



Charles N. Kahn III
President and CEO

May 6, 2024

The Honorable Ami Bera, MD
United States House of Representatives
172 Cannon House Office Building
Washington, DC 20515

Dear Representative Bera,

On behalf of the Federation of American Hospitals (FAH), thank you for the opportunity to respond to your request for information on Artificial Intelligence (AI) in health care. As a Member of Congress, cofounder of the Health Care Innovation Caucus, and member of the bipartisan House Task Force on AI, you are uniquely positioned to understand that, if developed and used responsibly, AI has the potential to transform the efficiency of patient care, improve health outcomes, lower costs, and make other advancements in the field. We appreciate your willingness to seek feedback from health care stakeholders to inform decision-making and policy development in this area.

The FAH is the national representative of more than 1,000 leading tax-paying hospitals and health systems throughout the United States. FAH members provide patients and communities with access to high-quality, affordable care in both urban and rural areas across 46 states, plus Washington, DC, and Puerto Rico. Our members include teaching, acute, inpatient rehabilitation, behavioral health, and long-term care hospitals and provide a wide range of inpatient, ambulatory, post-acute, emergency, children's, and cancer services.

The FAH and its members appreciate the promise of AI and the need to carefully balance oversight to ensure safe and appropriate development and use with the need to innovate and continue to advance this potentially transformational technology. Rather than responding to specific questions, we have provided feedback on the four areas included in the RFI: implementation; efficacy, accuracy, and transparency; ethical and regulatory considerations; and other considerations.

Implementation

In considering the development and use of AI tools in health care, it is important to understand how the technology is defined, and particularly whether a tool incorporates generative AI and machine learning (ML) approaches. This is because generative AI and ML generally learn from the data the tools are operating on. By contrast, tools that operationalize set rules are predictable, and in health care, are often used to implement evidence-based clinical decision support or well-

defined operational activities. We note that the Food and Drug Administration (FDA) has approved at least 700 devices with AI aspects. AI can be a stand-alone product or support other medical technologies (such as AI embedded in a drug pump or surgical robot).

Advances in generative AI and ML are creating opportunities and challenges. Health care organizations have begun to pilot the use of generative AI for a number of activities and are beginning to deploy them for many of the same business and operational tasks that other industries have deployed AI, such as supporting human resources functions, helping with scheduling, and offering chatbots to support customers in finding the answers to questions. In the coming years, we anticipate continued adoption for operations and administrative functions, such as coding, billing, and appeals of denied claims.

AI also has potential to reduce administrative and paperwork burdens. These tools are currently being tested to see how well they might support clinical documentation and transitions of care, such as summarizing and organizing patient data, or synthesizing records to ensure that when nurses end their shifts, the next team has all of the relevant information at hand. A number of health care organizations have piloted use of AI tools to recommend documentation for clinicians to review, modify as needed, and approve, as well as initial drafts of patient communications, again for the clinician to review and modify as needed. Use of these tools must still recognize the central role of physicians and other medical professionals in patient care decisions.

Use of generative AI and ML in clinical decision making will require greater scrutiny to ensure safety and efficacy. Health care organizations are in the process of deploying the oversight and governance functions necessary to ensure patient safety and privacy and security of information. Despite these promising trends, the use of generative AI and machine learning tools has challenges, including but not limited to the following:

- Limited access to the large amounts of complete, reliable, and representative data needed to train algorithms;
- Limited ability for the end-user of a tool to verify the accuracy of embedded algorithms;
- Significant effort needed to integrate new tools into existing systems and processes;
- Significant effort needed to identify and manage workflow changes to deploy new tools, including the costs and challenges presented by working with dominant technology vendors;
- Costs of deployment and training;
- Uncertainty about regulatory requirements and how they may change;
- Uncertainty about liability and how it may be shared across tool developers and end-users;
- Protecting intellectual property when working with large language model developers;
- Limited inter-institutional cooperation for necessary large-scale model validation; and,
- Limited expertise in creating models.

The health care industry also continues to struggle with sharing data across entities, a practice called interoperability, in order to have a more complete picture of the patient. The recent cyberattack on Change Healthcare demonstrates the risks of aggregating data under a single umbrella. Interoperability allows for more federated data systems that are less rich targets for bad actors.

Efficacy, Accuracy, and Transparency

Efficacy and Accuracy. It is challenging to test the efficacy and accuracy of generative AI and ML driven health care solutions given that many institutions do not have the architecture to continuously monitor AI solutions. However, the following steps can increase our understanding of and confidence in these tools:

- **Population-based outcomes research to identify benefits and challenges in actual use.** Traditional approaches of comparing training to sample data may not accurately predict the outcomes of use in clinical settings or for specific populations.
- **Assurance labs that “evaluate models using nationwide standards and best practices.”¹** Federal support for these types of organizations would be a positive step but may still lack representative results relevant to specific populations. Steps should be taken to ensure that a wide variety of health care organizations can benefit from such labs.
- **Increased availability of large volumes of accurate and representative health data to support the development of AI-driven health care solutions.** Making these types of resources available to a broader range of health care organizations will require adequate levels of funding and steps to maintain the privacy and security of the information.

Bias, Unintended Outcomes, and Transparency. The FAH shares concerns about the potential risks of AI tools that may inadvertently embed bias or lead to poor patient outcomes. We recognize the risks that automated solutions can pose, including unintended outcomes such as misdiagnosis, biased analyses, inappropriate denials of service by payers, or inappropriate use and disclosure of sensitive health information. For example, the Centers for Medicare & Medicaid Services (CMS) notified Medicare Advantage plans that they cannot use AI tools to deny care without taking into account the unique circumstances of the individual.² Responsible development of AI tools includes the identification and mitigation of risks. Commercial AI tool developers must evaluate the risk of bias in their tools, take appropriate steps to mitigate bias, and communicate the results of testing and any needed cautions to their customers.

