AMA
AMERICAN MEDICAL ASSOCIATION

James L. Madara, MD
CEO, EXECUTIVE VICE PRESIDENT
james.madara@ama-assn.org

July 7, 2023

The Honorable Arati Prabhakar
Director
Office of Science and Technology Policy
Executive Office of the President
Eisenhower Executive Office Building
1650 Pennsylvania Avenue
Washington, DC 20504

Dear Director Prabhakar:

On behalf of our physician and medical student members, the American Medical Association (AMA) appreciates the opportunity to provide comments on the Office of Science and Technology Policy's (OSTP) Request for Information on National Priorities for Artificial Intelligence. As you can imagine, health care technology is advancing rapidly for many different uses and within many different sectors of the health care industry. Ensuring the responsible, equitable, ethical, and transparent design, development, and deployment of high-performing augmented intelligence (AI)-enabled tools within the health care space is a key priority for AMA members and our patients. **We strongly encourage OSTP and the Administration to broadly ensure health care AI is considered as a sector of significant national concern and importance and to engage with health care stakeholders to ensure appropriate policies, standards and regulatory requirements are in place to protect patient safety, promote equity, and ensure the quality and performance of the AI-enabled tools in question.**

AMA Policy on Augmented Intelligence

In June 2018, the AMA's House of Delegates adopted its first policy on health care augmented intelligence.[1]This policy provides that the AMA will leverage its ongoing engagement in digital health and other priority areas for improving patient outcomes and physicians' professional satisfaction by helping set priorities for health care AI. First and foremost, that policy supports the use of AI systems where those systems advance the quadruple aim of health care. Specifically, AI systems should: 1) enhance the patient experience of care and outcomes; 2) improve population health; 3) reduce overall costs for the health care system; and 4) support the professional satisfaction of physicians and the health care team. The policy further provides that the AMA will promote the development of thoughtfully designed, high-quality, clinically validated health care AI that:

- Is designed and evaluated in keeping with best practices in user-centered design, particularly for physicians and other members of the health care team;

[1] American Medical Association. *Augmented Intelligence in Medicine.* https://www.ama-assn.org/system/files/2019-08/ai-2018-board-report.pdf (Accessed June 28, 2023).

- Is transparent;
- Conforms to leading standards for reproducibility;
- Identifies and takes steps to address bias, avoids introducing or exacerbating health care disparities while actively addressing the same, including when testing or deploying new AI tools on marginalized populations; and
- Safeguards patients' and other individuals' privacy interests and preserves the security and integrity of personal information.

Initial AMA policy contemplated the use of AI in direct patient care. However, with the introduction of technologies such as generative AI and large language models (LLMs), as well as expanding use in non-clinical settings, such as by payors in making claims and coverage determinations, AMA policy continues to evolve. At our most recent House of Delegates meeting in June 2023, new AMA policy has been adopted that urges additional advocacy and action on the potential for AI-enabled dissemination of mis- and disinformation,[2] false and misleading responses generated by AI,[3] and the use of AI by payors resulting in limiting access to health care.[4]

Health Care Augmented Intelligence Opportunities and Risk

The AMA is tremendously excited about the opportunities presented by AI. In health care, AI is showing remarkable promise for many clinical uses, such as diagnostics, clinical decision support (CDS), treatment planning, and patient education. Several AI-enabled medical devices have already been approved or cleared by the Food and Drug Administration (FDA) for use in humans, including several image-based algorithm-enabled devices that primarily serve diagnostic screening or triage-type functions. It is hoped that, over time, these types of medical devices will serve as important clinical tools with strong value propositions, potentially enabling some health care services to become more affordable and more accessible. Eventually, we anticipate that many AI-enabled medical devices will perform at exceptionally high levels, helping to ease physician burdens and allowing physicians to spend more time with patients.

While a significant amount of attention has been paid to AI-enabled devices used in direct clinical care, AI is also being deployed for non-device administrative-type functions in the health care industry. Uses of AI here vary widely, from use by insurers to make claims and coverage determinations; at a population health level within health systems and hospitals to predict occurrence of infection and/or readmissions; in patient charting and scheduling; in resource allocation; and among many other non-clinical operations. AI is also showing significant promise in R&D, where AI-enabled drug development is gaining traction and alongside machine learning applications to large genomic datasets.

