



601 Pennsylvania Avenue, NW T 202.778.3200  
South Building, Suite 500 F 202.331.7487  
Washington, D.C. 20004 ahip.org

July 7, 2023

Arati Prabhakar, Ph.D  
The White House  
Office of Science and Technology Policy

*Submitted via regulations.gov*

**RE: Request for Information (RFI) Response: National Priorities for Artificial Intelligence (OSTP-TECH-2023-0007-0001) —AHIP Comments**

Dear Dr. Prabhakar:

Artificial Intelligence (AI) has the potential to offer Americans great improvements in health care affordability, access, and outcomes. However, as the use of AI in health care grows, we agree there needs to be a robust regulatory framework to protect consumers and guard against bias, especially for high-risk applications. AHIP applauds OSTP for taking a proactive approach to advance responsible and trustworthy AI and appreciates the opportunity to provide feedback through this RFI. AHIP is the national association whose members provide health care coverage, services, and solutions to hundreds of millions of Americans every day. We are committed to market-based solutions and public-private partnerships that make health care better and coverage more affordable and accessible for everyone.

Our comments focus on the use of AI in the health care and health insurance sectors. AHIP's members are using AI to transform health care and administrative processes to benefit the people they serve. For example, identifying gaps in evidence-based care to share with providers, improving consumer experience, speeding claims processes, and detecting fraud. At the same time, our members are committed to ensuring the application of AI is safe, transparent, explainable, and ethical. AHIP and its members also seek to ensure biases are neither perpetuated nor introduced in the development and application of AI that could negatively impact certain subpopulations.

Based on that perspective, AHIP has been collaborating with public and private entities to lead the way in protecting consumers yet fostering AI. AHIP has joined forces with business and technology leaders as well as consumer advocates to advance principles, best practices, and industry standards. As part of these initiatives, health insurance providers are seeking ways in which they can allow consumers to direct how their information is used, improve privacy and security, mitigate potential implicit data bias, establish governance best practices, and achieve other shared objectives.

We have also been engaging with federal and state lawmakers and policymakers on legislation and regulation to ensure policies are effective at achieving the stated goals, consistent with our shared values, and reasonable to implement. We continue to be concerned that there are many different governmental entities tackling responsible and trustworthy AI implementation, with little public coordination among them. Moreover, we are concerned that without a cohesive federal approach, states will create their own patchwork of policies and requirements that will be difficult to implement for health insurance providers that offer products across state lines.

AI touches every industry and every facet of life. As OSTP develops a National AI Strategy, we urge the Administration to do so through robust collaboration and partnership. We believe OSTP could play an important role in coordinating efforts across the various federal departments and agencies as well as the states and the private sector to implement a cohesive and coordinated response to the evolving landscape. Policymakers should take an ‘all-of-government’ approach towards regulating AI by establishing cross sector guardrails but permitting innovation within those confines and establishing requirements on a use-case and sector-specific basis using a risk-based approach. These policies should leverage not supplant existing industry standards and regulatory frameworks where possible. Specifically, we recommend that OSTP consider including the following in a National AI Strategy:

- **Fostering public awareness of the uses of AI and measures to ensure safety, transparency, explainability, and fairness.** Consumers should know what AI is, how it is used, and what recourse they have if they believe it is misused. Resources should also be available on the potential benefits (e.g., clinical advancements) and drawbacks (e.g., secondary uses of data). Americans will be more trusting of AI-based services with easy to access, relevant information.
- **Maintaining U.S. leadership.** The strategy should seek to stay ahead of foreign actors and governments that attempt to leverage AI to outpace the U.S. from a competitive standpoint or for nefarious activities such as exploiting individuals’ privacy and security.
- **Balancing innovation with consumer protections.** The National AI Strategy should balance protecting consumers with encouraging innovations. Federal efforts should focus on disclosure, transparency, and auditing rather than registration, licensure, or accreditation. The use of AI is quickly becoming ubiquitous and one size-fits all reviews of all applications would be both time and cost prohibitive.
- **Implementing a risk-based oversight framework.** The National AI Strategy should rely on a regulatory framework that considers the potential risks and benefits of the application in determining the appropriate solutions. The use of autonomous AI has a different level of threat than automating clinical algorithms. Moreover, the use of AI for diagnostics has different implications than for chatbots on a website. Flexibility to right-size business practices and mitigation techniques based on risk is necessary to realize the potential of AI while avoiding overly restrictive, infeasible, or misaligned policies that risk stifling innovation.
- **Ensuring that policies are applicable to the context, scope, and data use of a specific use case.** OSTP may be an appropriate agency to develop baseline policies to regulate AI, but other agencies should build expertise and take the lead on sector-specific expectations. For example, the Department of Health & Human Services should further develop its capabilities and capacity on uses of AI in health care, how these tools can promote the health and wellbeing of Americans, and what risks exists to develop guidance and regulations specific to the health care industry.
- **Preventing the perpetuation and introduction of bias.** Being transparent about potential bias in AI data, methods, and applications is an important step in mitigating harmful unintended consequences. Identified “bias” in underlying data, for example, can encourage “corrections” such as seeking new data sources and elements. It can also foster initiatives designed to benefit

