



AMA AI state advocacy and policy priorities

Issue Brief



About the AMA

The American Medical Association is the physicians' powerful ally in patient care. As the largest and only national association that convenes 190+ state and specialty medical societies and other critical stakeholders, the AMA represents physicians with a unified voice in courts, to legislative bodies, and to other key players in health care. The AMA leverages its strength by removing the obstacles that interfere with patient care, leading the charge to prevent chronic disease and confront public health crises, and driving the future of medicine to tackle the biggest challenges in health care. This work is led by AMA's mission: to promote the art and science of medicine and the betterment of public health.

For more information, visit [ama-assn.org](https://www.ama-assn.org).

THIS REPORT is for informational purposes only. It is not intended as medical, legal, financial, or consulting advice, or as a substitute for the advice of a physician, attorney, or other financial or consulting professional. It does not imply and is not intended as a promotion or endorsement by the AMA of any third-party organization, product, drug, or service.

Last updated 2024-12-01.

© 2024 American Medical Association. <https://www.ama-assn.org/terms-use>

Table of contents

Purpose & Background	4
Health plan use of AI	6
Transparency requirements for AI tools	8
Physician liability	11
Looking ahead	13
Endnotes	14

Purpose & Background

Purpose

Provide state legislators and medical societies with an overview of state AI policy activity in 2024 and the AMA's perspective on priority AI policy issues at the state level.

Background

The use of Augmented Intelligence (AI) is rapidly expanding across industries and influencing the operations and clinical practice of medicine. Physicians and other health care professionals have been using AI tools for decades, including rules-based algorithms that rely on human-defined logic to produce outputs and, more recently, generative AI models that can generate text, images, audio and video. There is a vast range of both clinical and administrative AI use cases in health care, such as summarizing medical notes, detecting and classifying the likelihood of future adverse events, and predicting patient volumes and associated staffing needs.

Until recently, there was very little movement on the federal and state legislative and regulatory fronts related to AI, but there was a notable uptick of activity in 2024. Policymakers are tasked with establishing a legal and regulatory framework for AI that balances the immense potential of these tools while minimizing risks. They are also attempting to calibrate the legislative need with the level of risk posed by specific AI use cases; for example, lower-risk use cases (e.g., patient scheduling) may not require the same degree of regulation and oversight as higher-risk tools (e.g., clinical decision support). Furthermore, policymakers must determine if new legislation is necessary, or if existing state laws (e.g., consumer protections, payer regulations, privacy laws, etc.) comprehensively regulate novel AI use cases. Furthermore, all this legislative activity is occurring against a backdrop of rapid advances in AI technology, the continued development of new AI use cases, and increased adoption by physicians.

In 2024, states introduced AI-related bills that would govern a diverse set of stakeholders, including states,

Key Definitions

Augmented Intelligence (AI): Evidence-based, computational methods and systems that enhance human capabilities and decision-making. The American Medical Association (AMA) uses the term “augmented intelligence” rather than “artificial intelligence” to reflect its perspective that artificial intelligence tools and services support rather than explicitly replace human decision-making.

Artificial Intelligence: The ability of computers to perform tasks that are typically associated with a rational human being—a quality that enables an entity to function appropriately and with foresight in its environment.

Machine Learning: A subtype of AI in which complex algorithms are trained to make predictions about future outcomes. Machine learning can be supervised or unsupervised.

Generative AI: Artificial intelligence systems that are capable of generating novel text, images, videos, or other outputs, typically based on foundation models. Foundation models are models trained on large datasets—and thus broadly applicable—and can be adjusted for specific applications. Typically used for generative artificial intelligence; LLMs are one type of foundation model.

Automated Decision-Making: A type of AI in which data and algorithms are used to make decisions without human intervention.