While AI has significant potential to improve the quality and cost of health care, improve access and reduce physician burden, it is not without risks at this juncture—risks that could pose serious harm to patients and create significant liability for physicians. The AI developed to date has proven somewhat limited in its scope, ability, and performance. Bias and discrimination continue to be a very serious problem with AI algorithms, and there has been little progress in our ability to readily identify the source of bias with AI-enabled technologies and to mitigate it, especially in machine learning algorithms. The introduction of new generative AI and LLMs such as ChatGPT have shown that while AI can perform

[2] American Medical Association. Report of Reference Committee B. Pg. 72. https://www.ama-assn.org/system/files/a23-refcomm-b-annotated.pdf (Accessed June 28, 2023).
[3] American Medical Association. Report of Reference Committee B. Pg. 72. https://www.ama-assn.org/system/files/a23-refcomm-b-annotated.pdf (Accessed June 28, 2023).
[4] American Medical Association. Report of Reference Committee G. Pg. 22. https://www.ama-assn.org/system/files/a23-refcomm-g-annotated.pdf (Accessed June 28, 2023).

competently at some tasks, it can also make very egregious errors and readily generates false information in many circumstances. While some AI-related errors can be rectified, with decisions about an individual's or population's health care, any error could have a significant detrimental impact on a patient's health and wellbeing making the risks of AI in the health care space significant.

Regulatory Needs

The AMA is committed to ensuring that the design, development, and deployment of AI is ethical, equitable, responsible, and transparent and promotes the use of only high-quality, high-performing AI technologies. These attributes are critical when considering AI deployed in health care settings, given the significant risks to patient's health and wellbeing should poorly performing AI proliferate. High quality, high performing AI is of value for both clinical and administrative uses but is especially important when considering technologies used in direct patient care.

As you are well aware, the United States is currently lacking any meaningful policies, regulations, or standards for the oversight of AI other than that which is regulated as a medical device by the FDA. While AI-enabled products and tools that meet the definition of a medical device or are otherwise considered device CDS are reviewed by and must meet the regulatory standards of the FDA, the current FDA oversight structure is not an ideal fit for the regulation of software-based medical devices and AI-enabled devices and tools. The traditional regulatory structure for hardware medical devices has been tested by software-based devices, as they pose new challenges (frequent updates for most, potential for continuous learning and updates for machine learning-based AI) that are not easily accounted for by the traditional regulatory system. While FDA has made admirable progress in attempting to create an appropriate regulatory structure for these new technologies, little is finalized and significant regulatory gaps specific to software-based medical devices and Augmented Intelligence/Machine Learning (AI/ML)-enabled devices still exist. For non-device AI, there are no existing regulatory guidelines or requirements for health care AI.

Consistent Terminology

The field of AI is presenting numerous challenges to governments, industry, and individuals, many of which are exacerbated by a current lack of consistency and clarity around roles, responsibilities, and oversight. While we are likely in the very initial stages of pursuing initial solutions to some of these early challenges, there are some immediate actions that government policymakers could pursue that would provide some initial clarity to interested parties. Before engaging in any additional rulemaking in this space, government policymakers should work with stakeholders to put forth agreed upon terminologies and taxonomies to describe AI and its components. Currently, a number of different departments and agencies have engaged in policymaking regarding AI and descriptions and definitions of AI and its related components consistently vary. FDA, for example, deems AI-enabled technologies with a machine learning component as "AI/ML," and the Office of the National Coordinator for Health IT (ONC) has coined the "decision support intervention," while the Health and Human Services Office for Civil Rights (OCR) uses "clinical algorithm" as a catch all for any algorithm or algorithm-enabled tool, including AI. Meaningful progress on AI policy and integration must mean interested parties are all consistently speaking the same "language."

Further discussion of AI is potentially even more complex, when defining types of AI and how they are used. The AMA has consistently advocated for the development of agreed-upon terminology in this space and has played a role in providing some clarity through its support of the Current Procedural Terminology

(CPT) Editorial Panel's adoption of Appendix S to the CPT code set.[5] Appendix S represents a stakeholder-driven effort to define the levels of autonomy within different AI systems physicians may encounter, providing consistency for regulators and payors as we begin to consider the need for increased clinical integration and reimbursement of these new technologies.

