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

October 1, 2021

Lynne Parker, PhD
Co-Chair
National Artificial Intelligence Research Resource Task Force
National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22314

Erwin Gianchandani, PhD
Co-Chair
National Artificial Intelligence Research Resource Task Force
National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22314

Re: RFI Response–National AI Research Resource

Dear Drs. Parker and Gianchandani:

On behalf of our physician and medical student members, the American Medical Association (AMA) appreciates the opportunity to provide comments to the National Artificial Intelligence Research Resource (NAIRR) Task Force on their implementation roadmap. Appropriate development, regulation, and clinical integration of health care artificial, or "augmented," intelligence is a priority for the AMA given the rapid pace of innovation in this space and its impact on patient care.

New algorithmic-based health care tools in both the clinical and administrative space are coming to market nearly every day. Unfortunately, the regulatory system to ensure safety and efficacy of clinical tools and appropriateness of administrative tools have not yet been able to keep pace with the rapid rate of development, leaving physicians, patients, and health care facilities and systems without clear guidance on which tools are of sufficiently high quality and can be safely and appropriately integrated in clinical practice and administrative workflows. The AMA strongly supports development of new regulatory systems that are better suited towards meeting the unique challenges and considerations presented by health care augmented intelligence; however, new oversight structures must ensure that AI-based tools brought to market are safe, efficacious, include appropriate data privacy protections, and meet the goals of the quadruple aim—enhance the patient experience of care and outcomes, improve population health, reduce overall costs for the health care system while increasing value, and support the professional satisfaction of physicians and the health care team.

The AMA has developed policy to guide the development, regulation, and use of health care augmented intelligence and has also developed privacy principles to ensure appropriate protections for sensitive patient information are in place. **We strongly encourage NAIRR to consider the unique needs of patients, physicians, and other medical professionals when developing its implementation roadmap.**

We have attached the AMA's policies and principles to help guide the Task Force's work as it relates to health care augmented intelligence. If you should have any questions or would like to discuss further, please contact Shannon Curtis, Assistant Director of Federal Affairs at Shannon.Curtis@ama-assn.org.

Sincerely,

James L. Madara, MD

Attachments

AMA
AMERICAN MEDICAL ASSOCIATION

James L. Madara, MD
CEO, EXECUTIVE VICE PRESIDENT

james.madara@ama-assn.org

May 3, 2021

David Meyers, MD
Agency for Healthcare Research and Quality
5600 Fishers Lane
Rockville, MD 20857

Dear Dr. Meyers:

On behalf of the American Medical Association (AMA) and its physician and medical student members, I am responding to the Agency for Healthcare Research and Quality's (AHRQ) request for information "Use of Clinical Algorithms That Have the Potential To Introduce Racial/Ethnic Bias Into Healthcare Delivery." The AMA has long recognized that racial and ethnic health inequities are an unjust and major public health reality in the United States. Understanding that race is a social and political construct and not a biological risk factor for disease and death, the AMA has publicly acknowledged that racism impacts public health and is a barrier to effective medical diagnosis and treatment. We share AHRQ's concerns that clinical algorithms may inappropriately incorporate race or ethnicity into its recommendations and believe the efforts of AHRQ's Evidence-based Practice Center (EPC) in identifying algorithms of concern and potential solutions are extremely important in helping to advance equity in health.

Advancing equity in health requires the understanding and acceptance of the harmful impacts of historical and contemporary racism on our individual and collective ability to strive for and achieve a reality in which we all have the resources, conditions, opportunities, and power to thrive and achieve optimal health. The AMA is strongly committed to achieving these goals and addressing such issues is a top priority for our organization. Specifically, we recommend that clinicians and researchers focus on genetics and biology, structural racism, and other structural determinants; and collect, report, and use race data as a proxy for structural racism and not ancestry, when describing risk factors for disease and outcomes. Below, we address certain questions in AHRQ's request, numbered according to how they were numbered in the notice issuing the request.

1. What clinical algorithms are used in clinical practice, hospitals, health systems, payment systems, or other instances? What is the estimated impact of these algorithms in size and characteristics of population affected, quality of care, clinical outcomes, quality of life, and health disparities?

The AMA agrees with AHRQ that gathering additional information on the clinical algorithms in use today and whether they factor race and ethnicity into their calculations and the impact so doing may have on health care is of utmost importance. Collection of additional information on these specific algorithms is an essential early step towards identifying where racism and bias may exist in clinical decision-making tools and how they should be addressed to ensure clinical care and health of historically marginalized communities are not negatively impacted by their application. There are many clinical algorithms in use across health care and among the medical specialties. Given that the approaches in design and implementation, as well as the underlying data provenance, vary, it will be important to seek further input from medical specialty societies and other organizations that have expertise and direct experience with development and use of specific algorithms.