For additional definitions, please see the AMA's [Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care report](#).

organizations (e.g., hospitals, health insurance plans) and individuals (e.g., physicians) that use AI tools, as well as organizations that develop the tools. The majority of introduced bills would mandate states to study AI and/or establish AI working groups to inform future policy making. Several other bills focused on establishing transparency requirements across stakeholders who use or are affected by the outcomes of AI tools, prohibiting discrimination by AI tools, and/or regulating health plan use of AI. Of the approximately 20¹ bills that passed during 2024, the vast majority did not legislate specific AI activities, but rather established AI task forces to study the implications of AI from a policy perspective both broadly and specifically related to health and recommend approaches to mitigate potential risks. A few states passed transparency laws that other states are likely reviewing as potential models going forward: Utah [SB149](#) and California [AB3030](#) established certain disclosure requirements to consumers and patients, and Colorado [SB205](#) focused on transparency and disclosure between anyone who develops AI tools and those who “deploy” them (e.g., hospitals, physicians), as well as those who deploy them and the end user (e.g., physicians, patient). Notably, physicians can be developers, deployers, or end users. Notwithstanding substantial news coverage and a Senate [report](#) on AI use by health plans, only California passed a law ([SB1120](#)) detailing when a health plan can use AI to support medical necessity determinations.

The AMA is committed to ensuring the physicians’ perspective is heard when state lawmakers are considering legislative or regulatory approaches to the use of AI in health care. While a myriad of activity is occurring at the federal level, the AMA is also navigating the state legislative and regulatory landscape. Working alongside state medical associations and national specialty societies, this issue brief focuses on three key AI policy areas at the state level:

- **Health plan use of AI**—How health plans are permitted to use AI in eligibility determinations, medical necessity determinations, utilization management processes, and what/when physician intervention is critical;
- **Transparency**—What information must be communicated about AI use between relevant entities engaging in such use, such as between a physician and a patient; a physician and a hospital or state; and/or from a hospital, state, or AI developer to a physician; and,
- **Physician liability**—What liability, if any, does a physician potentially face for the use or non-use of AI tools in an administrative or clinical setting.

Below, this issue brief provides more detail and the AMA’s policy position on each of these priority policy areas:

¹ This number refers to AI bills that would implicate current (or potentially future) health care activities. There were many more AI bills introduced and passed related to other industries and use cases. For more information on 2024 health AI policy activity, please see Manatt Health’s Health AI Policy Tracker [here](#).

Health plan use of AI

Payors and health plans are increasingly leveraging AI, particularly in utilization management processes, medical necessity determinations, and eligibility or coverage decisions. While the use of AI tools can create efficiencies by automating processes and streamlining operations, the AMA is concerned that these tools are making automated decisions without considering the nuances of each individual patient's medical conditions and needs, increasing denials for medically necessary care, and creating access barriers (e.g., delays in care) for patients. In the past few years, there have been several well-publicized cases of health plans allegedly using AI to make medical necessity determinations to improperly deny patients coverage for services and/or deny prior authorization requests.^{1,2,3}

There are several key legislative and regulatory questions at stake. For example, when and how can AI be used by health plans—are certain activities acceptable for AI use while others are not? Furthermore, if AI is used to make medical determinations, must a physician review and confirm the AI tools' recommendations—and, if so, when or under what circumstances?

Only a small number of states introduced legislation during 2024 to address these questions. Oklahoma introduced language (SB1975) that would prohibit the use of AI to determine “who shall or shall not receive insurance coverage or the amount of such coverage.” Both Oklahoma (HB3577) and New York (AB9149) introduced language that would allow health plans to use AI in making positive coverage and eligibility determinations, but would require a physician to review any decision that would negatively impact coverage or access. California passed a bill (CA SB1120) which mandates that AI tools used for utilization review or utilization management decisions comply with a variety of requirements, including that the tool: does not supplant individualized health care provider decision-making; does not directly or indirectly discriminate; be fairly and equitably applied; and be open to audit for compliance. California's new law also specifies that an AI tool cannot deny, delay, or modify health care services based on medical necessity and that such decisions shall only be made by a physician or health care professional competent to evaluate the specific clinical issues involved.

In December 2023, the National Association of Insurance Commissioners developed a model bulletin for states to consider issuing to remind all insurance plans, including health insurance plans, that decisions or actions impacting consumers that were made or supported by AI tools must comply with applicable insurance laws and regulations; as of December 2024, 19 jurisdictions had adopted NAIC's model bulletin.^{4,5} And in February 2024, the Centers for Medicare & Medicaid Services (CMS) clarified that Medicare Advantage (MA) organizations may use AI and related technologies to assist in making coverage determinations, but the tools may not override standards related to medical necessity or rules related to how MA plans make coverage determinations decisions.⁶

AMA policy position

- **Payor use of AI for access and efficiency**
 - 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
 - Payors should never use automated decision-making systems to create or exacerbate access barriers