Transparency

The AMA firmly believes that responsible, equitable, and ethical deployment of health care AI must be transparent, regardless of whether the AI in question is part of a regulated medical device or is non-device AI. As health care AI potentially poses significant and disparate risks to the health and wellbeing of patients, certain elements of transparency must be mandated when AI-enabled tools are deployed. Transparency mandates must include mandated disclosure about the use of AI or other algorithms in health care decision-making at every level. Transparency mandates must also require the disclosure of certain information about AI-enabled tools that allow physicians and potentially patients to assess the performance and quality, including disparate impact, of the AI, or at least factor the use of AI into shared decision-making.

Currently, there are no mandated transparency requirements for any AI-enabled health care devices, systems, or tools. While FDA approval/clearance obviously requires certain disclosures about the device prior to marketing, there have been no efforts to update labeling requirements for these types of devices to ensure the labeling is fit for purpose and appropriate for these new software-based devices. Product labeling requirements for FDA-approved/cleared medical devices should be put in to place and should include new transparency requirements that help physicians appropriately evaluate the performance of the device, understand the intended use and conditions for deployment, provide information about the data set used to train the algorithm, address the risk for bias and disparate impact, assess data privacy issues, and provide information regarding clinical validation of the algorithm. As software-based medical devices, particularly those that are AI/ML-enabled, pose different regulatory challenges than traditional hardware-based medical devices, it is critical that updated approaches to labeling and transparency be put into place. Without appropriate disclosure of relevant information that allows a physician to evaluate the performance and suitability of the device in question, physicians should not engage with AI-enabled devices due to the risk of patient harm and resulting physician liability.

While the use of AI in a regulated medical device may be the foremost thought when considering health care AI, it must be noted that there is an overwhelming number of use cases for what we consider "non-device" AI. Non-device AI has the potential to touch almost every aspect of the health care industry and can include uses in non-device clinical decision support, patient education, population health, physician administrative burden reduction, health care resource allocation, and individual coverage determinations and payor utilization management decisions. While some of these uses pose significantly more risk to patient health than others, all potential uses have the ability to impact individual patient care.

Due to the continuous potential impacts on patients, the AMA supports industry- and government-wide transparency mandates when AI is used to make health care decisions or generate health care information. Physicians and patients must be informed when clinical decisions or recommendations include the use of AI-enabled decision-making tools, when AI is used to generate patient-facing health care information, and when payors utilize algorithms to make coverage determinations or issue claim denials, among other uses. Stakeholders and policymakers must also consider what additional types of information about AI must be

[5] American Medical Association. *Current Procedural Terminology Appendix S: AI Taxonomy for Medical Services and Procedures.* https://www.ama-assn.org/practice-management/cpt/cpt-appendix-s-ai-taxonomy-medical-services-procedures (Accessed June 29, 2003).

disclosed to physicians and patients to ensure the appropriateness and validity of AI-generated recommendations and decisions, to limit the spread of AI-enabled or generated mis- and disinformation, and to ensure payors are not utilizing AI and other algorithms to inappropriately and disparately deny coverage to patients without appropriate evidence-based criteria and medical review or based on faulty training data.

While there is a clear and concerning lack of any current health care AI transparency mandates from any government entity, the AMA is pleased to see proposals from ONC that would require certain disclosures from developers of certified health information technology (health IT), e.g., electronic health record (EHR) systems, regarding the use of AI systems in their EHR products. The AMA applauds ONC's efforts towards addressing the complexities surrounding transparency of rapidly evolving AI tools in health care. ONC's use of its Health IT Certification Program to surface information about AI development, training, risk, and fairness is likely the federal government's first major effort to establish regulatory guardrails on AI transparency. The AMA strongly supports ONC's vision to promote greater trust in health care AI and predictive models through the ONC Health IT Certification Program.

The AMA also appreciates ONC's efforts to make information about EHR-enabled or -interfaced AI available to an end user in plain language. It is likely many physicians will engage with AI while using their EHR. Physicians, for better or worse, spend hours each day in their EHR and are likely accustomed to finding and using information to inform their decisions. Providing physicians insight into AI tools through their EHR leverages the health IT ecosystem they are already familiar with.

