



Charles N. Kahn III
President and CEO

October 27, 2025

Arati Prabhakar
Director
Office of Science and Technology Policy
The White House
1600 Pennsylvania Avenue NW
Washington, DC 20500

Via electronic submission at: <http://www.regulations.gov>


RE: Request for Information; Regulatory Reform on Artificial Intelligence

Dear Director Prabhakar:

The Federation of American Hospitals (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, and cancer services. The FAH appreciates the opportunity to submit comments to the Office of Science and Technology Policy (OSTP) regarding the Request for Information (RFI); Regulatory Reform on Artificial Intelligence.

AI as applicable to healthcare can include algorithms, tools, and devices that process healthcare-related data, including Personally Identifiable Information (PII) and/or Protected Health Information (PHI), regarding a specific patient as well as from other populations that may be relevant. Key use cases for AI in healthcare focus on alleviating the administrative burden faced by clinicians, providers, and leaders. This allows them to reclaim valuable time to focus on patient care, critical decision making, and high-risk activities such as transitions of care.

In considering the development and use of AI tools in healthcare, it is important to first understand how AI technology is defined, and particularly whether a tool incorporates generative AI and machine learning (ML) approaches. This is because generative AI and ML 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 authorized at least 1,000 devices with AI aspects for marketing in the United States. 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. Healthcare 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 as used in other industries, 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.

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 transformational technology. AI has the potential to revolutionize healthcare delivery, including improving patient care, operational efficiency, and health outcomes. AI also could be used to improve the life cycle of existing systems by eliminating the need for interfaces to ensure data retention and accessibility requirements are met.

We emphasize the unique attributes of healthcare (e.g., size of the industry; impact on health and safety; health data privacy) and the needs of the multiple health sector actors that will interact with AI technologies, such as AI developers, healthcare technologists, healthcare organizations, healthcare providers, and patients/consumers. In addition, we appreciate the RFI's opportunity for engaging in the AI Action Plan development process, and invite further collaboration in developing specific regulations, to take advantage of our deep expertise and practical knowledge in this sector.

We note the need to consider the topic of liability, which is a new and challenging aspect of AI. While healthcare 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 the AI tool itself or a flaw in the tool's development.

We appreciate the RFI's emphasis on identifying "federal regulations that hinder AI development, deployment, or adoption," and with that tenet in mind we make the following recommendations to achieve that goal.

Federal Barriers to AI Development and Deployment

Establish a Uniform Regulatory Framework

The United States has the opportunity at the federal level to establish a uniform and practical framework to promote AI in healthcare, considering the context of existing laws and regulations that can be leveraged as a baseline to ensure safe and effective use of AI in healthcare. Where new laws and regulations are necessary, **we strongly recommend a singular federal regulatory framework, including a national federal privacy law, that expressly pre-empts state laws, for the use of AI in healthcare.** We caution against a patchwork of individual state AI regulations as this will severely limit the ability to deploy AI at scale in healthcare and limit the potential positive impact AI could have on an overall healthcare system's performance. A more uniform approach to AI regulation would catalyze rather than hamper AI development and adoption in healthcare.

This framework should establish standard definitions and benchmarking tools to support effective implementation of AI for healthcare. For example, the National Institute of Standards and Technology (NIST) has a long history of working across the public and private sectors to develop consensus-based, workable standards and tools to advance the use of technologies that improve the lives of individuals. **We recommend that NIST develop a single set of definitions within the area of AI that can**

be leveraged by others, avoiding contradictory legal and regulatory approaches. In addition, NIST has the appropriate technical knowledge and experience to develop benchmarking metrics and other tools that both AI developers and deployers would be able to leverage as they seek to ensure that AI solutions are safe, secure, trustworthy, and fit to purpose in healthcare. For example, NIST could lead development of standards for model testing, validation, benchmarking, and lifecycle monitoring of AI technologies used in healthcare. **We recommend maintaining a database of validated data sets used for AI technologies to facilitate greater explainability for these technologies.**

Regulate AI in Healthcare Utilizing Existing Risk Management Approaches

Risk management is a key aspect of ensuring that AI solutions, generative and rules-based, are appropriately developed, disseminated, and monitored over time. For AI solutions in particular, a risk management approach can help both developers and healthcare providers efficiently focus technical and organizational controls on higher risk deployment.


