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

September 22, 2023

Senator Bill Cassidy, M.D.
Ranking Member
U.S. Senate Committee on Health, Education, Labor and Pensions
Washington DC 20510-6300

Re: Exploring Congress' Framework for the Future of AI

Submitted via HELPGOP_AIComments@help.senate.gov

Dear Senator Cassidy:

AHIP appreciates the opportunity to provide feedback on the White Paper entitled "*Exploring Congress' Framework for the Future of AI: The Oversight and Legislative Role of Congress Over the Integration of Artificial Intelligence (AI) in Health, Education, and Labor*," ("White Paper") released on September 6, 2023. 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 making health care better and coverage more affordable and accessible for everyone. We applaud your proactive approach to advancing responsible artificial intelligence (AI) for Americans and thank you for the opportunity to comment.

Health insurance providers are leveraging AI to make great improvements in health care affordability, access, and outcomes through broader data access, new technologies, and solutions to benefit our enrollees. For example, insurers are using AI to improve consumer experience, research disease progression to identify patients at risk of decompensation, find opportunities to fill gaps in care, automate prior authorization approvals where possible to speed decisions, and identify fraud and abuse. 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 while 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. For example, AHIP worked with the Consumer Technology Association on its ANSI/CTA-2090, *The Use of Artificial Intelligence in Health Care: Trustworthiness* and its new *Best Practices and Recommendations for Bias Management*.

As the use of AI grows, Americans deserve the peace of mind of knowing that it is being used responsibly and for their benefit. Thus, we agree that there is a need to set guardrails to protect consumers and mitigate bias. However, those guardrails must balance the benefits of AI against the associated risks, and design policies that seek to mitigate those risks in a way that least interferes with beneficial innovation. As the Senate considers strategies to oversee AI, we urge you to consider ways to use a risk-based framework to develop a streamlined approach to oversight.

We believe any legislation addressing AI should permit programs, practices, and procedures to reflect the context, scope, and data use of a specific use case. Private sector organizations need the flexibility to tailor efforts to their unique circumstances given the breadth of potential uses. We agree that a one-size-fits all approach risks overregulation that could negatively impact innovation and jeopardize American leadership. We recommend focusing on promoting appropriate governance, transparency, explainability, privacy, and mitigation techniques through adherence to industry and federal standards. AI is a means to an end; the underlying technology should not be subject to outside reviews or audits, but rather the outcomes associated with its use should be measured and monitored to ensure patients are protected and unintended negative consequences are swiftly addressed.

Thank you for the opportunity to provide feedback. The use of AI holds great promise for transforming health care for all Americans. Engaging a diverse set of stakeholders is essential to success, and AHIP and our health insurance provider members are eager to work with you and other stakeholders on these important efforts. We include more detailed responses to your questions in the attached recommendations. If you or your staff have any questions, please reach out to me at either sgupta@ahip.org or 202-955-4384.

Sincerely,

Sohini Gupta

Sohini Gupta, J.D.
Executive Vice President, Government Affairs and Innovations

Supporting Medical Innovation

How can FDA support the use of AI to design and develop new drugs and biologics?

Artificial Intelligence (AI) could play a crucial role in identifying important new treatments. Reducing the time and cost of developing new drugs and biologics could facilitate access to effective new treatments and improve outcomes. However, the FDA must maintain its rigorous, evidence-based process to ensure that drugs and biologics are safe and effective regardless of whether AI is used to support their development or not. AHIP generally believes that risk-based frameworks that balance risks and benefits to consumers, along with existing protections or processes, should be used to determine the need for regulation and oversight of AI tools. AI used in the development of a drug that could be taken by millions poses a greater risk to consumers than, for example, AI that fills in what it thinks a consumer is searching for within a health care application. However, it could also bring revolutionary benefits to consumers through greater access to breakthrough medications. Further, medications developed with the use of AI tools are subject to existing safety and efficacy reviews by the FDA. The FDA should continue to follow their rigorous evidence-based, scientific process for reviewing and approving new drugs, to ensure that these drugs and biologics meet appropriate standards for safety, replicability, and reliability, while allowing flexibility that allows use of AI in the drug development process to promote innovation.

How can FDA improve the use of AI in medical devices?