2. Do the algorithms in question 1 include race/ethnicity as a variable and, if so, how was race and ethnicity defined (including from whose perspective and whether there is a designation for mixed-race or multiracial individuals)?

The usage of race and ethnicity as variables, and how both are defined, varies among the clinical algorithms in use today. This is attributable in part to changes in protocols over time, as some of the clinical data registries from which algorithms are derived are more than several decades old. There is also variation among multiple health data systems in *how* the data is collected (are race and ethnicity patient- or investigator/clinician reported) and the number of choices provided to the reporter including options such as reporting mixed-race, "other," or an individual's preference to not report. Furthermore, because race is a social construct, there is significant variability in how "races" are defined by society, lawmakers, and others. These definitions have changed and evolved in usage and application over time. Accordingly, their inclusion as variables creates challenges in developing meaningful consensus definitions, especially as our society diversifies over time, further clouding how we define these variables.

3. Do the algorithms in question 1 include measures of social determinants of health (SDOH) and, if so, how were these defined? Are these independently or collectively examined for their potential contribution to healthcare disparities and biases in care?

The AMA defers to the stewards of the data and data platforms (e.g., electronic health record systems; clinical data registries) to provide information as to the degree to which SDOH are collected, applied, and examined during the development and implementation of clinical algorithms. The AMA is actively engaged in the Gravity Project, which was created to develop FHIR-based standards to capture and exchange SDOH information.[1]

5. For the algorithms in question 1, what approaches are used in updating these algorithms?

Medical societies and related organizations are taking action to address concerns associated with the **potential for racist and/or biased** clinical algorithms. For example, the American Society of Nephrology (ASN) and the National Kidney Foundation (NKF) have formed a task force to review and reconsider the inclusion of race in the estimated glomerular filtration rate (eGFR).[2] The current algorithm for eGFR has been called into question for its inclusion of a "correction" for Black patients that estimates a higher kidney function than non-Black patients.[3] Not only is the science that this adjustment is based on questionable, but because the algorithm erroneously overestimates kidney function for Black patients, it may improperly lead to delays in and withholding of care. The stated goal of the ASN and NKF task force is to ensure that the eGFR tool provides an "unbiased assessment of kidney function so that laboratories, clinicians, patients, and public health officials can make informed decisions to ensure equity and personalized care for patients with kidney diseases." Efforts like this one, involving stakeholders from medical societies, patient organizations, and related specialists, can provide insights and a potential framework for the meaningful review of clinical algorithms and their potential for perpetuating medical racism and bias in clinical decision-making.

The AMA is also mindful of advances in scholarship (e.g., public health, critical race theory, and social epidemiology) which call for a shift from thinking of race as a biological risk factor for disease to a deeper understanding of racism as a determinant of health. Thoughtful reconsideration must include an

[1] http://www.hl7.org/gravity/.
[2] https://waysandmeans.house.gov/sites/democrats.waysandmeans.house.gov/files/documents/20.9.25%20ASN%20Response%20to%20Chairman%20Neal%20re%20Race%20and%20eGFR.pdf.
[3] Vyas DA, Eisenstein LG, Jones DS. Hidden in plain Sight—reconsidering the use of race correction in clinical algorithms. *N Engl J Med*. 2020 Aug 27;383(9):874-82.

examination of the underlying social conditions that contribute to health outcomes, including how systemic racism has created and shaped such social conditions as failure to address and remedy social risk factors will hinder efforts to reduce and eliminate the health inequities long associated with racism and discrimination. While this review is underway, clinicians should be encouraged to use their clinical judgment to determine if race correction is warranted, and thoughtfully consider if they may improve or exacerbate inequities for individual patients and populations.[4]

The AMA House of Delegates in November 2020 passed new policy directing our organization "to collaborate with appropriate stakeholders and content experts to develop recommendations on how to interpret or improve clinical algorithms that currently include race-based correction factors."[5] The AMA is currently undertaking an effort to convene a variety of organizations to gather more information about the use of clinical algorithms and create an action plan for how to address these problems. The AMA looks forward to supporting, encouraging, and coordinating its efforts with these organizations to both better understand the algorithms in use today and how they can be improved upon to ensure they help drive equitable care.

We believe that, in addition to efforts like our own, AHRQ is ideally situated to conduct and fund additional research into the use of race and ethnicity data in clinical settings and algorithms, their potential contribution to medical racism and/or bias in clinical decision-making, and the methods needed to eliminate such racism and/or bias.

6. Which clinical algorithms have evidence that they contribute to healthcare disparities, including decreasing access to care, quality of care or worsening health outcomes for Black, Indigenous, and other people of color? What are the priority populations or conditions for assessing whether algorithms increase racial/ethnic disparities? What are the mechanisms by which use of algorithms contribute to poor care for Black, Indigenous, and other people of color?