Risk management approaches are deeply integrated with healthcare systems, including both existing workflows and regulatory schemes. Hospitals and healthcare systems have extensive experience in, and have long deployed, risk management approaches to ensure the safety of healthcare services and the privacy and security of health information. We emphasize the importance of AI as an auxiliary tool to augment human actions, where a human in the loop has final decision-making authority over any actions involving, defining, or executing treatment plans or clinical decisions.

At the federal level, the existing risk management landscape includes a range of safety and privacy requirements, such as the Medicare Conditions of Participation and the Health Insurance Portability and Accountability Act (HIPAA) Privacy and Security Rules. In addition, healthcare technologies have established risk management for electronic medical record (EMR) and health record (EHR) workflows. 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 healthcare systems and organizations. **An AI regulatory framework should take into consideration the extent to which the AI model directly impacts patient care and should focus on processes to ensure algorithms are transparent, auditable, ethical, fair, non-biased, and safe by incorporating the existing risk management framework already established in healthcare.** To that end, as discussed above, it is critical that hospitals and healthcare systems are involved in the creation of AI policies for healthcare, including establishing standards around risk categories for AI use cases.

As a specific example, data privacy and security should be managed under the existing privacy framework established under HIPAA. If prescriptive mandates are pursued regarding transparency, disclosures, and opt-out rights for developers and deployers, we recommend excluding Covered Entities from such requirements by adding a robust and comprehensive safe harbor provision. This will minimize duplicative requirements and ensure alignment with existing regulations under HIPAA through the Department of Health and Human Services' (HHS) oversight.

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,¹ or inappropriate use and disclosure of sensitive health information. Responsible development of AI tools

¹ We agree with and support the Centers for Medicare & Medicaid Services (CMS) guidance to Medicare Advantage plans establishing that they cannot use AI tools to deny care without considering the unique circumstances of the individual. This is a sensible protection to prevent inappropriate denials for medically necessary care.



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 that leads to systematically inaccurate or misleading results, and communicate the results of testing and any needed cautions to their customers.

Clarification of FDA Oversight and SaMD Definition

The FDA's statutory authority under the Federal Food, Drug, and Cosmetic Act (FDCA) extends to products that meet the definition of a "device," which increasingly includes software-based tools. However, this broad definition may inadvertently encompass AI-enabled technologies in which the primary function is ancillary to direct medical diagnosis or treatment. To ensure regulatory clarity and avoid unnecessary burdens on innovation, we recommend that OSTP work with FDA to provide more practical and precise guidance on what constitutes Software as a Medical Device (SaMD). Clarification is needed to exclude certain Clinical Decision Support (CDS) software from the SaMD category.

Concerns Regarding FDA's Draft Guidance on AI/ML-Enabled Device Software Functions

The FDA's Draft Guidance on "AI/ML-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations" outlines expectations for design, validation, modification control, and lifecycle performance monitoring of AI/ML-based medical devices. While the intent to ensure safety and effectiveness is commendable, the draft status of this guidance introduces uncertainty for developers and stakeholders. Organizations must prepare for potentially stringent post-market surveillance, change management, and transparency requirements, which could significantly increase regulatory burden. This is particularly concerning for developers of AI tools that may be classified as SaMD under the current framework. We urge OSTP to support efforts that balance innovation with patient safety, including clear, risk-based criteria for regulatory oversight and streamlined pathways for low-risk AI applications.

Critical Tools for AI Governance

Collaboration with Providers

We urge federal agency collaboration with providers, including hospitals and healthcare systems, when developing an AI framework and regulations. Healthcare providers have strong expertise in matters related to direct patient care, clinician experience, and healthcare operations. Providers also have extensive experience in managing patient data under HIPAA regulations. It is therefore critical that the voices of the hospital and healthcare system community be represented during the creation of AI policies for healthcare.