- **Role of physician and physician oversight in payor use of AI**
 - Any automated decision making tool that recommends limitations or denials of care 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) 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
 - Use of automated decision-making should not replace the individualized assessment of a patient's specific medical and social circumstances
- **Disclosures for payor use of AI**
 - 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
 - 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, and disclosed to both patients and their physician in a way that is easy to understand; patients and physicians should be informed and empowered to question a payor's automated decision-making
 - 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
- **Avoiding discrimination in payor use of AI**
 - Payors should provide clear evidence that their systems do not discriminate, increase inequities, and that protections are in place to mitigate bias
 - 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 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
- There should be stronger regulatory oversight, transparency, and audits when payors use these systems for coverage, claim determinations, and benefit design
- 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.

Transparency requirements for AI tools

AI in health care has been used for decades, from diagnostic imaging to predictive analytics. More recently, generative AI models have been used to generate personalized treatment plans based on patient medical history and convert open-ended clinical notes and data into standardized formats. Many AI models are often considered “black boxes,” in which users cannot discern the specific steps the algorithm takes to arrive at its final output. Additionally, generative AI models are known to “hallucinate,” generating outputs that misrepresent the training data, or are false but in many cases appear accurate. There are also concerns about “model drift” in AI tools, which describes the notion that a model may change over time (sometimes becoming less accurate) in ways that are hard to track—this may happen due to AI models learning and adapting over time, and/or because new data inputs can change how a model functions and learns. Similarly, there are concerns around algorithm or data bias, which can lead to inaccurate findings by the AI tool. Whether the AI tools are rules-based algorithms or generative models, the concept of transparency focuses on ensuring that appropriate information is shared between those who develop the tools, those who use the tools (e.g., physicians), and those who the tool impacts (e.g., patients).

Legislative and regulatory activity has largely focused on ensuring that each party that engages with an AI tool has the necessary information to assess the risks, identify appropriate use, and ideally garner overall trust in the AI tool. The key tension in these laws is how to appropriately balance sharing necessary information between parties without exacerbating administrative burden on all parties. For example, tools that pose limited risk to patients (e.g., appointment scheduling) and/or tools for which a physician independently validates the model’s output (e.g., ambient documentation tools that summarize a patient visit for physicians to review and validate) may not require the same level of disclosure to the patient as tools that operate more autonomously without physician oversight or review (e.g., a patient portal message that answers a patient’s question and was solely drafted by AI). Similarly, developers may be required to share more information about higher-risk models (e.g., those that inform treatment plans) with physicians or administrators than about models that support clinical operations (e.g., staffing predictions). From the physician perspective, there are two key questions when evaluating an AI transparency bill:

A note on patient privacy: Patient privacy is paramount. There are concerns that AI tools may: (1) unintentionally re-identify patient information, (2) use sensitive patient data to train new AI products without patient consent, and/or (3) not conform to HIPAA-standards. The AMA advocates that AI developers and their tools should conform to the AMA’s Privacy Principles to ensure the security of private patient information.

1. What information does a physician need to know about an AI tool in order to use it appropriately and effectively?

Several states (e.g., VA [HB747](#), VT [HB710](#), IL [HB5322](#), among others) introduced general AI legislation (not specific to health care) that would require AI tool developers to share documentation that describes the AI’s intended uses, training data, data collection practices, and risk and discrimination mitigation strategies with “deployers” (e.g., physicians, health systems) of AI tools. Colorado and California passed laws related to the disclosure of AI systems. Colorado [SB205](#) establishes transparency and disclosure requirements for developers and deployers of “high-risk” AI systems, defined as “any artificial intelligence

system that, when deployed, makes, or is a substantial factor in making, a consequential decision.”²

The law mandates developers of “high-risk” AI systems to disclose training data, purpose of the AI tool, intended benefits and uses of the AI system, and risks of algorithmic discrimination. Governor Polis signed SB205 into law, but not without reservations; in a [letter](#) to the Colorado General Assembly he noted the law’s extended implementation date of Feb. 1, 2026, and urged lawmakers to address his concerns during the 2025 legislative session.

California [AB2013](#) focuses on disclosure by developers of generative AI systems that are made publicly available to consumers in California. The law broadly defines developers to include “a person, partnership, state or local government agency, or corporation that designs, codes, produces, or substantially modifies an artificial intelligence system or service for use by members of the public.” Under this law, developers are required to publicly post information on, among other provisions, the training data used, the number of data points in the datasets, and whether the datasets include personal information. Since the Colorado and California laws are not specific to health care, it is as-yet unclear how these requirements may integrate into a physician’s daily workflow.