Perhaps the most well-known example of an algorithm that contributed to inequitable care is the Optum algorithm which used health care cost as a proxy for health, resulting in lower risk scores being assigned to Black patients who were equally sick to similarly situated white patients.[6] While assigning higher risk scores to patients who have higher health care costs may have seemed reasonable to the developers because higher health costs are often associated with greater needs, doing so failed to account for the systemic and long-standing inequities in care that have resulted in fewer expenditures on Black patients. This resulted in further inequitable care, including excluding black patients from care management programs that dedicate additional resources to coordinate care for higher risk programs.

The use of race or ethnicity in clinical algorithms used in cardiology, nephrology, obstetrics, and urology, among others, have been questioned and subjected to close scrutiny. The recent reviews of the questionable use of race and ethnicity in these algorithms has led to efforts to reassess the use of such data by a variety of related groups.[4] While these reviews are ongoing, it is clear that a comprehensive assessment of the use of race and ethnicity data in clinical algorithms is vital to understand the extent of current use and ensure that their inclusion does not re-enforce pre-existing inequities in care.

Importantly, highlighting the questionable use of race and ethnicity in clinical algorithms has led to an opportunity to re-consider the impact of clinical algorithms broadly and focus attention on where

---

[4] Id.
[5] Racial Essentialism in Medicine D-350.981, https://policysearch.ama-assn.org/policyfinder/detail/algorithm?uri=%2FAMADoc%2Fdirectives.xml-D-350.981.xml.
[6] Obermeyer et al, "Dissecting racial bias in an algorithm used to manage the health of populations," *Science* (Oct 25, 2019), https://science.sciencemag.org/content/366/6464/447.

additional research is needed. In addition to gathering known examples of clinical algorithms, it is essential to identify where gaps in knowledge continue to exist. It is clear, for example, that investigation into algorithmic bias can be hampered by the fact that many algorithms used in artificial intelligence (AI)-driven platforms are proprietary, with a lack of transparency on the data sets and information underlying their output.[7] It is also necessary to consider the role training data plays in contributing to clinical algorithmic bias. Researchers are starting to identify errors in commonly used machine learning training datasets. One study conducted by computer scientists from the Massachusetts Institute of Technology found that 3.4 percent of the data was inaccurate or mislabeled.[8] Discrepancies in the quality and accuracy of training data can manifest itself in algorithmic system use. Implicit bias may contribute to human mislabeling practices and subsequent training data sets. This is particularly concerning when using trained algorithms in health care settings or in conjunction with other AI tools or services.

As described above, and in response to resolutions passed by the AMA House of Delegates in November 2020, the AMA is undertaking an effort to convene a variety of organizations to gather more information about the use of clinical algorithms and create an action plan for how to address these problems. We believe that in addition to efforts like our own, AHRQ is ideally situated to conduct and fund additional research into the use of race and ethnicity data in clinical settings and algorithms, their contribution to health inequities, and the methods needed to eliminate bias.

7. To what extent are users of algorithms, including clinicians, health systems, and health plans, aware of the inclusion of race/ethnicity or other variables that could introduce bias in these algorithms and the implications for clinical decision making? What evidence is available about the degree to which the use of clinical algorithms contributes to bias in care delivery and resulting disparities in health outcomes? To what extent are patients aware of the inclusion of race/ethnicity or other variables that can result in bias in algorithms that influence their care? Do providers or health systems communicate this information with patients in ways that can be understood?

It is likely that some but not all users of algorithms, including clinicians, health systems, and health plans, are aware of the inclusion of race/ethnicity or other variables. The AMA does not have any evidence of its own to know the degree to which users of these algorithms are aware that inclusion of those variables could introduce bias and have implications for clinical decisions. Clinicians who input the data required by the algorithm to produce an output may or may not realize that this could contribute to biased results. Clinicians depend on the developers, some of which are the specialty societies to which they belong, to validate the clinical algorithms.

Patients have a fundamental right to know the risk, benefit, indications, and alternatives, including to not proceed, of any health care intervention that they are considering, and physicians have a fundamental obligation to ensure their patient's consent to care is well-informed. Yet, most patients have no idea when an algorithm is being used to inform their care or provide possible treatment options. They are not aware of the types of variables—including race or ethnicity—that go into clinical algorithms, what those variables truly represent, and what impact those variables may have on their care. Furthermore, they may not be aware of if and when their data is used to contribute to the development of AI and machine learning tools. Additionally, as medicine grapples with questions of whether and how to best inform individuals of the use of algorithms, AI, and machine learning tools in their care delivery, we note that the AMA's Privacy Principles state that individuals should have the right to know whether their data will be used to develop and/or train machines or algorithms.[9] The opportunity to participate in data collection for

[7] Id.
[8] https://www.engadget.com/mit-datasets-ai-machine-learning-label-errors-040042574.html.
[9] https://www.ama-assn.org/system/files/2020-05/privacy-principles.pdf.

these purposes must be on an opt-in basis. We encourage AHRQ to urge Congress to adopt these concepts into any forthcoming federal privacy legislation.

