AMA
AMERICAN MEDICAL ASSOCIATION

James L. Madara, MD
CEO, EXECUTIVE VICE PRESIDENT

james.madara@ama-assn.org

March 14, 2025

Lynne Parker, PhD
Principal Deputy Director
Office of Science and Technology Policy
Eisenhower Executive Office Building
1650 Pennsylvania Avenue, NW
Washington, DC 20504

Re: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

Dear Dr. Parker:

On behalf of the physician and medical student members of the American Medical Association (AMA), I am writing to provide perspectives on high-priority policy actions for health care augmented intelligence (AI). As you know, the use of AI-enabled technologies in health care is growing at a rapid pace, often without a clear, consensus-driven framework for oversight to ensure AI is performing as intended and protects patient safety. Health care AI has the potential to transform the practice of medicine by enabling highly personalized care, reducing physician burdens, and strengthening the patient-physician relationship. However, to fully realize these benefits, we must establish effective standards and safeguards that ensure AI used in patient care is of high quality, consistently reliable, and prioritizes patient well-being.

The AMA has engaged in policy development work on AI dating back to 2018, when our House of Delegates passed its foundational policy on "augmented intelligence." This policy aimed to help guide AMA's engagement with a broad range of stakeholders and policy makers, and recognize the assistive role that AI-enabled technologies has in the practice of medicine.[1] This policy acknowledges the critical role physicians should play in the design, development, and deployment of health care AI, and emphasizes the importance of safety, transparency, and bias mitigation. In 2024, the House of Delegates expanded on this work and adopted policy that provides more robust and nuanced recommendations regarding clinical and administrative use cases, including enhanced focus on AI transparency, generative AI, AI accountability and liability, propagation of mid- and dis-information by AI, patient data privacy and security, and AI use by payors. [2]

The AMA believes there are several policies the Administration could advance in the near term that would help promote the use of AI in health care while ensuring patient safety and safeguarding access to care. Generally, the AMA advocates for a whole government approach to health care AI that ensures the design, development and deployment of AI-enabled technologies are transparent, responsible, ethical, and equitable. The AMA strongly believes that a framework of oversight needs to be in place to ensure the quality and performance of AI-enabled technologies in health care settings, but acknowledges that

[1] https://www.ama-assn.org/system/files/2019-08/ai-2018-board-report.pdf.
[2] https://www.ama-assn.org/system/files/ama-ai-principles.pdf.

oversight for AI should likely exist on a risk-based scale. Where the use of AI-enabled technology poses a greater risk for patients, such as with AI-enabled medical devices and clinical decision support, strong oversight needs to be in place. Where AI-enabled technologies pose a lower risk to patients, less rigorous oversight may be appropriate. However, to build trust in these new technologies and the shared decision making that results from their use, more progress needs to be made. The AMA outlines several key policy recommendations for your consideration below.

General Governance Principles for Health Care AI

The AMA believes that there are a number of guiding principles for the design, development, and deployment of health care AI that must promote a clear national strategy for oversight of health care AI. These principles seek to promote policies that will build trust among physicians and ensure the safety of our patients, both of which are critically important to recognizing the opportunity that AI presents to improve patient care. The idea that AI developers will simply agree to be good actors absent other guardrails is not sufficient to build the necessary trust and ensure performance. AMA policy outlines several principles for AI governance that are of top priority to physicians and that should underly the Administration's approach to AI oversight and regulation:

- Health care AI must be designed, developed, and deployed in a manner which is ethical, equitable, responsible, accurate, transparent, and evidence based.
- 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.
- 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.
- AI systems should be developed and evaluated with a specific focus on mitigating bias to ensure that the deployment of these technologies does not widen the existing gap in health care access, treatment, or outcomes.
- 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.
- AI risk management should minimize potential negative impacts of health care AI systems while providing opportunities to maximize positive impacts.
- 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.
- 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.

- 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.

AI Transparency

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 is one of the most critical elements to build trust in new AI-enabled technologies and serves to help both physicians and patients understand when and what they are engaging with when AI is involved in the care experience. At a time when excitement surrounding health care AI is high, but trust in these technologies is somewhat low, new policies are necessary to help bolster understanding and to ensure appropriate selection and use. The AMA strongly urges the Administration to move quickly to require transparency for AI-enabled technologies with clinical applications that contribute to medical decision making, such as by finalizing FDA guidance including new recommendations for pre-market submission for AI-enabled medical devices.[3] This guidance includes many new recommendations that are critical to providing physicians with key information regarding AI-enabled medical devices, such as updated labeling recommendations and recommendations for user interface design. We also urge the Administration to develop transparency frameworks for other uses of AI, such as AI with administrative applications.