We also recognize that ONC's AI transparency and risk management efforts may not ensure information is uniform or consistent, and that the utility of the information may not always be user centric. While this may ultimately become the goal, we believe that surfacing information about AI tools is the first step in providing physicians meaningful knowledge about the trustworthiness of AI in health care. The AMA encourages the Administration to continue evaluating novel approaches in utilizing its health IT policies to increase AI transparency and better inform AI users.

As AI-enabled tools are rapidly coming to market, physicians, clinicians, and patients will need education on AI's impact on health care and the practice of medicine. Physicians, particularly, will need training to better maximize the use of AI and to enhance their confidence in utilizing health care AI. Education should focus on assisting physicians in interpreting if AI tools are fair, appropriate, valid, effective, and safe. All AI users will need plain language descriptions to better understand how AI decisions or clinical recommendations are made. The AMA is encouraged by ONC's efforts to promote trustworthy AI through transparency and believes this will help establish a basis for education. The Administration should coordinate with public and private AI experts and develop educational materials to help physicians and other users interpret and act on information provided by AI and health IT developers.

Consensus Standards and/or Guidelines for AI Design, Development, and Deployment

The AMA understands that, as a whole, it will be immensely challenging, if not impossible, to broadly regulate all specific and individual uses of health care AI. Oversight of FDA-regulated AI-enabled devices presents significant challenges to the traditional system of medical device regulation and will continue to evolve and constantly require new methods and thinking. While it is clear that AI-enabled technologies meeting the definition of a medical device or considered device CDS must continue to be individually evaluated and regulated to ensure safety and efficacy, it is also clear that doing so will likely require movement towards a standards-based approach. Clear, consistent regulatory standards are essential to not only provide clarity and certainty to developers, but to engender trust among physicians and patients.

It is also obvious that to ensure safe, responsible, equitable, and ethical deployment of high-quality non-device AI in health care broadly, a standards-based approach should be utilized, if not mandated. A number of AI-enabled technologies and tools currently fall well outside of the regulatory authority of FDA, including generative AI tools such as ChatGPT, and will, thus, never be reviewed or otherwise regulated. Physicians and patients alike will have no assurances of quality, safety, or performance and may have little information by which to evaluate the product. Development of standards applicable to these types of AI will be necessary to limit patient harm and promote health.

The AMA strongly urges the federal government and other policymakers to consider how we move closer towards consensus standards for responsible, equitable, and ethical design, development, and deployment of health care AI. Currently, a number of organizations, coalitions, and efforts are active in trying to develop standards application to health care AI and AI more broadly. However, no existing efforts are comprehensive enough to ensure the safety and performance of health care AI and there is no broad consensus about which efforts should be controlling. While the AMA is pleased to see the broad engagement in this space, several competing standards development efforts are not necessarily helpful, and can potentially hinder efforts to develop clear and consistent standards and guidelines in the health care AI space. Federal policymakers should consider how we best develop standards and guidelines to ensure the safety and performance of health care AI and should work closely with other health care stakeholders to help either drive consensus standards or mandate a process for doing so. It is also critical that any standards process be open to participation by any interested parties and provided for public comment. Standards processes driven only by those with a vested financial interest in the outcomes do not always align with the best interests of patients. We must ensure that any attempt to regulate or otherwise create standards for AI design, development, and deployment be patient-centered, not profit-centered, at its core.

In both device and non-device AI, compliance with appropriate standards for design, development, and deployment should be a mandated element of AI transparency. Once appropriate standards for AI-enabled technologies are in place, developers should be mandated to disclose to consumers which standards their product meets. This type of transparency will provide end users with valuable information about the technology in question and will assist in the ability to make more informed decisions and assess the quality of the AI in question.

Patient Data Privacy and Cybersecurity

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 falls with the data holders. In other words, third parties who access an individual's data should act as responsible stewards of that information, just as physicians promise to maintain patient confidentiality. The AMA has developed a set of Data Privacy Principles that outlines essential principles for patient privacy in digital health. The Principles provide individuals with rights and protections from discrimination and shift the responsibility for privacy from individuals to data holders. The Principles also call for robust enforcement of penalties for violation of rights to help patients develop and maintain trust in digital health tools, including health care AI.

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 EHRs are interoperable across the health care system and, therefore, accessible by AI that has been trained box" nature of AI. 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.