8. What are approaches to identifying sources of bias and/or correcting or developing new algorithms that may be free of bias? What evidence, data quality and types (such as claims/utilization data, clinical data, information on social determinants of health), data sources, and sample size are used in their development and validation? What is the impact of these new approaches and algorithms on outcomes?

To ensure appropriate care for patients who may not have received it, perhaps because health care decisions were based at least in part on the output of a clinical algorithm, an individual's clinical status and therapeutic options can be reviewed and revised as appropriate during their follow-up evaluation and management visits. If risk calculators are used, clinicians could adopt an approach similar to the example given by the authors of the *New England Journal of Medicine* article[10] for coronary bypass surgery, changing what is entered for race and ethnicity into the algorithm, and discussing with the patient any observed differences in absolute risk that might be based on race. In most clinical situations, what matters most to patients, their families, and physicians is the absolute risk of any proposed intervention. A similar exercise applied on a population basis could provide insights as to sources of bias.

Medical specialty societies and other organizations that have expertise and direct experience with development and use of specific algorithms will be critical to developing recommendations on how to identify, interpret, and improve clinical algorithms that currently include race-based correction factors. Changing the types of data used to train algorithms and the labels of such data may be one way to reduce racial bias in clinical algorithms; however, ensuring such labels are consistent with patient self-identification and do not exacerbate inequities requires "in-depth understanding of how structural discrimination operates in society," which may not be front of mind—or within the expertise of—all health researchers and clinicians.[11]

Some scholars note that attempts to make algorithms "race neutral" by eliminating race as a variable are insufficient; rather, researchers must "anticipate the structural bias in a dataset or the social implications of a product" and take a "proactive, explicitly anti-racist approach to data collection, analysis and prediction."[12] As an overarching matter, to best prevent and combat the influences of racism and bias in clinical algorithms, we specifically recommend that genetics and biology, the experience of racism, and social determinants of health—not race—be used in clinical algorithms contemplating disease risk factors. As has been found in the reassessment of measures of renal function, the substitution of one biomarker, serum creatinine, by another (cystatin) has the potential to eliminate the inherent bias when serum creatinine levels are inappropriately adjusted based on race. It would advance efforts to eliminate race from clinical algorithms if AHRQ is able to identify similar alternative variables.

10. What are existing and developing standards (national and international) about how clinical algorithms should be developed, validated, and updated in a way to avoid bias? Are you aware of guidance on the inclusion or race/ethnicity, related variables such as SDOH, prior utilization, or other variables to minimize the risk of bias?

---

[10] Vyas DA, Eisenstein LG, Jones DS. Hidden in plain Sight—reconsidering the use of race correction in clinical algorithms. *N Engl J Med*. 2020 Aug 27;383(9):874-82.

[11] Owens, K., Walker, A. Those designing healthcare algorithms must become actively anti-racist. *Nat Med* 26, 1327–1328 (2020). https://doi.org/10.1038/s41591-020-1020-3.

[12] Id.

AMA Privacy Principles

As Congress continues discussions around federal privacy legislation, the AMA seeks to ensure that resulting privacy law protects the sacred trust at the heart of the physician-patient relationship. Specifically, the AMA is working to ensure that as health information is shared—particularly outside of the health care system—patients have meaningful controls over and a clear understanding of how their data is being used and with whom it is being shared. Above all, patients must feel confident that their health information will remain private. Preserving patient trust is critical.

These principles, derived primarily from AMA HOD policy, will serve as the foundation for AMA advocacy on privacy. They are meant to apply to entities other than those already considered covered entities under HIPAA—in other words, physicians generally would not be subject to additional regulation. The principles take into consideration that some data historically not considered "personal" may in fact be personally identifiable (e.g., IP addresses, advertising identifiers from mobile phones). Accordingly, the Principles' use of the term "data" includes information that can be used to identify an individual, even if it is not descriptive on its face.

The Principles provide individuals with rights and protections from discrimination and shift the responsibility for privacy from individuals to data holders other than HIPAA-covered entities (collectively referred to in this document as entities). 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 Principles also call for robust enforcement of penalties for violation of rights to help patients develop and maintain trust in digital health tools, including the use of smartphone applications (apps) to access their own health information.