Current AMA policy provides a framework for transparency actions to ensure they meet the needs of physicians and patients. This framework contemplates both *when* the use of the AI should be transparent to those that engage with it and *what* should be disclosed to users about the technologies they are engaging with, outlined below:

- 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 disclosure should align and meet ethical standards or norms.
- 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.
- 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.
- 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.
- 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.
- 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.

[3] https://www.fda.gov/regulatory-information/search-fda-guidance-documents/artificial-intelligence-enabled-device-software-functions-lifecycle-management-and-marketing.

- 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.

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:

- Regulatory approval status.
- Applicable consensus standards and clinical guidelines utilized in design, development, deployment, and continued use of the technology.
- Clear description of problem formulation and intended use accompanied by 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.
    - Description of how impacted populations were engaged during the AI lifecycle.
    - Detailed information regarding data used to train the model:
        - Data provenance.
        - Data size and completeness.
        - Data timeframes.
        - Data diversity.
        - Data labeling accuracy.
    - Validation Data/Information and evidence of:
        - Clinical expert validation in intended population and practice setting and intended clinical outcomes.
        - Constraint to evidence-based outcomes and mitigation of "hallucination"/"confabulation" or other output error.
        - Algorithmic validation.
        - External validation processes for ongoing evaluation of the model performance, e.g., accounting for AI model drift and degradation.
        - Comprehensiveness of data and steps taken to mitigate biased outcomes.
        - Other relevant performance characteristics, including but not limited to performance characteristics at peer institutions/similar practice settings.
        - Post-market surveillance activities aimed at ensuring continued safety, performance, and uniform access.
    - Data Use Policy:
        - Privacy.
        - Security.
        - Special considerations for protected populations or groups put at increased risk.
    - Information regarding maintenance of the algorithm, including any use of active patient data for ongoing training.
    - Disclosures regarding the composition of design and development team, including diversity and conflicts of interest, and points of physician involvement and review.
    - 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.

Liability and Accountability for AI Errors

Questions of liability that may arise in conjunction with the use of AI in patient care are a significant concern to physicians. The use of AI presents end users with new questions regarding what happens if the AI they were supposed to be able to trust ends up being wrong. The legal questions presented by the use of these tools are novel and complex. They have not been legislated, nor has liability been apportioned through regulation, with the exception of the prohibition of discrimination by clinical algorithms included in the Office for Civil Rights Section 1557 Non-Discrimination Rule.

AI introduces a new universe where physicians and machines will engage in shared decision making, with the physician not necessarily being able to have full insight into the logic or reasoning behind the decision of the machine. Likewise, in the case of algorithmic drift or error, the physician may not be able to understand if and when the technology has become inaccurate. Should physicians ultimately be held liable for their reliance on AI, it begs the question of why a physician would engage these tools at all if they will be held liable for the machine's errors. If technologies that are supposed to be trustworthy but end up erring in ways that could cause harm to patients, yet physicians are liable for trusting those decisions, it follows that we are creating liability risks for physicians that they may be hesitant to bear. This hesitancy has been demonstrated in the AMA's Physician Sentiments of AI survey, where 82 percent of the survey respondents noted liability for AI errors as a top concern when considering engaging with these tools.[4]

While FDA-reviewed AI-enabled medical devices should provide physicians with a reasonable level of assurance regarding their quality and performance, a number of AI-enabled health care technologies are not regulated. Without standards or regulations, physicians have no assurances of safety and quality, creating a higher risk of liability when engaging with them. For use of health care AI to flourish, appropriate apportionment of liability and accountability for AI errors needs to be in place and in-line with AMA policy regarding physician liability for use of AI outlined below:

- 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.

    - 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.
    - 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.
    - 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.

[4] https://www.ama-assn.org/system/files/physician-ai-sentiment-report.pdf.

- 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.
- Liability protections for physicians using AI-enabled technologies should align with both current and future AMA medical liability reform policies.

Mitigating Misinformation in AI-Enabled Technologies

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. 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.

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. The AMA urges policymakers to consider the risks of AI-enabled propagation and dissemination of misinformation and incorporate the following policy principles into any oversight structure for health care AI:

- 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.
- Algorithms should be developed to detect and flag potentially false and misleading content before it is widely disseminated.
- Developers of AI should have mechanisms in place to allow for reporting of mis- and disinformation generated or propagated by AI-enabled systems.
- Developers of AI systems should be guided by policies that emphasize rigorous validation and accountability for the content their tools generate and 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.
- 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.
- 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.

AI Data Privacy and Cybersecurity

The AMA views data privacy and cybersecurity as bedrock elements for safeguarding patient rights and sustaining physician trust in AI-powered tools. The integration of AI into health care is profoundly transformative, yet it poses equally profound risks if not managed responsibly. Ensuring the protection, confidentiality, and integrity of data used to power AI systems is paramount, especially when the data involved are deeply personal, revealing patient health conditions, demographic details, and other potentially sensitive information.