July 7, 2023  
Page 3

specific groups or populations that have been historically marginalized, which might be considered “good bias.” Care should be taken to ensure that efforts to avoid harmful bias do not impede deliberate and beneficial efforts to identify and benefit an underserved group or population. For example, integrating race and ethnicity into predictive analytics could aid in ferreting out disparities and developing mitigating techniques.

The use of AI holds great promise for improving health care for all Americans. AHIP believes that through public-private partnership we can address the challenges posed by using AI while promoting innovation and maintaining American leadership. Engaging a diverse set of stakeholders is essential to this success. Thus, we thank you for the opportunity to provide feedback on the RFI and attach our detailed comments in response to the specific questions posed. AHIP and our health insurance provider members look forward to working with you and other stakeholders on these important efforts. If you have any questions, please reach out to me at either [dlloyd@ahip.org](mailto:dlloyd@ahip.org) or 202-778-3246.

Sincerely,

Danielle A. Lloyd  
Senior Vice President, Private Market Innovations & Quality Initiatives

## AHIP Attachment

**Question: What specific measures – such as standards, regulations, investments, and improved trust and safety practices – are needed to ensure that AI systems are designed, developed, and deployed in a manner that protects people’s rights and safety? Which specific entities should develop and implement these measures?**

Artificial Intelligence (AI) usage by health insurance providers occurs today against an important backdrop of existing laws and regulations at both the federal and state levels such as those touching on privacy, security, discrimination, and interoperability. There are also numerous ongoing voluntary activities among industry participants, such as the development and adherence to principles, best practices, and standards. Any further action by the federal government should take into consideration what is already in place and underway to complement rather than duplicate these activities.

To date, the National Institute of Standards and Technology (NIST) has been a federal leader in this area and should serve as a resource for individuals and entities as substantive and technical policies are developed and adopted. For example, NIST has developed AI standards<sup>1</sup>, in conjunction with the private sector and academia, which address societal and ethical issues, governance, and privacy policies and principles. They have also conducted several AI standards-related efforts<sup>2</sup> that together create fair and comprehensive resources that are in wide use by the health insurance industry. In particular, health insurance providers rely on NIST’s AI Risk Management Framework<sup>3</sup> as it employs a risk-based approach and provides some consistency across industries.

The private sector has also created governance, ethical, and practice standards for organizations developing and deploying AI. For example, AHIP worked with the Consumer Technology Association (CTA) and the American National Standards Institute (ANSI), on the recent development of the consensus-driven standard, ANSI/CTA-2090, The Use of Artificial Intelligence in Health Care: Trustworthiness,<sup>4</sup> considers three elements for how trust can be created and maintained:

- **Human Trust**— focuses on fostering humanistic factors that affect the creation and maintenance of trust between the developer and users. Specifically, human trust is built upon human interaction, the ability to easily explain, user experience and levels of autonomy of the AI solution.
- **Technical Trust**— focuses on the technical execution of the design and training of an AI system to deliver results as expected. Technical trust can also be defined by considerations for data quality and integrity including issues of bias, data security, privacy, source and access.
- **Regulatory Trust**— is gained through compliance by industry based upon clear laws and regulations. This trust can be based upon information from regulatory agencies, federal and state laws and accreditation boards and international standardization frameworks.

AHIP participates in CTA’s efforts to develop additional standards building off the Trustworthiness standard that focuses on avoiding bias in the development and use of AI.