As with medications, procedures, or other health care interventions, transparency regarding the action, expected use, and possible cautions about an AI solution will inform appropriate use. Guardrails to ensure responsible development of commercial AI and ML tools should include transparency measures on issues such as how a model works, the data used to train it, appropriate and inappropriate uses of the tool, and results of any testing that has been done to assess bias or other adverse outcomes. This transparency could be in the form of model cards or nutrition labels that provide information about key attributes.

It is important, however, to balance these transparency and risk management approaches with innovation and the risk of unnecessary burden. A health care practitioner will not realistically be able to individually evaluate AI tools and their output in the midst of patient treatment. Careful attention will be needed to ensure that appropriate information is available in the workflow, without creating significant disruptions to the care process.

¹ [A Nationwide Network of Health AI Assurance Laboratories - PubMed \(nih.gov\)](#)

² [CMS Announces Guidance for Medicare Advantage Criteria \(natlawreview.com\)](#)

Ethical and Regulatory Considerations

Shared Responsibility and Developer Accountability. There is a shared responsibility between the developers and end-users of AI tools to build and deploy them in a way that is safe, effective, and secure. However, end-users are reliant on technology developers, who must be held accountable for the safety and reliability of their products and required to be truthful in marketing their products. In addition, commercial AI and ML developers must provide end-users of their tools with guidance on ethical use, such as when it is necessary to have “a human in the loop,” and the limits of their models. End-users also will need guidance on how to provide oversight of AI tools that are in use to ensure that they are functioning appropriately over time.

As health care embraces AI and other analytic tools, it will be important for clinicians and other users to be able to explain the role of algorithms to individuals, including patients, affected by AI-assisted decisions, using information provided by the developer. Explanations provided by the developers should be meaningful and useful, tailored to the audience and calibrated to the level of risk. It may not be realistic to inform patients about all uses of AI in their care, particularly if it is embedded in a medical device or other tool.

Risk management approach. The unique capabilities of generative AI and the automation of repetitive tasks in health care pose unique risks. Risk management is a key aspect of ensuring that AI solutions, generative and rules-based, are appropriately developed, disseminated, and monitored over time. Risk management approaches also would support trustworthy, safe, appropriate, and equitable design.

However, health care providers already deploy risk management approaches to ensure the safety of health care services and the privacy and security of health information. At the federal level, this includes a range of safety requirements, such as the Medicare Conditions of Participation, as well as the Health Insurance Portability and Accountability Act (HIPAA) Privacy and Security Rules. There are many other federal, state, and local laws that address safety, privacy, and security in health care. Some hospitals and health systems deploy and develop AI solutions for use within their own organizations and are deploying governance and risk management processes to support responsible use.

To be most successful at realizing the promise of AI and protecting against negative outcomes, the health care sector will need tools, standards, and guidance to incorporate the use of AI-enabled tools into existing risk management structures. Any AI regulatory requirements that conflict with existing risk management processes will slow down progress in realizing the benefits of technology and could inadvertently result in less effective risk management of complex health care organizations.

Equal access to technology. Access to AI and ML tools may not be equally distributed across types of health care providers. As with the deployment of Electronic Health Records (EHRs), there is a risk that small and rural providers and their patients may be left behind due to the cost and complexity of deploying AI tools. Therefore, it is crucial to encourage and incentivize the dissemination and validation of models to as many systems as possible.

Privacy considerations. “Covered Entities” under HIPAA have existing physical, administrative, and technical approaches to maintaining the security and confidentiality of protected health information (PHI). They are challenged, however, by the continued proliferation of state-level privacy laws that are not pre-empted by HIPAA. We are concerned that the growing use of AI and ML could accelerate the trend toward multiple, sometimes conflicting, requirements across states. Compliance with multiple sets of requirements poses significant challenges for organizations that operate in all states, regional organizations, and local organizations whose patients come from multiple states. This variable treatment of the same health data also fails to serve patients who would be better served by a clearly understood, strong, implementable privacy standard.

We believe that privacy would be enhanced by a true national standard that preempts all state privacy laws that apply to PHI. In addition, any national privacy law outside of HIPAA should exempt all HIPAA-covered entities (not just PHI). The FAH has long advocated for strong national privacy standards and continues to believe that this is the best approach for both PHI and personal health data that fall outside of HIPAA.

Other Considerations

As Congress considers whether and how to best balance the benefits and risks that generative AI and ML pose, it should also consider definitions. Standardized definitions created by a body with appropriate expertise and processes to consult widely with stakeholders, including health care providers, would facilitate a common understanding of the varied technologies that build from advanced computing techniques.

The FAH recommends that Congress continue to engage with end-users of these technologies to gain an understanding of on-the-ground deployments and the realities of a shared responsibility between those who develop these tools and those who use them in the care of patients. Health care providers already operate under a risk management approach in complying with myriad privacy, security, and patient safety rules. The FAH recommends that Congress separate requirements for commercial developers of AI used in health care from those that apply to end-users, keeping in mind that end-users must receive adequate information from the developers. Any new obligations for end-users should be carefully considered and build from the existing risk management approaches that are used in healthcare today.

Congress also will need to give careful consideration to the topic of liability, which is a new and challenging aspect of AI. While health care providers bear responsibility for the care they provide, the developers of commercial AI products must also be accountable if safety, bias, or other harms are caused by a flaw in the tool itself.



The FAH appreciates the opportunity to comment on the Request for Information and we look forward to continued dialogue and partnership as we strive to advance the use of health IT, including AI, to improve our nation’s health care system. If you have any questions regarding our comments, please do not hesitate to contact Ryann Hill at (202) 624-1514.

Sincerely,

Ryana Hill