The Critical Importance of Data Privacy Measures

AI development relies on enormous repositories of health data for training, validating, and improving algorithms. Because nine out of 10 patients believe privacy is a right—and nearly 75 percent of people express concern about the privacy of their health data—there is a compelling need for robust privacy frameworks that go beyond traditional Health Insurance Portability and Accountability Act (HIPAA) rules to account for the "black box" nature of AI systems.[5] This includes:

- **Minimizing Data Collection**: AI developers should collect only the minimum amount of data necessary for a defined purpose. Guardrails must be in place to limit data usage to legitimate, specific aims (e.g., clinical improvement) rather than ambiguous or open-ended scenarios.
- **Anonymization and Pseudonymization**: Advanced data-handling techniques (e.g., deidentification, secure enclaves, federated learning, and differential privacy) must be incorporated by AI developers to reduce the risk of reidentification. These measures are especially critical in generative AI applications, where personal or otherwise confidential data might inadvertently surface in the course of the model's outputs.
- **Transparency and Patient Rights**: AI developers must fully inform patients and physicians about how data will be collected, stored, and reused. Users must have the right to opt out, request the deletion of their data, and revise or revoke consent as appropriate. Without clear and consistent transparency measures, patients may lose confidence in AI and be reluctant to share data critical to advancing these technologies.

Cybersecurity and Threat Mitigation

Cyberattacks such as ransomware, phishing schemes, and adversarial or "model poisoning" attacks pose serious risks to health care infrastructure, given the sector's reliance on interconnected devices, including electronic health record systems, telemedicine platforms, and mobile health applications. In 2017, 83 percent of physicians surveyed had already experienced some form of cyberattack, and, more recently, high-profile disruptions—such as the February 2024 ransomware attack on Change Healthcare—demonstrate the severity of the threat landscape. A single successful cyberattack can:

- **Compromise Large-Scale Data**: AI systems demand continual data exchange across complex networks. If these networks are breached, attackers can gain access to sensitive patient information, practice-specific data, and even the proprietary elements of an AI model itself.
- **Undermine Clinical Integrity**: When threat actors engage in data poisoning or adversarial manipulation, the accuracy and reliability of AI models degrade significantly. An altered or

[5] https://www.ama-assn.org/system/files/ama-patient-data-privacy-survey-results.pdf

corrupted AI model can yield inaccurate predictions or recommendations, jeopardizing patient safety and clinical outcomes.
- **Facilitate Intellectual Property Theft**: Model-stealing attacks can allow malicious actors to reproduce highly valuable AI frameworks. In health care, where solutions are costly to develop and refine, such theft can blunt innovation and deter investment.

Because AI is particularly sensitive to data quality, the AMA emphasizes policies requiring AI developers to secure all points of vulnerability—from training data and data feeds to user interfaces and system endpoints. Specifically:

- **Secure Data Inputs**: Entities developing or deploying AI must validate and scrutinize the quality of data inputs to reduce the risk of data poisoning.
- **Advanced Cyber-Defenses**: Developers should employ tools such as encryption protocols, multifactor authentication, routine system audits, and continuous monitoring of AI outputs to identify malicious intrusions or anomalous behavior.
- **User Education and Accountability**: Cybersecurity is not solely a technology issue—it is also a human one. Health care professionals, administrators, and all end users of AI tools must receive training to recognize potential threats, report suspicious activities, and adhere to best practices that protect data integrity.

Accountability and Collaboration

As AI models proliferate, the burden of data security and privacy should not rest solely on physicians or individual health care providers. Many AI developers enter legal arrangements (e.g., business associate agreements) bringing them under HIPAA; however, significant accountability gaps remain. The AMA recommends:

- **Distribution of Liability**: Developers, implementers, and all third-party technology vendors need to shoulder appropriate responsibility for safeguarding and monitoring AI-driven systems.
- **Monitoring and Rapid Incident Response**: When breaches occur or anomalies arise, timely notification to both health care organizations and individuals is critical. This must extend beyond the narrow bounds of mandatory HIPAA requirements to ensure all patients and clinicians potentially affected are promptly informed.
- **Aligned Public-Private Efforts**: Ongoing dialogue among federal agencies (including the Assistant Secretary for Technology Policy/Office of the National Coordinator, the Federal Trade Commission, and other regulators), health care institutions, and AI developers is essential to fostering an environment in which AI innovation can thrive without sacrificing patient trust or safety.

