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

April 29, 2022

Submitted Via Email to: AIframework@nist.gov

National Institute of Standards and Technology (NIST)
100 Bureau Drive
Stop 2000
Gaithersburg, MD 20899

Re: NIST Artificial Intelligence (AI) Risk Management Framework (RMF): Initial Draft (Issued March 17, 2022)

Dear NIST Representatives:

American patients and consumers can see great improvements in health care affordability, access, and outcomes through new technologies and solutions that incorporate AI. AHIP[1] appreciates the opportunity to submit comments in response to the Initial Draft of the AI RMF. This release is timely given the national focus on AI development priorities announced by the White House Office of Science and Technology Policy, State activities and the National Association of Insurance Commissioners (NAIC) focus in this area, and the burgeoning use of and AI applications in various sectors of the economy, including health care. NIST's work can help inform the national efforts and the work emerging in the States.

Engaging and Informing Consumers About the Benefits and Use of AI

While some Americans may be unfamiliar with AI systems and processes, others recognize and welcome the advantages that AI can bring to help them improve their health and well-being. For example, a recent news article[2] described how AI can more accurately predict and diagnose heart events when compared to traditional methods. Such tools can assist health care providers and patients with conversations about courses of care, identifying appropriate actions, and avoiding unnecessary issues or surgeries.

To the extent that AI can improve a person's health and outcomes, NIST can help by highlighting practical examples. NIST also can take a leading role in educating consumers about

[1] AHIP is the national association whose members provide coverage for health care and related services to hundreds of millions of Americans every day. Through these offerings, we improve and protect the health and financial security of consumers, families, businesses, communities, and the nation. We are committed to market-based solutions and public-private partnerships that improve affordability, value, access, and well-being for consumers.

[2] *See,* https://www.fastcompany.com/90740876/this-ai-can-prevent-your-death-10-years-from-now-so-how-does-that-work.

the benefits and potential risks that AI can present. Education can include easy-to-understand examples of AI use.

Earning the Trust of American Consumers Is Essential to Success

Consumers deserve assurance that AI systems are trustworthy and reliable. We agree with and support the fairness, accountability, and transparency principles included in this Framework. For example, we support stakeholder sign-off and model activities, and we look forward to NIST continuing stakeholder discussions in these areas.

AHIP appreciated being a participant in the NIST Virtual Workshop that was held in March and included a discussion of these principles. We ask NIST to leverage public workshops and continue to provide feedback opportunities. Gathering feedback specific to each industry is essential, as concerns and questions vary greatly – and in health care, decisions have a very personal impact on every American.

In other areas to further earn Americans' trust and comfort with AI, AHIP has been working with the Consumer Technology Association (CTA) on a standard that will complement NIST's work. While the CTA standard has not yet been publicly released, it is expected to include a significant focus on risk management for AI. We believe that NIST will find this standard useful for its own framework, and we suggest that NIST consider referencing this work in future AI RMF versions. We will be happy to share the standard when it is approved for final release.

AI systems and policies should be based on ethical principles that respect the rights of individuals and set ethics best practices for public and private entities. AHIP and our members are committed to advancing the ethical implementation of AI through public and private collaborations. AHIP is participating in the Center for Practical Bioethics AI Project to proactively identify ways in which to ensure ethical use. Ethics are an integral component to AI systems, which will strengthen individual and organizational trust in the software techniques methods and outcomes.

Support for Organizations Will Help Speed Effective Adoption

AHIP strongly supports the AI RMF's goal to provide organizations with the opportunity to evaluate, define and manage their risk thresholds based on their risk-based tolerances. Each organization is unique, and they can define and manage their AI risks based on their system and physical environments, customers' needs, business lines and capabilities, and a corporate culture that assesses and addresses risks in line with legal and compliance responsibilities. Leadership for each organization should be empowered to make decisions based on needs, vulnerabilities, risks, and capability to insure against and to respond to perceived risks. This approach has

worked well under the Health Insurance Portability and Accountability Act (HIPAA), and a similar process can be used for the NIST AI RMF.

We support the flexibility of the framework, and we encourage NIST to continue to allow for flexibility, which will encourage collaboration as AI developments continue. The AI RMF does not - and cannot - cover every potential use case. Therefore, a tiered risk management approach is wise, given that not all AI systems pose the same risks. We believe that guidance and recommendations should be balanced and not overly prescriptive, as each industry has unique considerations across AI use cases. “Lessons learned” across industries should be explained and leveraged, as this area continues to grow and develop.

Our comments below correspond with the themes and concepts laid out in the NIST AI RMF. We understand that the intent of the draft framework will address risks in the design, development, use, and evaluation of AI to assist organizations with understanding and determining risk. The AI RMF will be used in addition to the NIST Cybersecurity and the Privacy frameworks. We agree that the AI RMF should be used as a companion to and should complement other NIST and private sector frameworks. Where possible, NIST can benefit the public's understanding by cross-referencing the concepts and components of the frameworks wherever possible. An alternate approach could be for NIST to issue supplemental guidance to further explain how the various frameworks interrelate and support common goals. We recommend that NIST continue working with stakeholders to provide more clarity and practical examples of real-life uses cases, whenever possible.

Thank you for the opportunity to comment on this important topic. Please contact me at (202) 492-5492 or mzluke@ahip.org if you require additional information.