Patients, too, are increasingly concerned about data misuse. In a 2022 survey of 1000 patients, 92 percent considered data privacy a right. Nearly 75 percent expressed concern about protecting the privacy of their health data. While patients are most comfortable with their physician or hospital having access to their data, patients are least comfortable with social media sites, employers, and big technology companies receiving access to their health data. In fact, 94 percent of patients believe that companies that collect, store, analyze or use health data should be held legally accountable for data misuse or inappropriate access.[6]

Yet, data privacy relies on strong data security measures. There is growing concern that cyber criminals will use AI to attack hospitals and health care organizations. AI poses a new complexity to health IT operations. AI-operated ransomware and AI-operated malware can be targeted to infiltrate health IT systems and automatically exploit vulnerabilities. We are already aware of attackers using ChatGPT to craft more convincing or authentic emails to entice people to click on links—giving them access to electronic health information.

AI is particularly sensitive to the quality of data. Data poisoning is the introduction of "bad" data into an AI's training set, affecting the model's output. AI requires a great deal of data to build logic and patterns used in clinical decision-making. Protecting this source data is critical. Threat actors with access to input data could introduce data sets that pollute the training information used in AI. Failure to secure and validate these inputs, and corresponding data, can contaminate AI models—resulting in patient harm.[7] More needs to be done to educate AI end-users on identifying high-quality, safe, and secure AI systems. Lack of understanding and awareness may result in acquiring AI-based products that are misconfigured, easily exploitable, or maliciously designed to impact operations, steal sensitive data, or disrupt clinical decision-making. The complex, interdependent nature of AI can also create multiple avenues for cyber-attacks. Ransomware groups may use new tactics—recognizing that health care organizations can be exploited through the dozens of vendors and other third parties who have access to their health IT systems.

The promise of AI comes at a price. It is important that physicians, administrative, and clinical leaders are empowered to make informed decisions about the purchase and implementation of AI and to be prepared to utilize these tools in a safe and effective way. In addition, manufacturers that build AI devices must understand the downstream impact of flaws or vulnerabilities in their software's data privacy and security protections. **The AMA urges the Administration to consider federal policies geared to promote robust data privacy and security risk mitigation practices, and to increase the information available to AI users to make informed decisions.**

Liability and Accountability for Use of AI-Enabled Technologies

In order to promote the clinical integration of promising AI-enabled technologies, physicians and patients must trust that they can appropriately rely on the performance of the technology and that it will not create bias, exacerbate health inequities, and cause undue harm to patients, in addition to creating additional liability concerns for the physicians. Liability and accountability for use of AI and other algorithm-enabled technologies in clinical practice presents very novel and complex legal questions which have not yet been litigated in court and are very far from settled. This new potential liability for physicians presents a significant potential barrier to clinical integration of AI-enabled technologies, since if the liability risks

[6] Id.
[7] https://healthsectorcouncil.org/wp-content/uploads/2023/02/Health-Industry-Cybersecurity-Artificial-Intelligence-Machine-Learning_1.pdf

are too severe, there will be no incentive for a physician to incorporate these new technologies. Since 2018, AMA policy has stated that those best situated to mitigate the risks from AI should be held accountable for any resulting harms. In other words, the risk of poorly designed or poorly performing AI should rest with those designing, developing, and deploying the technology, not with the physician. Additionally, when considering liability for poorly performing AI, regulators and courts should consider a knowledge-based standard for assessing who is responsible for any harm—liability should only result if a physician knew or should have had reason to know that the AI in question may perform poorly or was performing poorly and utilized it anyway. Physicians should likely remain liable if they inappropriately or negligently utilize an AI tool and these instances should follow well-held existing standards for determining medical liability.