As another example, the Association for the Advancement of Artificial Intelligence (AAAI) is a nonprofit scientific society devoted to advancing the scientific understanding of the mechanisms underlying thought

---

<sup>1</sup> <https://www.nist.gov/artificial-intelligence/plan-federal-ai-standards-engagement>

<sup>2</sup> <https://www.nist.gov/artificial-intelligence/technical-ai-standards>

<sup>3</sup> <https://doi.org/10.6028/NIST.AI.100-1>

<sup>4</sup> <https://shop.cta.tech/products/the-use-of-artificial-intelligence-in-healthcare-trustworthiness-cta-2090>

and intelligent behavior and their embodiment in machines. AAAI aims to promote research in, and responsible use of, AI and increasing the public's understanding of AI. AAAI developed a Code of Professional Ethics and Conduct to inspire and guide the ethical conduct of all AI professionals in an impactful way, while serving as a basis for remediation when violations occur.

A National AI Strategy should focus on building international leadership, fostering awareness, promoting transparency, and monitoring outcomes. This work should build off voluntary, risk-based, consensus-driven standards that focus on policies and procedures, training, risk assessment, monitoring, response, and other facets of development and deployment. The federal government should not micromanage the day-to-day operations, but rather set basic expectations and establish boundaries within which industry can innovate. Future policy based on a National AI Strategy should focus on transparency to the end user (not publicly) of data sources, basic methods employed, and its intended purpose. A National AI strategy should focus on oversight to ensure compliance with consumer protections such as consumer safety and ensure that bias has not led to disparities in outcomes. The federal government's role is to monitor the world-wide environment, track trends in uses and vulnerabilities, and conduct large-scale research to safely advance the field in the US. Building public awareness of such research along with compliance requirements and consumer protections should be included in the National AI Strategy.

To accomplish all of this, OSTP should collaborate across the disparate federal departments and agencies, states, and private sector organizations like AHIP to chart a cohesive course forward that best fosters American innovation while ensuring our values of transparency, consumer protections and privacy are met.

AI will touch every industry and every facet of life. A National AI Strategy should consider ways to use a risk-based framework to develop a streamlined approach to oversight. OSTP may be an appropriate agency to develop baseline policies to regulate AI but other agencies should develop expertise specific to their subject matter. For example, the Department of Health & Human Services should work towards developing subject matter expertise on AI in healthcare and how these tools can be used to promote the health and wellbeing of Americans (and relatedly how they should be regulated for the healthcare industry). AI will impact every industry, and policymakers should take an 'all-of-government' approach towards regulating these tools on a use-specific basis, leveraging existing industry standards and regulatory frameworks to the fullest extent possible.

**Question: How can the principles and practices for identifying and mitigating risks from AI, as outlined in the Blueprint for an AI Bill of Rights and the AI Risk Management Framework, be leveraged most effectively to tackle harms posed by the development and use of specific types of AI systems, such as large language models?**

AI has potential to improve the health care Americans receive as well as their overall health. However, we must balance the quick pace and scope of innovations that will advance the economy with consumer protections including safety, privacy, transparency and other American values. AI systems and policies should be based on ethical principles that respect the rights of individuals and protect populations and set best practices for public and private entities to follow ethical approaches. AHIP and our members are committed to advancing the ethical implementation of AI through public and private collaborations. For example, AHIP is participating in the Center for Practical Bioethics AI Project to proactively identify ways in which to ensure ethical development and use. Ethics are an integral component to AI systems, which will strengthen individual and organizational trust in the software techniques, methods, applications, and outcomes.

To that end, consumers deserve assurances that AI systems are trustworthy and reliable. We agree with and support the fairness, accountability, and transparency principles included in the AI Risk Management Framework. For example, we support stakeholder sign-off and model activities. We also agree with the AI Bill of Rights that systems must be safe and effective and with the AI Risk Management Framework that trustworthy AI depends upon accountability. Earning consumer trust will be essential to the successful use of AI in all areas, but particularly in health care.