At the federal level there is also recognition that physicians will need practical and timely information related to AI tools. In December 2023, the Health and Human Services (HHS) Office of the National Coordinator for Health Information Technology (ONC) finalized its “Health Data, Technology and Interoperability” (HTI-1) rule,⁷ imposing requirements for developers of ONC-Certified Health IT that use predictive decision support interventions (DSI), including the requirement that the Health IT developer must make detailed information about the predictive DSI available to users of the software (e.g., hospital systems, physicians). Precisely how these requirements will be implemented is yet to be seen, but it is possible that continued regulation from HHS (and states who have developed preliminary language on these topics) will inform future state activities.

2. What information does a physician need to communicate to a patient about the use of AI in clinical practice?

States also introduced bills that would require “deployers” (e.g., physicians) to notify “end users” (e.g., patients) if an AI tool is being used. These bills generally fall into two categories: (1) requiring that, if an AI tool is used to support a “consequential decision” related to an individual, such as determining that individual’s access to health care services, the individual must be notified that an AI tool was used to make the decision (e.g., VT [HB710](#)); or (2) requiring that, if an individual is interacting with a generative AI tool such as a chatbot or if generative AI is used to generate written or verbal communication, they be notified upfront (e.g., NY [SB9381](#)). Utah and California passed laws addressing both these areas: Utah’s [SB149](#) is a general consumer protection law that requires “regulated occupations” to disclose conspicuously the use of generative AI before any oral or written communication with the end user (e.g., patient); this law is not specific to health care, but over 30 different health care professions, including physicians, are listed under the definition of “regulated occupation.” California’s [AB3030](#) is specific to health care, requiring a health facility, clinic, physician’s office, or group practice that uses generative AI to generate written or verbal patient communication pertaining to patient clinical information that is not reviewed by a human licensed or certified health care provider, to ensure the communication includes (1) a disclaimer to the patient stating that the communication was generated by generative AI, and (2) clear instructions describing how a patient can communicate directly with a human health care provider, among other provisions.

² “Consequential decision” is defined as “a decision that has a material legal, or similarly significant, effect on the provision or denial to any consumer of, or the cost or terms of: [...] health-care services; [...] insurance.”

AMA policy position

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. AI disclosures should meet ethical standards and norms. In addition, 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.

- **Transparency between developers and providers**
 - Developers should disclose the following to allow the physician to appropriately evaluate the system or technology prior to purchase³ or utilization:
 - Clear description of problem formulation and intended use
 - Clear and detailed instructions for use
 - Intended population and intended practice setting
 - Clear description of any limitations or risks for use, including possible disparate impact and potential biases (e.g., populations not included in the training data or for which the output is less accurate)
 - Detailed information regarding data used to train the model (e.g., data provenance, data size and completeness, data timeframes, data diversity, data labeling accuracy)
- **Disclosure and record-keeping of AI use in clinical care**
 - When AI is used in a manner which directly 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.
 - 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
- **Physician approval of AI tool outputs**
 - 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

³ Note: The AMA's AI Principles outline additional policy recommendations related to what information a developer should disclose to a purchaser and/or user (physician) prior to the *purchase* of an AI tool.

Physician liability

Liability for the use of AI in health care remains largely undefined. For example, who is liable if a physician relies on an AI tool to support clinical decision-making and the AI tool's output is inaccurate or wrong (e.g., wrong dosage of prescription, incorrect diagnosis, inaccurate charting or coding) remains unclear. Furthermore, while the use of AI in clinical decision making is not yet widespread, state legislators and regulators may want to consider who would be liable (e.g., developer, physician, other) in a future-state where AI tools are regularly used to support clinical decision-making, or indeed make a clinical decision. The liability for physicians from a medical malpractice (tort) context and regulatory oversight (i.e., from state medical boards) may vary but are both critically important. From a tort context, the use of AI in medical decision making is largely untested and will likely be decided in the courts using existing and novel legal theories.