We believe the FDA will continue to play a critical role in ensuring medical devices that incorporate AI are safe and effective. While we are generally supportive of the Software as a Medical Device framework the FDA developed, we are concerned that FDA Clinical Decision Support (CDS) Software Guidance released a year ago stifles innovation and may contravene efforts by health plans to effectively manage the health of their members. We strongly recommend that the FDA update this guidance to clarify that only those applications made available for commercial use are subject to requirements applicable to Software as a Medical Device. As noted above, the FDA could also foster the safe application of AI in medical devices by ensuring the use of tools is documented in the submission documents and labeling of the device and ensuring manufacturers provide appropriate education regarding the inclusion, risks, and benefits of the AI embedded in the device are understood by clinicians and consumers.

Moreover, FDA could make optional recommendations that premarket submissions include developers' considerations to address equity, inclusion, and representation issues in their development process, and if there are inclusion features that allow for access for individuals with hearing, visual, and other impairments that may inhibit use. This optional submission will incentivize developers and sponsors to share information that would be useful for those considering potential adoption of the device, while still providing necessary flexibility in development and submission requirements. Additionally, for AI and Machine Learning (ML)-enabled devices, FDA could leverage its "Good Machine Learning Practice for Medical Device Development: Guiding Principles" document and recommend sponsors include how they "promote safe, effective, and high-quality medical devices that use artificial intelligence and

machine learning (AI/ML)", including, at a high-level, steps taken to identify and mitigate risks and potential limitations of the model(s) used in the software device.

What updates to the regulatory frameworks for medical devices should Congress consider to facilitate innovation in AI applications while also ensuring that products are safe and effective for patients?

Health insurance providers support the use of evidence-based digital tools such as AI with sufficient data on real-world applications to improve patient care and access and to promote the use of innovative treatments. However, there has been a rapid proliferation of digital therapies and the FDA only regulates a small subset of such therapies. Given this proliferation, health insurance providers, patients, and clinical providers need to understand which therapies are safe and effective and how these tools can help with individual treatment plans or needs. We need independent, credible organizations to assist with evaluations of these various technologies.

There are already efforts underway to develop an evidence-based evaluation framework for digital therapies.[1] For example, the Agency for Healthcare Research and Quality, the Peterson Center on Healthcare, and the American Psychiatric Association have worked on models that seek to assess and evaluate the safety and functionality of digital tools in health care. We encourage the government to harness the work that is already happening and commit to working with stakeholders, including health insurance providers, to improve the landscape for the continued development and use of digital therapies.

However, the government should also provide clarity on which products will require FDA review. Precise definitions will be key to preventing overregulation. For example, the FDA should clarify and potentially modify its final policy on clinical decision supports and how the agency is interpreting the Criteria in in Section 520(o)(1)(E) of the FD&C Act as the current interpretation may be overly broad and inappropriately subjecting products to FDA review.

How can FDA harness external expertise to support review of products that are developed using AI or that incorporate AI?

We appreciate the need to ensure the FDA has the appropriate expertise to support review of products that are developed using AI or that incorporate AI. To leverage external expertise to bolster the agency's reviews, the FDA could establish an advisory committee focused on AI that is structured similarly other advisory committees to provide technical reviews and advise on AI policy and technical developments.

What are the potential consequences of regulating AI in the United States if it remains unregulated in other countries?

Overregulating AI risks stifling innovation and jeopardizing American leadership. However, we agree there is a need to protect consumers while promoting innovation. Patient safety must be

[1] Anthony Watson, Richard Chapman, Gigi Shafai, and Yuri A. Maricich; "FDA regulations and prescription digital therapeutics: Evolving with the technologies they regulate;" Frontiers in Digital Health (Apr. 17, 2023); available at: https://www.ncbi.nlm.nih.gov/pmc/articles/PMC10150093/.

foremost as the use of AI in health care grows. AI must share health care's fundamental ethos to avoid harm and protect patients. For example, users of AI should be mindful of harm from bias and AI developers should strive to avoid harmful bias in their products.

While we agree there is a need to protect consumers from potential harms from AI, we also urge Congress to avoid actions that could prevent patients from benefiting from AI. Potential legislation and regulation should use a risk-based framework that focuses on promoting transparency about the potential harms and benefits of AI rather than restricting its use.