We furthermore agree with the AI Bill of Rights and the AI Risk Management Framework that known adverse biases, for example a lack of inclusiveness in data used to train AI, should be made transparent and that there should be consumer protections from algorithmic discrimination. However, a National AI Strategy must be careful not to unduly slow modernization by creating a one-size fits all, highly regulated environment for the use of AI. A risk-based framework will be essential to balancing the potential harms and benefits of the use of AI for a specific function. Just within health care there are dramatically different levels of risk depending on intended use. For example, AI that offers clinical decision-support services involves greater risks to the patient than AI that helps with routine administrative functions such as appointment reminders.

Finally, we agree that generative AI based on large language models warrants a closer look than many other forms of AI. However, it again depends on the specific use case. For example, generative AI filling in what it thinks a consumer is searching for within a health care application is fairly low risk. Generative AI pre-filling a prior authorization request that requires physician sign off might be of moderate risk. Whereas generative AI authoring a recommended course of treatment in a medical record, even with physician sign off, would be of higher risk. Thus, we urge OSTP to consider base expectations across sectors at a high level but leave specific policies to the relevant federal departments and agencies for the sector.

**Question: Are there forms of voluntary or mandatory oversight of AI systems that would help mitigate risk? Can inspiration be drawn from analogous or instructive models of risk management in other sectors, such as laws and policies that promote oversight through registration, incentives, certification, or licensing?**

As a general matter, policies and programs such as registration, certification and/or licensing risk stifling innovation, preventing the beneficial use of AI, and hampering America's ability to compete on the international stage. Satisfying such requirements for each use of AI would not be feasible given its ubiquitous nature. Nor would it be practical given there are no federal policies against which to assess. Finally, it would not be reasonable given the expense of such efforts at that scale. Federal requirements for disclosure, transparency, and auditing might better balance risk and rewards. For example, transparency to demonstrate risk mitigation may be feasible while allowing customers to opt-out of services that use AI in any capacity could cause adverse consequences.

The Office of the National Coordinator for Health Information Technology (ONC) recently proposed a criterion for the Health Information Technology Certification Program<sup>5</sup> (Certification Program) that represents an initial effort by the federal government into overseeing AI. We support ONC's approach of requiring transparency rather than external review of clinical decision support technologies that engage or interact with certified health IT. However, we are concerned about the potential downstream implications for decision support technologies used by health insurance providers for administrative functions, not

---

<sup>5</sup> <https://www.healthit.gov/topic/certification-ehrs/about-onc-health-it-certification-program>

clinical care. Ultimately, organizations should be able to make their own risk-based decisions. Federal requirements should focus on transparency of those decisions, the existence of appropriate governance and risk management controls, the measurement of outcomes to detect unintended consequences, and establishment of mitigation techniques when necessary.

**Question: What will the principal benefits of AI be for the people of the United States? How can the United States best capture the benefits of AI across the economy, in domains such as education, health, and transportation? How can AI be harnessed to improve consumer access to and reduce costs associated with products and services? How can AI be used to increase competition and lower barriers to entry across the economy?**

AI has many advantages that include operational efficiencies, cost reduction, technical innovation and error reduction. Health insurance providers currently use AI to benefit patients and consumers, improving care and creating efficiencies while minimizing opportunities for fraud. However, we believe the full extent of the benefits of AI applications have yet to be realized. Some examples of how health insurance providers use AI include:

- Cleaning, normalizing, and labeling data for use in various programs.
- Clinical models to understand health conditions and disease progression through research.
- Identification of gaps in the provision of evidenced-based care.
- Predictive analytics to identify patients who may benefit from improved access to services (excluding factors such as historical spending that may introduce disparate impact or bias).
- Physician performance to identify high-value care for use in consumer choice and network design.
- Service models to enhance the customer experience.
- Prior Authorization to identify data included in electronic medical records and streamline requests and approvals (with clinicians involved throughout the process).
- Actuarial analysis to help identify utilization patterns (not individually identifiable) for an employer or other health plan sponsor to understand usage trends both now and for the future.
- Claims analysis to increase efficiency and to identify potential fraud and abuse.
- Provider directories to improve the accuracy of included data elements.
- Market research regarding prospective employer sponsors of health insurance to help determine which employers might align with a company's product offerings, value, areas of access, or health care networks.

Health insurance providers are using AI to identify clinical risk and improve customer experience. AI has the potential to both improve patient access to care and reduce administrative costs. However, AI is used to augment, not replace human decision making and expertise. For example, AI can help facilitate a streamlined prior authorization process by using algorithms to issue approvals. In these tools, AI is not used to deny prior authorization requests, simply to follow the same algorithms that a person would use to approve a request. This allows health care providers and patients to receive approvals more quickly, while health insurance providers can focus their experts on cases that do not allow for a straightforward decision.