To-date, there has been essentially no state or federal legislative or regulatory activity related to physician liability in the use of AI tools, with the exception of the Office for Civil Rights (OCR) 1557 Non-Discrimination Regulation, in which OCR rejected suggestions that liability be shared with developers or that developers be subject to strict liability.⁴ A select few states introduced legislation that would have required health care providers to review any health care decisions made by or with the use of an AI tool (e.g., Louisiana HB916, Georgia HB887—*Note: neither of these bills progressed to a floor vote*). Notably these bills were silent on liability, but Georgia's bill included language related to the potential for disciplinary action by the state medical board for physicians who fail to review decisions made by an AI tool.

The Federation of State Medical Boards (FSMB), which represents and advocates on behalf of state medical boards, published a [report](#) in April 2024 outlining recommendations for state medical boards to navigate AI, including medical board oversight of physician's using AI. The report includes language that physicians using AI "accept responsibility for responding appropriately to the AI's recommendations" and that "failure to apply human judgement to any output of AI is a violation of a physician's professional duties."⁸ This language is focused on the medical boards regulatory oversight of physicians, which is distinct and separate from theories of liability in the tort context. Nevertheless, this strong language is not out-of-step with traditional theories of a physician's ethical and professional duties that have held physicians responsible for care delivered to patients. FSMB explains that the level of accountability will vary based on how the AI is used, including the risk to patient safety. It is important to note that this is simply guidance, and it remains to be seen if and how this will inform state medical board activity and physician regulation.

AMA policy position

Issues of liability are complex and will likely be determined by the courts for years to come. As legal theories of liability and accountability for AI use evolve, the AMA will continue to advocate to ensure that physician liability for the use of AI-enabled technologies is limited and adheres to current legal theories applicable to medical malpractice. Overall, the AMA believes that individual(s) or entity(ies) that are (1) best positioned to know the AI system's risks and (2) best positioned to avert or mitigate harm need do so through design, development, validation, and implementation. Notably, many if not most of those activities happen before an AI tool is used by a physician in clinical practice.

⁴ "A covered provider's liability under section 1557 is not contingent on or related to a developer's potential liability under this rule or this provision." 89 Fed. Reg. 37649 (May 6, 2024)

- **Physicians**
 - Physicians should not be held liable for the performance of the technology in question if they do not know or have reason to know that there are concerns about the accuracy, quality, and safety of an AI-enabled technology
 - Physicians should not be penalized if they do not use AI systems while regulatory oversight, standards of care, clinical validation, clinical usefulness, and standards of care are in flux
- **Payors, Hospitals, Health Systems, Government**
 - Payors, hospitals, health systems, or governmental entities should be prohibited from mandating the use of health care AI systems as a condition of licensure, participation, payment, or coverage
- **Developers**
 - Developers of autonomous AI systems with clinical applications (screening, diagnosis, treatment) are in the best position to manage risks 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
 - 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

Looking ahead

The AMA anticipates continued state AI legislative activity in 2025. In addition to the policy areas listed above, states will likely continue introducing bills imposing transparency requirements across all stakeholders, prohibiting discrimination by AI tools, and establishing state AI task forces to study the potential legislative and regulatory levers to support the appropriate use of AI.

The AMA remains committed to supporting state medical associations and national specialty societies as they advocate for the adoption of state laws and regulations that provide sufficient patient and physician protections while still embracing the potential for AI to support workflows and reduce burden experienced by physicians.

Endnotes

1. STAT. [Denied by AI: How Medicare Advantage plans use algorithms to cut off care for seniors in need.](#) March 2023.
2. ProPublica. [How Cigna Saves Millions by Having Its Doctors Reject Claims Without Reading Them.](#) March 2023.
3. ProPublica. [“Not Medically Necessary”: Inside the Company Helping America’s Biggest Health Insurers Deny Coverage for Care.](#) October 2024.
4. National Association of Insurance Commissioners. [NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers.](#) December 2023.
5. NAIC. [Implementation of NAIC Model Bulletin: Use of Artificial Intelligence Systems by Insurers as of December 1, 2024.](#)
6. Department of Health and Human Services, Centers for Medicare and Medicaid Services. [Frequently Asked Questions related to Coverage Criteria and Utilization Management Requirements in CMS Final Rule \(CMS-4201-F\).](#) February 6, 2024.
7. Department of Health and Human Services. 45 CFR Parts 170, 171. [Health Data, Technology, and Interoperability: Certification Program Updates, Algorithm Transparency, and Information Sharing.](#) December 2023.
8. FSMB. [Navigating the Responsible and Ethical Incorporation of Artificial Intelligence into Clinical Practice.](#) April 2024.