Payor Use of Augmented Intelligence and Automated Decision-Making Systems

Payors, including health insurers and health plans, are increasingly employing AI-driven or algorithm-based systems to make claim determinations, define coverage limitations, and engage in benefit design. This development holds both promise and peril for patients and physicians. On the one hand, these systems can expedite administrative workflows, reduce paperwork, and streamline prior authorization processes. On the other hand, multiple reports indicate that certain automated decision-making tools are being used to rapidly deny care without meaningful oversight, often based on incomplete or narrowly

tailored data. These concerns have sparked growing state-level scrutiny, heightened interest from Congress, and ongoing discussions about the appropriate federal regulatory framework.

Emerging Patterns and Concerns

Recent accounts illustrate why new guardrails are urgently needed. Reporting by ProPublica found that one health plan used an automated tool to deny approximately 300,000 claims in just two months, with each claim receiving a review that averaged 1.2 seconds. Separately, two class action lawsuits were filed in 2023 against United Health Care and Humana, alleging that an AI model, nHPredict, denied medically necessary care to elderly and disabled Medicare Advantage enrollees. The lawsuits further assert that 90 percent of the tool's denials were faulty—and that payors continued to rely on the tool despite this knowledge.

Such instances underscore a troubling trend wherein automated systems may rely on generalized or "similar patient" profiles rather than patient-specific clinical data. This lack of individualized review can lead to inappropriate denials, truncated coverage periods, and the imposition of coverage limitations that hinder patients' access to critical services or force them to pay out of pocket.

Potential Benefits and the Need for Balance

Notably, AI has the potential to reduce administrative burdens on physicians if implemented responsibly. Automated systems can standardize documentation requests, speed up claim adjudication, and minimize the time physicians spend on administrative tasks. This improved efficiency could lead to lower health care costs and improved patient experiences, provided that:

- **Clear Guidelines Exist**: Payors should use automated decision-making only in well-defined, lower-risk situations or as a first pass to facilitate routine approvals. Human oversight remains essential for complex or borderline cases.
- **Transparent, Evidence-Based Criteria**: When coverage determinations are delegated to an AI system, the criteria used should be fully disclosed to patients, physicians, and regulators.
- **Opportunity for Override**: There must always be a pathway to override an automated decision when unique clinical or social circumstances warrant a more nuanced review. Payors should empower treating physicians to discuss care needs directly with qualified clinicians making final determinations.

Insufficient Statutory and Regulatory Protections

There are currently no explicit federal statutory directives—and only minimal regulatory requirements—governing payors' use of AI-driven automated decision-making in coverage and benefits. While the Biden Administration called attention to the need for guidance, neither the Centers for Medicare & Medicaid Services nor other agencies adopted broad, binding rules specifically addressing payor AI use. This regulatory gap has led some states to consider legislation or regulations that would impose additional transparency and accountability requirements. Congress, too, has convened hearings on the topic, although no subsequent laws have been enacted.

Key Recommendations for Responsible Payor Use of AI

To ensure that AI-driven decisions do not inadvertently restrict care or harm patients, the AMA recommends the following:

- **Public Reporting and Disclosure**: Payors should publicly report any use of automated decision-making systems for coverage determinations. They should base their decisions on clearly documented, easily accessible, evidence-based clinical guidelines rather than opaque or proprietary criteria.
- **Limitations on Use**: AI tools should only be used to enhance efficiencies, reduce administrative complexity, and diminish workflow burdens. They must never replace individualized assessments of patient-specific medical and social factors. Patients and physicians should be empowered to challenge automated decisions that appear medically inappropriate.
- **Human Review for Adverse Determinations**: Any algorithmic recommendation denying or limiting care must be subject to physician review prior to finalization, particularly when it involves serious or high-risk clinical scenarios. When denial is considered, the treating physician must have an opportunity to speak directly with the reviewing physician who oversees AI-based determinations.
- **Transparency in Data Sources and Model Training**: Payors should disclose the sources of data used to train their AI systems, including relevant demographic attributes. Such transparency helps ensure the model has been assessed for bias.
- **Regular Audits and Monitoring**: Payors must conduct and publish the results of regular audits to evaluate patterns of denials, approval rates, and appeals—broken down by key demographic variables. Where disparities emerge, payors should take corrective measures, including model retraining or modification.

The AMA appreciates the Administration's focus and attention on recognizing the opportunities presented by AI. While AI is a growing consideration across several sectors, we urge the Administration to ensure that oversight and implementation of AI in health care continues to be of high priority. Health care AI is unique in that its risks to the health and wellbeing of our patients are potentially high. It also, however, presents us with a significant opportunity to provide our patients with high quality, high value, extremely personalized care if we are careful to get it right. The AMA looks forward to working with you to prioritize focus on health care AI and to promote policies that ensure the safety and wellbeing of our patients. Please do not hesitate to contact Shannon Curtis, Assistant Director of Federal Affairs, at Shannon.Curtis@ama-assn.org with any questions or to discuss further.

Sincerely,

James L. Madara, MD