**Question: What are the opportunities for AI to enhance equity and how can these be fostered? For example, what are the potential benefits for AI in enabling broadened prosperity, expanding economic and educational opportunity, increasing access to services, and advancing civil rights?**

AI is an important tool for advancing health equity and reducing disparities in the healthcare system. As noted above, health insurance providers can use predictive analytics to identify patients at risk. AI could

also power clinical models to help identify potential disparities. Improving access to care will be essential to reducing disparities. Technology powered by AI can play an important role in improving the availability and accessibility of care, two key dimensions of access.<sup>6</sup> For example, AI-powered chatbots can assist enrollees at any time of the day from any location where the internet is available. Remote-patient monitoring systems paired with AI can help health care providers track their patients' progress and intervene if necessary.

**Question: What are the unique considerations for understanding the impacts of AI systems on underserved communities and particular groups, such as minors and people with disabilities? Are there additional considerations and safeguards that are important for preventing barriers to using these systems and protecting the rights and safety of these groups?**

Being transparent about potential bias in AI data, methods, and applications is an important step in mitigating harmful unintended consequences. Identified “bias” in underlying data, for example, can encourage “corrections” such as seeking new data sources and elements. It can also foster initiatives designed to benefit specific groups or populations that have been historically marginalized, which might be considered “good bias.” Care should be taken to ensure that efforts to avoid harmful bias do not impede deliberate and beneficial efforts to identify and benefit an underserved group or population. For example, integrating race and ethnicity into predictive analytics could aid in ferreting out disparities and developing mitigating techniques. Rather than prohibiting the use of variables in a model that could be beneficial in certain cases, inclusion of such variable should be allowed but with safeguards such as appropriate governance and documentation of the use of such variables. For example, race was an important variable to include in efforts in tracking and trending COVID-19 rates and in assisting people with vaccination access. Engaging a diverse set of stakeholders who will be impacted by the AI in the design of the use cases is also a good way to better understand AI programs, promote beneficial advancements, and mitigate harmful unintended outcomes as much as possible.

It is possible, and sometimes necessary to “tune” machine learning (“ML”) models to have good bias so they work exceptionally well for specific groups of people who need intervention or support, or perhaps it is not realized that a targeted intervention is needed until the ML uncovers the need. These groups could be defined by having certain diseases, living in certain geographies, or of a certain socioeconomic status. ML, and the AI behind it, provides us the opportunity to optimize models to perform best among populations that are most vulnerable or historically suffer from implicit bias and resulting disparate impacts.

While regulations may aim to prohibit algorithmic bias, these same types of rules could unintentionally limit the ability of entities to use AI to combat systemic inequality by targeting health interventions at specific underserved communities – including at-risk and historically marginalized groups. As policymakers and regulators consider what safeguards should exist to protect underserved communities, they should be mindful of potential unintended consequences and limiting “good bias” that may aim to serve these communities.

**Question: What additional considerations or measures are needed to assure that AI mitigates algorithmic discrimination, advances equal opportunity, and promotes positive outcomes for all, especially when developed and used in specific domains (e.g., in health and human services, in hiring and employment practices, in transportation)?**

---

<sup>6</sup> Penchansky R, Thomas JW. The concept of access: definition and relationship to consumer satisfaction. *Med Care*. 1981 Feb;19(2):127-40. doi: 10.1097/00005650-198102000-00001. PMID: 7206846.



*Improve the Availability and Quality of Demographic Data.* AI depends on its underlying data. In order for AI to function correctly, the set or sets of data and other elements to calculate and achieve what has been programmed can be relied on as part of the function and methodology. Improving demographic data standards and collection will allow us to build AI that works for more people and prevent unintended bias that can result from not representing all people in the data used to build and train AI. Health insurance providers obtain consumer demographic data in two ways: (1) indirect data estimations and (2) direct data collection.