Individual Rights:

1. Individuals have the right to know exactly what data of theirs an entity is accessing, using, disclosing, and processing—and for what purpose—at or before the point of collection.
2. Individuals have the right to control how entities access, use, process, and disclose their data, including secondary (and beyond) uses.
3. Individuals should be notified within a reasonable period of time following a material change in the entity's data access, use, disclosure, and processing practices.
4. Individuals have a right to direct entities to not sell or otherwise share data about them.
5. Individuals and entities should be able to protect and securely share pieces of information on a granular, as opposed to a document, level.
6. Individuals have a right to direct an entity to delete their data across the entity's ecosystem of services, including when the entity goes out of business or is bought out by another entity (with potential narrowly delineated exceptions, as determined by regulatory bodies and consistent with stakeholder input.
7. Individuals have the right to access and extract their data from a platform in a machine-readable format.
8. Individuals should have the right to know whether their data will be used to develop and/or train machines or algorithms. The opportunity to participate in data collection for these purposes must be on an opt-in basis.
9. Individuals should have a private right of action against entities that are subject to these requirements if the FTC and/or state Attorney General declines to pursue enforcement.
10. Privacy rights should be honored unless they are waived by an individual in a meaningful way, the information is appropriately de-identified (using techniques that are demonstrably robust, scalable, transparent, and provable), or in rare instances when strong countervailing interests in

public health or safety justify invasions of privacy or breaches of confidentiality and, in such case, to the minimum extent necessary.

11. Disclosures of an individual's data should be limited to that information, portion of the medical record, or abstract necessary to fulfill the immediate and specific purpose of disclosure.
12. Individuals who access their medical records using apps should have mechanisms to annotate—but not change—the copy of the record they hold. These mechanisms should track who made the annotation, when, how, and why.

Equity:

1. Privacy protections should promote equity and justice.
2. Health care information is one of the most personal types of information an individual can possess and generate—regardless of whether it is legally defined as "sensitive" or protected health information under HIPAA—and individuals accessing, processing, selling, and using it without the individual's best interest at heart can cause irreparable harm.
3. Individuals should be protected from discrimination, stigmatization, discriminatory profiling, and exploitation occurring during collection and processing of data, and resulting from use and sharing of data, with particular attention paid to minoritized and marginalized (vulnerable) communities. Similarly, individuals should be protected from discrimination, stigmatization, profiling, and exploitation based on inferences drawn from a refusal to use or cessation of use of an app or digital health tool.
4. Because low-income individuals and other vulnerable populations have fewer resources and tools at their disposal to effectively assert their privacy rights, purchase technology with the most advanced and up-to-date privacy and security technology, and recover from harmful invasions of privacy, privacy frameworks (legal or otherwise) must advance policies to benefit individuals of all income levels. For example, the AMA would not support a policy in which paid apps provided greater privacy protections than free apps.
5. Law enforcement agencies requesting medical information should be given access to such information only with a court order and if the law enforcement entity has shown, by clear and convincing evidence, that the information sought is necessary to a specific, legitimate law enforcement inquiry; that the needs of the law enforcement authority cannot be satisfied by non-identifiable health information or by any other information; and that the law enforcement need for the information outweighs the privacy interest of the individual to whom the information pertains. Any applicable legal requirements for law enforcement access to medical information imposed by federal, state, or local laws shall apply in addition to this principle.
6. Employers and insurers should be barred from unconsented access to identifiable medical information to assure that knowledge of sensitive facts does not form the basis of adverse decisions against individuals, such as non-coverage of stigmatized health conditions.
7. Privacy legislation should provide robust and comprehensive protections against genetic discrimination and misuse of genetic information.

Entity Responsibility:

1. All entities that maintain an individual's health information should have an obligation or "duty of loyalty" to the individual, including the duty to maintain the confidentiality of that information.
2. An entity must disclose to individuals exactly what data it is collecting and the purpose for its collection. Such information should not be used for a materially different purpose than those disclosed in the notice at the point of collection of such information. For example, an entity that collects location data to provide weather should not use that data for advertising.

3. Entities should only collect the minimum amount of information needed for a particular purpose, in accordance with regulation and/or federal guidance. For example, a weather app may need general location data (e.g., zip code), but not precise location data (e.g., GPS coordinates).
4. Entities should establish and make publicly available a data retention policy with established protocols for retaining information for operational or regulatory compliance needs.
5. Entities should be required to disclose to individuals what specific elements of data they collect, why, how often, for what purpose, and specifically with whom they are sharing the data.
6. Privacy policies should be written to promote understanding by individuals with elementary school levels of reading comprehension. Terms should be clearly defined and unambiguous. For example, statements such as, "We may share this data with our partners to improve quality" are vague and should not be permitted.
7. Entities should be prohibited from using health data to discriminate against individuals, including creation of "risk scores" that could hinder patients and their families from receiving health, disability, or life insurance; housing; employment; or access to other social services.
8. Entities should make their de-identification processes and techniques publicly available.