The AMA has serious concerns with initial federal efforts attempting to assign broad AI liability. In particular, we are concerned with an initial proposal by OCR that would assign individual physicians' liability for using algorithms, including AI, that ultimately resulted in discriminatory harm to patients. While we agree that there should be liability for these types of harms, we do not agree that liability for poorly designed and discriminatory AI should rest with the individual physician. Compared to an individual liability approach which tends to mask system problems and leave underlying system issues unaddressed, a quality and safety approach is more effective for surfacing and addressing system problems. If liability for AI systems is not appropriately apportioned, it will represent a very critical loss of opportunity as integration and uptake of these new technologies will suffer and quality improvement and patient safety efforts will be undermined. **We urge the Administration to work with its departments and agencies to ensure that regulatory oversight of the development and use of AI does not inadvertently disincentivize the use of high-quality AI with continuous quality improvement. Regulatory systems must appropriately ensure safety and efficacy without creating excess liability for end users and enforcement agencies and offices must ensure the liability is appropriately placed with those most responsible for any resulting harm.**

Promoting Health Equity and Mitigating Algorithmic Bias and Discrimination

While the processes of many AI/ML-enabled technologies and tools are automated, they rely on human policies and institutional procedures, human expertise, and human-generated data, which are all subject to the innate biases and inequities that exist among institutions, people, and data. As in other areas of society, this risk of biased outputs based on biased inputs poses ethical challenges for health care institutions and clinicians, as well as potential disparate health harms to historically marginalized patients and communities. Though these issues predate AI, ignoring them will likely increase health inequities. Rather, we should remain conscious of the risks and opportunities throughout the lifecycle of AI and continually engage those most impacted by its use.

Diversity and inclusion in the development of AI are lagging. Lack of gender diversity persists in scientific authorship of AI-related publications, and women hold fewer than one third of data science jobs, mirroring lack of ethnic and racial workforce diversity. Underrepresentation narrows the research agenda, resulting in lost opportunities to develop “equity-centered technologies.” The “leaky pipeline,” where “marginalized groups progressively leave science, technology, engineering, and mathematics (STEM) subjects in the period between preschool and college” and lack of emphasis on diversity within the workplace both contribute to underrepresentation in the AI industry, with “good professional traits recognized in the [majoritized] group… valued differently in [minoritized] groups” (e.g., ambition) and a

"minority tax" diverting minoritized professionals from their own career advancement via research and teaching to "diversity initiatives" (e.g., committee service).[8]

Preliminary evidence shows team diversity including designers, coders, clinicians, and end users can help with "better anticipating the likely impacts of certain model choices on [specific] groups and modes of failure…. averaging out... prediction errors across developer [specific] groups…, [and] broadening the actual questions being addressed by AI…."[9] Efforts around AI/ML can draw on the broader Principles for Equitable Health Innovation developed by the AMA in collaboration with others as In Full Health.[10] These lessons and principles could be extended to government, with priority placed on establishing diverse teams that are creating and implementing regulatory frameworks, as well as building these aspects into government contracting and grant-making, supporting equity advances in the private sector.

As had been previously seen with a number of algorithms, AI has the capacity to amplify biases resulting in disparate impact. Convolutional neural networks "that provide high accuracy in skin lesion classification are often trained with images of skin lesion samples of white patients" with datasets only including about 5-10 percent Black patients, resulting in only half the diagnostic accuracy among Black patients compared to original claims, especially troubling when Black patients having the highest mortality rate for melanoma. "AI algorithms used health costs as a proxy for health needs and falsely concluded that Black patients are healthier than equally sick white patients, as less money was spent on them. As a result, these algorithms gave higher priority to white patients when treating life-threatening conditions, such as diabetes and kidney disease, even though Black patients have higher severity indexes."[11] These examples highlight the need for standardized approaches to addressing known and potential inequities related to AI product testing.

Until recently, guidelines and regulatory frameworks related to AI/ML have focused on technical and operational issues, including only a cursory mention, if at all, rather than embedding equity considerations. Specific examples including "SPIRIT-AI" and "CONSORT-AI" mention but do not formally include disaggregating by specific groups, while "TRIPOD-ML" and "STARD-AI guidelines for model reporting do not allude to these issues." Meanwhile, "the engagement of prominent regulatory bodies with Machine Learning for Health care remains at a preliminary stage, and engagement with fairness tends to be either minimal or vague."[12]

The guideline landscape has begun to change with the release of documents such as the National Institute of Standards and Technology (NIST) publication, *Towards a Standard for Identifying and Managing Bias in Artificial Intelligence*, which takes a socio-technical approach to addressing systemic, statistical, and human sources of bias with preliminary guidance.[13] Authors suggest "analyzing for disparities between less and more socially advantaged populations across model performance metrics…, patient outcomes,