Sincerely,

Marilyn Zigmund Luke
Vice President

AHIP

Attachment A

Specific Comments in Response to the NIST AI Risk Management Framework

Reference / Section	AI Risk Management Framework	Comments and Recommendations
Technical characteristics 5.1.1 Accuracy	The text explains that, "[a]ccuracy indicates the degree to which the [Machine Learning (] ML [)] model is correctly capturing a relationship that exists within training data . . . examined via standard ML metrics (e.g., false positive and false negative rates, F1-score, precision, and recall), as well as assessment of model underfit or overfit (high testing errors irrespective of error rates in training). AI risk management processes should take into account the potential risks to the enterprise and society if the underlying causal relationship inferred by a model is not valid, calling into question decisions made on the basis of the model. Determining a threshold for accuracy that corresponds with	"Accuracy" in the context of the AI RMF describes a specific model using general methodology using an accuracy test. It is unclear how this methodology could apply within processes that need to identify incorrect or fraudulent results, rather than an accuracy function.[3] For example, high accuracy (or low error) models could be deemed to have "accuracy" because they predict a majority classification (i.e., a percentage of total claims processed). However, in some fraud and abuse detection methods that use AIs, in general terms, a fraudulent claim could still be deemed "accurate" in a system without further investigation. Likewise, AI programs that seek out an incorrect or unusual "outlier" would not be looking for accuracy, and instead would be looking for a fraudulent claim. In a non-programming context, emphasizing the appropriateness of metrics specific to each of the models would be better in the Framework rather than emphasizing "accuracy" which

[3] For reference, "Greater Accuracy Does Not Mean Greater Machine Learning Model Performance," available at: https://towardsdatascience.com/greater-accuracy-does-not-mean-greater-machine-learning-model-performance-771222345e61.

	acceptable risk is fundamental to AI risk management and highly context-dependent."	may be outside the scope of the function or process. We recommend more discussion to highlight this concept and the potential applications. In addition, using the term "accuracy" to establish guidelines can inadvertently convey that the metric "accuracy" should assess the performance of models across all cases. The term may have different meanings for policy makers, AI developers, and practitioners, and thus we recommend NIST consider careful crafting of definitions when determining appropriate terms for the AI RMF technical characteristics. A more precise definition for "predictive accuracy**"** may be better to capture correctness or usefulness depending on intended outcome. **We recommend more discussion to highlight this concept and the potential applications as well as to ensure understanding of definitions. Overall, the text in this section as written is complex and dense. We recommend using more understandable text and using appropriate terms to help non-technical users and readers (e.g., to describe whether a model measures what it intends to measure.)**
5.2.3 Privacy	This section discusses privacy as the norms and practices that help to safeguard values such as	**We recommend expanding this section. The U.S. Constitution, laws and regulations, and a variety of ethics and policies are in**

	human autonomy and dignity such as freedom from intrusion, limiting observation, or individuals' control of facets of their identities (e.g., body, data, reputation). AI systems may promote privacy but are contextual and vary among cultures and individuals.	**place to protect individuals' privacy, which is a right rather than a characteristic, norm, or practice. This section should emphasize the importance of individual privacy and how privacy as a foundational construct must be built into the process.** In addition, referencing the NIST Privacy and Cybersecurity Frameworks and other guidance documents could be helpful for this section.
6. Core	The AI RMF Core provides outcomes and actions that enable dialogue, understanding, and activities to manage AI risks. The Core is composed of three elements: functions, categories, and subcategories.	We support the outline for AI core functions. We note that Figure 6 (AI Lifecyle) does not include AI decommissioning. **We recommend NIST add "decommissioning" to the figure and offer baseline recommendations for risk management considerations when removing AI from use.**
6.1 Map and 6.2 Measure	The Map function establishes the context and applies the attributes of the AI RMF taxonomy to frame risks related to an AI system. The Measure function provides knowledge relevant to the risks associated with attributes of the AI.	**Overall, we support the Map categories and subcategories.** The categories and subcategories help create a consistent baseline of recommendations for evaluating AI across the lifecycle. Typically, solutions have non-AI components that work in conjunction with AI components to create a final AI decision tool. The NIST discussion in this section can be interpreted to mean that the Framework applies only to AI

		components, or alternatively that all components in a system are AI and / or become AI, even if used on a stand-alone basis. **We recommend that NIST work with stakeholders to develop recommendations for how to evaluate non-AI components, which can be integrated or separate functions.**
6.3 Manage	This function addresses risks which have been mapped and measured and are managed in order 3 to maximize benefits and minimize adverse impacts.	**We recommend defining "impact" and "scale" as used in this section. In addition, we support mechanisms for disengaging or deactivating AI that demonstrates outcomes inconsistent with intended use. We recommend expanding this section to include a discussion for creating a contingency plan for the deactivation of AI, and to include resulting business considerations (i.e., ensuring no halt in services during disengaging or deactivation processes).**
6.4 Govern	The Govern function cultivates and implements a culture of risk management within organizations developing, deploying, or acquiring AI systems.	**We support strong governance processes and believe the Framework should contain this important section. While we support sharing relevant details regarding creation and use of algorithms, we also recommend NIST and AI-related entities consider privacy, security, and intellectual property concerns.** Considerations should be made to prevent unintended consequences such as breaches or algorithm corruption by external parties.

		Governance processes may want to evaluate the "who, what, and when" if a breach occurs or an algorithm is corrupted by an external source. AI checklists in alignment with industry requirements could help set expectations to further facilitate standards relating to governance structures.