Applicability:

1. Privacy legislation should apply to entities that access, use, transmit, and disclose data, including HIPAA business associates, with exceptions for HIPAA-covered entities given their obligations under existing HIPAA regulation. We believe this framework would lead to enhanced transparency around the use of business associates in health care, particularly now that entities not traditionally associated with health care are more active in the health care industry.
2. Local, state, and federally sponsored registries, as well as medical specialty-run registries, should be deemed in compliance with new privacy legislation if they establish a Data Governance Council. The Data Governance Council must include patient representatives and establish practices around sharing registry data. Note that some health conditions (such as HIV or substance use disorder) may have additional, more restrictive privacy safeguards, including through state law. This principle is not intended to replace those protections.
3. Privacy legislation should be adaptable to many different organizations, technologies, sectors, and uses to promote competition. It should also be scalable to organizations of all sizes and be platform- and technology-agnostic and customizable.
4. We recognize the potential need for accommodations for small businesses in certain scenarios, but overall privacy principles should apply to them as they do to larger businesses. For example, an entity with fewer than 10 employees may not need a full-time privacy officer but must still be able to satisfy responses to individuals with questions about the entity's data practices.
5. Privacy legislation should promote data access needed for narrowly delineated medical or public health research or quality improvement and accreditation activities by clinicians and researchers, including open access to appropriate machine-readable public data, while prioritizing the development of a culture that informs individuals about the potential benefits and risks of sharing data with external partners, explicit communication of allowable use with periodic review of informed consent, and protections against using data to deny or limit access to coverage.

Enforcement:

1. Individuals should not be responsible for costs of enforcement unless they are exercising their private right of action (in permitted instances where the Federal Trade Commission (FTC) and the individual's State Attorney General (AG) do not enforce).
2. Federal privacy legislation should serve as a federal floor, not a ceiling.
3. Legislation should not weaken any state's laws or regulations regarding privacy.

4. State Attorneys General (AGs) should be permitted to bring an action in federal court to enforce these requirements on behalf of their states' residents.
5. Federal privacy legislation should authorize funds for FTC to investigate violations of an individual's privacy protections, with a report back to Congress identifying investigation outcomes and trends.
6. Federal legislation should expand Section 5 of the FTC Act to include "manipulative", "abusive", and/or "coercive" behaviors (i.e., behaviors that aren't outright deceptive or causing significant harm, but nevertheless designed to convince people to act against their best interest for the benefit of the entity—for example, dark patterns).
7. Legislation should provide the FTC with Administrative Procedures Act (APA) rulemaking authority, specifically including the ability of FTC to define:
    a. Unfair data processing practices (e.g., processing biometric or geospatial data that are not required for use of the app);
    b. Additional safeguards for certain categories of information (contemplates future-gazing scenarios like human augmentation, cloning, etc.);
    c. Boundaries of data systems;
    d. Minimum privacy and security standards for products that process or use an individual's data (can help with privacy/security being built into the design of apps/products – known as "privacy by design");
    e. The minimum data elements needed for particular purposes;
    f. To the extent appropriate, narrowly delineated exceptions to data deletion rights;
    g. Matters related to patient consent (how to define, what is informed and meaningful, etc.). We firmly believe that "all or nothing" consent is meaningless and would not support such consent acting as a safe harbor from an entity's responsibilities under the statute and regulations; and
    h. Mitigating and aggravating factors for establishing fine/penalty amounts (for example, penalties would be steeper for reckless disregard and knowing/willful conduct). FTC should have authority to impose penalties on both the entity and its officers.

There are many existing and developing standards. Specifically, three existing standards, CONSORT, SPIRIT, and TRIPOD-ML are referenced by a developing standard, MINIMAR (MINimum Information for Medical AI Reporting), by Hernandez-Boussard, et al.[13] CONSORT is Consolidated Standards of Reporting Trials (25-item) for clinical trials. PIRIT is Standard Protocol Items: Recommendations of Intervention Trials (33-item) and is a checklist for interventional trials. These two standards will be extending their checklist to include guidelines for machine leaning (ML) AI components, which will complement a new initiative from TRIPOD, TRIPOD_ML, the Transparent Reporting of a Multivariable Prediction Model for Individual Prognosis and Diagnosis for Machine Learning. MINIMAR is feeding into these initiatives and is proposed to be adopted as a standard. It will help the dissemination of algorithms across health care systems and provide transparency to address potential biases and unintended consequences.