We appreciate the white paper recognizes the need for a clear definition of AI. We urge Congress to clearly define AI when considering potential legislation. Congress should also work with federal agencies to ensure definitions of AI used in regulations are clear and precise. Overly broad definitions could be misinterpreted or misapplied, and risk regulating products that simply follow algorithms but do not use human-like intelligence or have the capability to learn without being specifically programmed by a human. We recommend alignment with established terms that rely on existing standards. For example, such as National Institute for Science and Technology's (NIST) AI Risk Mitigation Framework's use of "AI System:" an engineered or machine-based system that can, for a given set of objectives, generate outputs such as predictions, recommendations, or decisions influencing real or virtual environments. AI systems are designed to operate with varying levels of autonomy (Adapted from: OECD Recommendation on AI:2019; ISO/IEC 22989:2022).[2]

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

Medical Ethics and Protecting Patients

What existing standards are in place to demonstrate clinical validity when leveraging AI? What gaps exist in those standards?

Tools that leverage AI should be held to the same standards for clinical validity as those that do not use AI. Existing standards on the safety and ethical use of AI could be leveraged to ensure transparency on the inclusion of AI in a health care setting.

What practices are in place to mitigate bias in AI decision-making?

AI is an important tool for advancing health equity and reducing disparities in the health care system. For example, health insurance providers can use predictive analytics to identify disparities in care and connect patients at risk to additional services such as case management. AI can also power clinical models to help identify potential disparities. Improving access to care

[2] https://www.nist.gov/itl/ai-risk-management-framework

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.[3] As noted in the white paper, patient-facing AI can enhance care and improve outcomes. For instance, 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.

One source of biased outcomes from AI results from issues from the underlying data used to program its algorithms.[4] No data set will ever be complete or free of potential biases. Health care data can be particularly challenging given the lack of integration across data systems.[5] An additional source of bias can include human errors or misunderstandings of model impact that results in disparate outcomes. The application of AI-generated information or tools within programs must also be monitored for impact on outcomes. Furthermore, existing societal biases can be further deepened by AI. Being transparent about potential bias in the data used to train AI, methods, and applications is an important step in mitigating harmful unintended consequences. Identified “bias” in underlying data, for example, can encourage mitigation 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

[3] 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.
[4] Nelson, G. S. (2019). Bias in artificial intelligence. North Carolina medical journal, 80(4), 220-222.
[5] Ibid.

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.

"Good bias" is also important in clinical care, when it is clinically appropriate to target certain selected populations to deliver health services within accepted standards of care. Medical care decision-making requires special consideration because it is often appropriately biased towards certain types of patient populations in accordance with accepted standards of care. For example, many types of cancer screening begin at specified ages, and certain types of tests and therapies are offered based on sex identified at birth (e.g., cervical cancer or prostate cancer screening). We caution against any type of regulatory framework that could impede clinical care and the appropriate, intentional selection of population cohorts in accordance with accepted standards of care.

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 aimed at proactively identifying 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.

We appreciate the federal government's work to date to promote safe and ethical AI use. 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 White House Office of Science and Technology Policy (OSTP) AI Bill of Rights that systems must be safe and effective and with the NIST 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 White House Blueprint for an AI Bill of Rights and the NIST 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 unintended algorithmic discrimination. We are concerned, though, that at present there is no agreement on how to measure model bias and how to decide which subgroups for which potential bias is relevant (racial/ethnic groups, age, gender, other), and there is no standardized method to evaluate whether an AI tool is used in a fair and unbiased manner. Thus, we see a major need and opportunity for advancing the science behind identifying, measuring, and mitigating unintended biases in AI tools.

The Office of the National Coordinator for Health Information Technology (ONC) recently proposed a criterion for the Health Information Technology Certification Program[6] (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. ONC has focused on providing information on the populations used to train an algorithm to support consumer and provider decision-making. However, we encourage ONC to continue work with other agencies

[6] https://www.healthit.gov/topic/certification-ehrs/about-onc-health-it-certification-program

and standard development entities to ensure alignment of AI best practices across health care tools.

The private sector has also started to create 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,[7] 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.

In September 2023, CTA released voluntary standards outlining ways to identify and manage AI bias in health care.[8] These standards outline types of biases and suggests strategies to mitigate bias throughout the development process.

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 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.

Additionally, the Brookings Institute recently published the Critical Algorithmic Systems Classification (CASC). The CASC aims to outline a flexible regulatory framework that protects civil and consumer rights that addresses the unique challenges of using AI without changing the structure of the federal government.

As these efforts, and additional ones, continue to evolve and mature, the health care industry and appropriate federal government agencies need to work together to consolidate and coalesce towards a set of common national standards that can be adopted consistently across organizations.