---

[8] de Hond AA, van Buchem MM, Hernandez-Boussard T. Picture a data scientist: a call to action for increasing diversity, equity, and inclusion in the age of AI. Journal of the American Medical Informatics Association. 2022 Nov 14;29(12):2178-81.
[9] Id.
[10] The Principles for Equitable Health Innovation. In Full Health. https://infullhealth.org/our-principles/ (Accessed June 28, 2023)
[11] Norori N, Hu Q, Aellen FM, Faraci FD, Tzovara A. Addressing bias in big data and AI for health care: A call for open science. Patterns. 2021 Oct 8;2(10):100347.
[12] Wawira Gichoya J, McCoy LG, Celi LA, Ghassemi M. Equity in essence: a call for operationalising fairness in machine learning for healthcare. BMJ Health Care Inform. 2021 Apr;28(1):e100289. doi: 10.1136/bmjhci-2020-100289. PMID: 33910923; PMCID: PMC8733939.
[13] Schwartz R, Vassilev A, Greene K, Perine L, Burt A, Hall P. Towards a standard for identifying and managing bias in artificial intelligence. NIST Special Publication. 2022 Mar 15;1270:1-77.

and resource allocation and then identify[ing] root causes of the disparities…" throughout the AI lifecycle.[14] Guidelines for algorithms could additionally build on specific equity approaches that have been described for observational studies and randomized controlled trials, moving from case studies to more robust evaluation.[15,16] Additional algorithmic-specific frameworks could provide approaches for guidelines and regulations. These include the Toronto Declaration human rights approach,[17] an operationalizing fairness approach,[18] Open Science tools for reducing bias in AI,[19] and a fairness-in-design approach with notions under individual and group domains.[20]

AI holds promise for the practice of medicine and public health. **To address the equity and ethics challenges for all interested parties, including health care systems, health workers, and patients and communities, a robust framework is needed to prevent, mitigate, and redress algorithmic bias, at the individual, institutional (in the form of policies and procedures), and societal (legislation and regulation) levels in the public and private arena.**

The AMA is strongly committed to the advance of AI in health care. However, to fully realize the promise of these new advancements, we must ensure that deployment does not outpace our ability to ensure its safety. All stakeholders must firmly commit to the responsible, equitable, ethical, and transparent design, development, and deployment of health care AI and federal and state governments bear responsibility for helping to ensure appropriate standards are in place to do so. The United States is well positioned to remain a global leader in health care AI, but unless we can ensure that physicians and patients can trust that health care AI is safe, equitable, high-quality, and improves value and access instead of increasing costs and limiting coverage, we will have squandered a significant opportunity. We look forward to continuing working with you to advance the integration of health care AI. 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

---

[14] Rojas JC, Fahrenbach J, Makhni S, Cook SC, Williams JS, Umscheid CA, Chin MH. Framework for integrating equity into machine learning models: a case study. Chest. 2022 Feb 7.
[15] Antequera A, Lawson DO, Noorduyn SG, Dewidar O, et al. Improving Social Justice in COVID-19 Health Research: Interim Guidelines for Reporting Health Equity in Observational Studies. *Int J Environ Res Public Health.* 2021;*18*(17), 9357. doi: 10.3390/ijerph18179357.
[16] Welch VA, Norheim OF, Jull J, et al.: CONSORT-Equity 2017 extension and elaboration for better reporting of health equity in randomised trials. *BMJ*. 2017;359:j5085. doi: https://doi.org/10.1136/bmj.j5085.
[17] The Toronto Declaration: Protecting the rights to equality and non-discrimination in machine learning systems - Amnesty International
[18] Wawira Gichoya J, McCoy LG, Celi LA, Ghassemi M. Equity in essence: a call for operationalising fairness in machine learning for healthcare. BMJ Health Care Inform. 2021 Apr;28(1):e100289. doi: 10.1136/bmjhci-2020-100289. PMID: 33910923; PMCID: PMC8733939.
[19] Norori N, Hu Q, Aellen FM, Faraci FD, Tzovara A. Addressing bias in big data and AI for health care: A call for open science. Patterns. 2021 Oct 8;2(10):100347.
[20] Zhang J, Shu Y, Yu H. Fairness in design: a framework for facilitating ethical artificial intelligence designs. International Journal of Crowd Science. 2023 Mar;7(1):32-9.