The clinical algorithms should be created based on data collected in a "datasheet for datasets." A new format to collect such data is presented as a "Data Nutrition Label format" by Holland, et. al.[14] Data is a fundamental ingredient in algorithms, clinical decision support and AI; the quality of a dataset used to build a model directly influences the results it produces. The nutrition label format caters to a wide range of requirements for and information available from a specific dataset. During label generation and subsequent updates, it also accommodates data specialists of different backgrounds and technical skill levels to select and to prompt data analysis, development, and validation. The label is built with scalability in mind, and with an eye towards standardization. It provides flexibility for dataset authors and publishers to identify the "right" kind and amount of information to include in a label; over time, this could become a set of domain-specific best practices.

11. To what extent are users of clinical algorithms educated about how algorithms are developed or may influence their decision making? What educational curricula and training is available for clinicians that addresses bias in clinical algorithms?

The AMA believes it is vital that all providers understand how the clinical algorithms they rely on to provide appropriate and equitable care in practice are developed. The need for such understanding is particularly acute as to how algorithms developed using artificial intelligence are trained in order to understand the appropriate uses for and limitations of such algorithms. Having this understanding will help ensure appropriate utilization of algorithms and encourage effective oversight by regulators, providers, and others. Over-reliance on any algorithm, particularly without an understanding of what its most effective uses are, can create a risk for amplifying and perpetuating biases that are present in the data, including any bias based in race or ethnicity.

The AMA expects that physicians will turn to their usual trusted sources of clinical information in their field, most commonly their specialty societies. The AMA already has formal agreements with several of the societies to jointly develop educational content and can reach out to them for this purpose. Through the societies and the AMA, physician awareness can be increased as to the strengths and limitations of clinical algorithms, how they can be best applied to patient care, and how to communicate nuanced information to patients. As specialty societies and physicians caring for patients focus on risk assessment of individuals, important contributions will be made by the public health community as it continues to focus on measuring risk at the population health level using instruments such as community health needs

[13] Hernandez-Boussard, Tina, Bozkurt, Selen, Ionaidis, John P. A., Shah, Nigam H, "MINIMAR (MINimum Information for Medical AI Reporting): Developing reporting standards for artificial intelligence in health care," Journal of the American Medical Informatics Association, Volume 27, Issue 12, December 2020, pages 2011-2015, https://academic.oup.com/jamia/article/27/12/2011/5864179.

[14] Holland, S., Hosny, A., Newman, S., Joseph, J., & Chmielinski, K. (2018). The Dataset Nutrition Label: A Framework To Drive Higher Data Quality Standards. *ArXiv, abs/1805.03677*, https://arxiv.org/ftp/arxiv/papers/1805/1805.03677.pdf.

assessments. Social epidemiologists will be focusing on measuring health inequities (e.g., excess deaths, mortality rate ratios). Estimates of individual, community and population level risk will all benefit when race and ethnicity data are no longer used as proxies for the actual contributors to risk such as racism, health inequities, and social risk factors.

Sincerely,

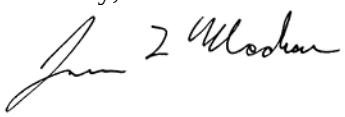

James L. Madara, MD

Policy

The American Medical Association House of Delegates has adopted policies to keep the focus on advancing the role of augmented intelligence (AI) in enhancing patient care, improving population health, reducing overall costs, increasing value and the support of professional satisfaction for physicians.

Foundational policy *Annual 2018*

As a leader in American medicine, our AMA has a unique opportunity to ensure that the evolution of AI in medicine benefits patients, physicians and the health care community. To that end our AMA seeks to:

- Leverage ongoing engagement in digital health and other priority areas for improving patient outcomes and physician professional satisfaction to help set priorities for health care AI
- Identify opportunities to integrate practicing physicians' perspectives into the development, design, validation and implementation of health care AI
- Promote 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
  - Is transparent
  - Conforms to leading standards for reproducibility
  - Identifies and takes steps to address bias and avoids introducing or exacerbating health care disparities, including when testing or deploying new AI tools on vulnerable populations

> "Medical experts are working to determine the clinical applications of AI—work that will guide health care in the future. These experts, along with physicians, state and federal officials must find the path that ends with better outcomes for patients. We have to make sure the technology does not get ahead of our humanity and creativity as physicians."
>
> —Gerald E. Harmon, MD, AMA Board of Trustees

  - Safeguards patients' and other individuals' privacy interests and preserves the security and integrity of personal information
- Encourage education for patients, physicians, medical students, other health care professionals and health administrators to promote greater understanding of the promise and limitations of health care AI
- Explore the legal implications of health care AI, such as issues of liability or intellectual property, and advocate for appropriate professional and governmental oversight for safe, effective, and equitable use of and access to health care AI

Regulation, payment, liability and other key policies
Annual 2019

Our AMA supports the use and payment of AI systems that advance the quadruple aim. 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 while increasing value, and (4) support the professional satisfaction of physicians and the health care team.