[7] https://shop.cta.tech/products/the-use-of-artificial-intelligence-in-healthcare-trustworthiness-cta-2090
[8] Artificial Intelligence in Health Care: Practices for Identifying and Managing Bias (CTA-2116)

What should be the federal role, if any, in addressing social and/or political bias?

The role of the federal government should focus on building international leadership, fostering awareness, promoting transparency, and monitoring outcomes. Any actions by the federal government should balance both the risks and benefits and set guardrails while allowing organizations to innovate within those guardrails. The federal government's role should be 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. The federal government should facilitate the advancement of the availability of demographic data. More complete data sets will lead to algorithms that are trained on data that is representative of more Americans and permit outcomes research by those employing AI.

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 US 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 mitigating potential harm from AI.

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 to use for assessment. 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 methods may be feasible whereas requiring that customers be allowed to opt-out of any services that use AI in any capacity could cause adverse consequences in health care as it becomes part of the medical standard for care.

Adopt a Whole of Government Approach

AI will touch every industry and every facet of life. Regulation and legislation should use a risk-based framework to develop a streamlined approach to oversight. AI governance and oversight should not be siloed, but rather integrated across government agencies that can consider how to provide the necessary safeguards within the context of industry-specific use cases. AI will impact every industry, and policymakers should take an 'all-of-government' approach towards regulating these tools on a use-specific basis, fully leveraging existing industry standards and regulatory frameworks.

For example, the Department of Health & Human Services should work towards developing subject matter expertise on AI in health care and how these tools can be used to promote the

health and wellbeing of Americans (and relatedly how they should be regulated for the health care industry). Federal agencies should have embedded AI experts and could convene advisory panels and technical expert panels to support agency staff with private sector and other outside experts.

Focus on Transparency

The federal government should focus on promoting transparency rather than regulating how or when AI can be used. A focus on transparency gives stakeholders the information they need to make informed decisions about using AI without preventing innovation or the leveraging of AI by organizations to improve processes and workflows. 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 should focus on transparency to the end user of data sources, basic methods employed, and its intended purpose and audience for the use of the model. However, this should not be required to be made publicly available to balance the user having sufficient information while dampening the risk of exposing proprietary methods. The federal government could focus on providing oversight to ensure compliance with consumer protections such as consumer safety and privacy protection. However, regulations should not be overly prescriptive about what data sets are used to train AI or when and how AI is used. In some cases, it may be appropriate to train AI on a curated dataset to ensure the algorithm meets its intended purpose (e.g., a health insurance provider training AI on data representative of its specific population to predict case management needs), while it other times, a broader dataset could be more useful (e.g., a large language model.) Allowing flexibility while promoting transparency will protect consumers while supporting the development of AI that can meet unique needs and support its intended purpose.

Support Efforts to Improve Data and Advance Equity Research

The federal government could also support broader efforts to improve the data used to train AI and support research into safe and ethical applications of AI. Accurate AI algorithms need broad, diverse, and representative data. The federal government could work with the private sector to improve the collection and accuracy of demographic data to support the development of AI that can support all patients and mitigate bias.

AI is driven by and learns from underlying data to achieve its programmed objectives. Improving demographic data standards and collection will allow us to build AI that works for more people and mitigate 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 which are currently in the process of being updated.[9] 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.[10]

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

[9] https://www.whitehouse.gov/omb/briefing-room/2023/01/26/initial-proposals-for-revising-the-federal-race-and-ethnicity-standards/
[10] https://www.ahip.org/documents/AHIP-Letter-on-Demographic-Data-Standards-with-Appendix.pdf

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.

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.

OMB, CDC, as well as with established private sector efforts such as the Gravity Project[11] should work together improve the collection and accuracy of demographic data to support the development of AI that can support all patients and mitigate bias.

Educate Consumers on AI and its Benefits and Risks

The federal government could also play a leading role in educating consumers about use of AI. Consumers will need education to understand what AI is, how it works, and the potential risks and benefits of the technology. Successful AI use will depend on understanding, transparency, and trust.

How can AI be best adopted to not inappropriately deny patients care?

AI has the potential to revolutionize health care and dramatically decrease the burden administrative tasks placed on clinicians. Ambient AI could facilitate a physician's charting process, saving them time and allowing them to focus on the patient. Natural language processing could find data in unstructured data in an EHR to reduce the burdens of quality measurement and generative AI could write a prior authorization request for a clinician.

