ¹⁶ <https://jamanetwork.com/journals/jamainternalmedicine/article-abstract/2808091>.
- ¹⁷ <https://www.jmir.org/2023/1/e46924>.
- ¹⁸ Obermeyer, Ziad, et al. "Dissecting racial bias in an algorithm used to manage the health of populations." *Science* 366.6464 (2019): 447-453. <https://www.science.org/doi/10.1126/science.aax2342>.
- ¹⁹ Ross, C., Herman, B. (2023) "Medicare Advantage Plans' Use of Artificial Intelligence Leads to More Denials." <https://www.statnews.com/2023/03/13/medicare-advantage-plans-denial-artificial-intelligence/> (Accessed September 14, 2023).
- ²⁰ Rucker, P., Miller, M., Armstrong, D. (2023). "Cigna and Its Algorithm Deny Some Claims for Genetic Testing, ProPublica Finds." <https://www.propublica.org/article/cigna-pwdx-medical-health-insurance-rejection-claims> (Accessed September 14, 2023).
- ²¹ Ross, C., Herman, B. (2023). "Medicare Advantage Algorithms Lead to Coverage Denials, With Big Implications for Patients." <https://www.statnews.com/2023/07/11/medicare-advantage-algorithm-navihealth-unitedhealth-insurance-coverage/> (Accessed September 14, 2023).