Regulation, payment and deployment

Our AMA advocates that:

- AI is designed to enhance human intelligence and the patient-physician relationship rather than replace it

- Oversight and regulation of health care AI systems must be based on risk of harm and benefit accounting for a host of factors, including but not limited to: intended and reasonably expected use(s); evidence of safety, efficacy and equity, including addressing bias; AI system methods; level of automation; transparency; and conditions of deployment
- Payment and coverage for all health care AI systems must be conditioned on complying with all appropriate federal and state laws and regulations, including but not limited to those governing patient safety, efficacy, equity, truthful claims, privacy and security, as well as state medical practice and licensure laws
- Payment and coverage for health care AI systems intended for clinical care must be conditioned on
  - Clinical validation
  - Alignment with clinical decision-making that is familiar to physicians
  - High-quality clinical evidence
- Payment and coverage for health care AI systems must
  - Be informed by real-world workflow and human-centered design principles
  - Enable physicians to prepare for and transition to new care delivery models
  - Support effective communication and engagement between patients, physicians and the health care team
  - Seamlessly integrate clinical, administrative and population health management functions into workflow
  - Seek end-user feedback to support iterative product improvement
- Payment and coverage policies must advance affordability and access to AI systems that are designed for small physician practices and patients and not limited to large practices and institutions
- Government-conferred exclusivities and intellectual property laws are meant to foster innovation, but constitute interventions into the free market, and therefore should be appropriately balanced with the need for competition, access and affordability

Penalties and mandates

Our AMA advocates that:

- Physicians should not be penalized if they do not use AI systems while regulatory oversight, standards, clinical validation, clinical usefulness and standards of care are in flux

Our AMA opposes:

- Policies by payers, hospitals, health systems or governmental entities that mandate use of health care AI systems as a condition of licensure, participation, payment or coverage
- The imposition of costs associated with acquisition, implementation and maintenance of health care AI systems on physicians without sufficient payment

Liability

Our AMA advocates that:

- Liability and incentives should be aligned so that the individual(s) or entity(ies) best positioned to know the AI system risks and best positioned to avert or mitigate harm do so through design, development, validation and implementation
- 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

Role of physician organizations

Our AMA advocates that our organization, national medical specialty societies and state medical associations:

- Identify areas of medical practice where AI systems would advance the quadruple aim
- Leverage existing expertise to ensure clinical validation and clinical assessment of clinical applications of AI systems by medical experts

> "To realize the benefits for patient care, physicians must have the skills to work comfortably with augmented intelligence in health care. Just as working effectively with electronic health records is now part of training for medical students and residents, educating physicians to work effectively with AI systems or, more narrowly, the AI algorithms that can inform clinical care decisions will be critical to the future of ai in health care."
>
> —Bobby Mukkamala, MD, AMA Board of Trustees

- Outline new professional roles and capacities required to aid and guide health care AI systems
- Develop practice guidelines for clinical applications of AI systems

National and state collaboration and strategic planning

Our AMA advocates that:

- There should be federal and state interagency collaboration with participation of the physician community and other stakeholders to advance the broader infrastructural capabilities and requirements necessary for AI solutions in health care to be sufficiently inclusive to benefit all patients, physicians and other health care stakeholders.

Education, professional development and other key policies *Annual 2019*

Citing the potential to improve both the quantity and quality of patient care, the AMA House of Delegates has adopted policy that further examines AI and its potential to benefit physicians, including those who are in training.

The AMA House of Delegates has directed that our AMA encourage:

- Accrediting and licensing bodies to study how AI should be most appropriately addressed in accrediting and licensing standards
- Medical specialty societies and boards to consider production of specialty-specific educational modules related to AI
- Research regarding the effectiveness of AI instruction in medical education on learning and clinical outcomes
- Institutions and programs to be deliberative in the determination of when AI-assisted technologies should be taught, including consideration of established evidence-based treatments, and including consideration regarding what other curricula may need to be eliminated in order to accommodate new training modules
- Stakeholders to provide educational materials to help learners guard against inadvertent dissemination of bias that may be inherent in AI systems
- The study of how differences in institutional access to AI may impact disparities in education for students at schools with fewer resources and less access to AI technologies
- Enhanced training across the continuum of medical education regarding assessment, understanding and application of data in the care of patients
- The study of how disparities in AI educational resources may impact health care disparities for patients in communities with fewer resources and less access to AI technologies
- Institutional leaders and academic deans to proactively accelerate the inclusion of non-clinicians, such as data scientists and engineers, onto their faculty rosters to assist learners in their understanding and use of AI
- Close collaboration with and oversight by practicing physicians in the development of AI applications

And, the AMA House of Delegates reaffirmed AMA Policy D-295.328, "Promoting Physician Lifelong Learning."

For more information: ama-assn.org/ai