Health plans use indirect estimation methods of race and ethnicity, such as Geographic Assignment based on Census or American Community Survey Data or the Bayesian Indirect Surname Geocoding, to support the identification of disparities within communities that they serve. These indirect methods estimate percentages of racial and ethnic groups in populations based on Census percentages in certain areas or, individual race and ethnicity assignment is based on an individual's surname and their geographic location based on Census estimates. Health plans do not use indirect data methods for determining patient level care needs given their reduced accuracy and lack of trust with consumers and other stakeholders, especially as the U.S. population becomes more diverse.

Underscoring the challenges is the need to address issues of trust about the sharing of potentially sensitive data among consumers and clinicians with health insurance providers. Consumers are never obligated to provide demographic data to their health insurance providers—it is always voluntary. Consumers' demographic data is never used to determine benefit or enrollment eligibility and an individual's decision to not disclose demographic data will not impact their ability to seek or obtain health care or insurance.

In addition, health insurance plans follow current racial and ethnic data standards required by OMB and CMS. Some health plans collect more granular demographic data beyond what is required by OMB and/or CMS to better understand their populations and provide more culturally and linguistically appropriate and tailored care.

Robust, accurate, actionable, and standardized demographic patient data is fundamental to advancing health equity. Collecting consistent demographic data allows health care entities to better understand the populations they serve and informs more culturally and linguistically appropriate patient-centered care. It also allows health care entities to better identify disparities in care and outcomes as well as understand the social drivers of health to better promote equitable care, devise innovative solutions, operationalize telehealth, and measure the effectiveness of interventions for continuous improvement.

To improve upon existing demographic data standards, AHIP convened diverse groups of health insurance providers and other stakeholders (e.g., patients representing different communities, providers, community-based organizations, and others) for over 18 months from 2020 – 2022 and employed an evidence-based and stakeholder-driven process with goals to:

- Align with national standards as much as possible while improving national standards when necessary;
- Standardize data at a high-level while allowing for local customization and granularity; and
- Aim for actionability while minimizing data burden.

As a result of this work, AHIP developed a set of revised demographic data standards designed to facilitate ease of response while allowing someone to report more granular information if desired. These standards are also designed to use more inclusive language for the collection of data on sexual orientation and gender identity (SOGI) while collecting a more holistic set of information on factors should as disability, language preference, and veteran status.

Enabling the electronic exchange of demographic data through standardized content is pivotal to successful equity efforts. Having interoperable patient demographic data would allow the health care ecosystem to collect this data when most appropriate and convenient for the patient and share the information with other partners with patient consent to inform patient care and population health management efforts as well as to more effectively address disparities in access to care and outcomes. To promote interoperability, AHIP also mapped the demographic data standards to standardized codes (e.g., LOINC, SNOMED, ICD-10) and developed a data documentation that provides guidance on how frequently each question should be asked and how various responses should be coded, particularly when an individual selects “I do not know,” “I choose not to respond,” or when the individual leaves the question blank.

Alignment of demographic data standards at an ecosystem level through such policy changes is crucial to efforts advance equity, particularly through those based on AI. An aligned and standardized approach to interoperable demographic data will empower the health care ecosystem to collaborate on shared health equity goals, measure progress towards those goals, and better serve individuals and communities. With consistent and interoperable data standards, great strides can be made in reducing inequities and addressing social drivers of health while improving outcomes and minimizing the data burden placed on individuals and on the larger health care ecosystem. This data could also be used to develop and train AI to better identify and address potential disparities.

We encourage OSTP to work with OMB, CDC, as well as with established private sector efforts such as the Gravity Project<sup>7</sup> to improve the collection and accuracy of demographic data to support the development of AI that can support all patients and reduce the risk of bias.

*Educate Consumers on AI and its Benefits and Risks.* Education should be done for individual consumers to understand what AI is, how it works, and the advantages/disadvantages of the technology so that trust can be built for practical AI uses.

*Leverage Existing Laws and Regulations.* As policymakers consider if additional safeguards are necessary to protect against algorithmic discrimination, we encourage them to consider how existing laws may already sufficiently protect people. There may be situations where existing frameworks may need to be supplemented to account for AI, but there are robust regulatory frameworks towards addressing discrimination in the U.S. today and regulators have already expressed their intent to leverage them when regulating AI applications. Rather than develop numerous, conflicting laws and regulations, the Federal government and the States should work together to leverage existing policies to foster transparency while preventing harm from AI.

---

<sup>7</sup> <https://thegravityproject.net>