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, creating efficiencies, and minimizing 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.

[11] https://thegravityproject.net

- 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.
- Analyzing 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 and to supplement existing clinical data to facilitate better decision-making. For example, AI can help facilitate a streamlined prior authorization process by using algorithms to issue approvals. In these tools, AI is generally only used to approve prior authorization requests by simply applying the same clinical algorithms that a person would use to approve a request. Denials are only automated in very clear-cut situations such as a prior authorization request for a patient that is not covered by that insurer or for a service that does not require prior authorization and, even then, appeals processes are available just in case. The focused use of AI in prior authorizations allows health care providers and patients to receive approvals more quickly, while health insurance providers can focus their experts on complex cases that could result in a denial.

Is the current HIPAA framework equipped to safeguard patient privacy with regards to AI in clinical settings? If not, how not or how to better equip the framework?

Health insurers have a long-standing commitment to data protection. Our enrollees' privacy and their trust are foundational to all that we do. AHIP's member health insurers are subject to the requirements of the HIPAA; the Privacy and Security Rules promulgated under HIPAA (45 C.F.R. Parts 160, 164); the Health Information Technology for Economic and Clinical Health ("HITECH") Act (Pub. L. No. 111-5) and regulatory changes made pursuant to HITECH; and 42 CFR Part 2, and other state and federal privacy laws. This legal framework tightly restricts an insurer's use and handling of an individual's protected health information (PHI) and imposes parameters on permitted uses and disclosures for essential functions such as treatment, payment, and healthcare operations. Even when data is collected outside of the designated record set, health insurers generally treat those data with the same robust protections as if it were HIPAA-covered. While these protections are robust, in the unfortunate situation of a breach of unsecured PHI, covered entities must report it to the Department of Health and Human Services, which in turns posts those breaches affecting 500 or more individuals for complete consumer

transparency. The Privacy Rule also provides consumers with broad rights related to their health information, including the right to request correction or amendment of incorrect information, data portability, and the right to request restrictions or disclosures of their information.

We appreciate the white paper recognizing the risks from third-party applications that currently reside outside of HIPAA. While technology and third-party applications present new opportunities, they also present new risks to patient privacy. New entity types, such as app developers, that are now common in the marketplace were not contemplated, let alone included, as covered entities within the traditional privacy laws and regulations. These new entities can be collecting, using, disclosing, and disseminating consumer health data for the purpose of monetizing it, and yet consumers are generally unaware that the data are outside these regulatory protection frameworks. In particular, it is concerning that AI can compound these risks through its ability to aggregate large-scale and disparate data permitting re-identification of previously de-identified data under HIPAA.

AHIP and its members believe that personal health information should be protected no matter who holds the data. Defining a federal approach for the privacy and security of all healthcare data would help ensure consistent protection of health information in a variety of situations and avoid a patchwork of state approaches that could result in gaps and vulnerability. Digital tools not currently subject to HIPAA should be subject to similar robust privacy law giving consumers peace of mind that all of their healthcare data is safe, secure, and private.

Who should be responsible for determining safe and appropriate applications of AI algorithms?

Consumers deserve assurances that AI systems are trustworthy and reliable. The federal government could play an important role in setting guardrails and monitoring for unintended consequences. AHIP and its members support the fairness, accountability, and transparency principles included in the NIST AI Risk Management Framework such as the stakeholder sign-off and model activities. We also agree with the White House Blueprint for an AI Bill of Rights[12] 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.

However, 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. We do not believe that external reviews of AI technology should be required. Currently there are no standards to review AI and decision support technologies against, there's no current entity that could perform such a review, and such reviews could be cost prohibitive. Rather, any potential future regulations should focus on promoting transparency about how and when an AI algorithm is trained and mitigating harm from potential rather than requiring external reviews.

Who should be liable for unsafe or inappropriate applications of AI algorithms? The developer? A regulating body? A third party or private entity?

[12] https://www.whitehouse.gov/ostp/ai-bill-of-rights/

AI is an emerging technology and determining potential liability will be complicated and could depend on how and why a harm resulted. For example, a developer could be liable if someone was harmed because their AI product was defective while a physician could be liable for harm caused by an inappropriate use of an otherwise safe AI-enabled device. Current laws should be leveraged as a start while the law adapts to emerging technologies. 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.