Technical Standards

We recommend that federal regulations enable interoperability standards that prevent vendor lock-in and foster competition while maintaining security. Federal regulations should support auditability and model validation frameworks to ensure AI outputs are dependable, even if the models are not always fully interpretable. We believe internal governance models to be used by hospitals and healthcare systems and other providers should be in place to monitor, refine, and ensure continuous compliance, while also ensuring there are auditable development practices that actively track and reduce bias, using standard metrics to measure efficacy.

Protecting Intellectual Property

We recommend that federal regulations establish clear guidelines for intellectual property (IP) ownership related to AI-generated content, ensuring that businesses retain rights to innovations supported by AI. Existing IP laws that disproportionately benefit large tech companies and potentially restrict access to AI advancements for other organizations need to be reconsidered. We urge OSTP and NITRD NCO to support equitable licensing frameworks for AI models and training data, allowing access to foundational technologies while respecting proprietary rights. We believe that AI innovation may be discouraged by actions such as automatically assigning AI ownership to AI itself (legal ownership should remain with the human developer or organization using AI) and overly restrictive licensing rules that limit AI model improvements or prevent collaboration in AI development.

Open-Source Development

We urge OSTP to support open-source AI development as a driver of innovation and establish guidelines for responsible use, including open-source or otherwise publicly available guidelines for how AI systems should be developed, implemented, and monitored. For example, an open-source AI framework will democratize innovation and allow smaller innovators to participate in AI development. OSTP also could encourage voluntary security and bias assessments for widely used open-source AI models to improve reliability. We discourage the implementation of policy actions requiring licenses or government approval for publicly available AI tools.


Regulatory Structures and Approaches for Optimal Deployment

Flexible AI Model Development and Accountability

We recommend policies that promote flexible, industry-driven AI development practices rather than government-imposed technical constraints. An AI model governance process for use by hospitals and healthcare systems and other providers should focus on the principles of transparency, explainability, and appropriate monitoring. The AI Action Plan should promote auditability and transparency in AI development while allowing flexibility in how organizations implement these principles. For example, AI tools that augment clinical decision-making should be transparent to the underlying data and/or sources used to support suggestions or recommendations, allowing the “human in the loop” to exercise judgment in relying on outputs from AI tools. In addition, there should be safeguards to ensure that the models are sustainable and avoid model degradation. We urge that the goal should be to operationalize mechanisms for post-deployment monitoring rather than pre-deployment approval processes for every algorithm. We caution against strict limits on model adaptation, which could prevent AI systems from learning and improving over time; or requiring AI models to be fully interpretable in every case — some advanced models (e.g., deep learning) have inherent complexity that cannot always be easily explained.

Security Against AI Model Attacks

AI technologies, for several reasons including how data is received, processed, and generated, have greater security exposure than non-AI systems. **We recommend industry-driven AI security standards which could expedite responsiveness to rising threats, such as by scaling with risk, imposing stricter controls on AI handling sensitive personal data while allowing more flexibility for non-sensitive applications.** Guidelines should be developed to support robust defenses against adversarial attacks such as model poisoning, evasion, and data leakage threats. The AI Action Plan should empower hospitals and



healthcare systems and other providers to implement automated monitoring and anomaly detection systems to identify AI model manipulation in real-time.

Additionally, we ask OSTP and NITRD NCO to advocate for a framework to support secure AI model training and deployment using encryption, federated learning, and permissive training techniques. The AI Action Plan should include AI cybersecurity and resilience protocols. For example, the AI Action Plan should emphasize de-identifying data where possible and only using PHI where necessary; security-by-design approaches; measures to monitor and prevent model poisoning; and privacy and security measures throughout the technology lifecycle. These protocols would establish national security standards for AI in healthcare to protect against cyber risks that could compromise patient safety.

Support Patient Data Rights and Maximize Benefits of AI in Healthcare

Patient autonomy is essential, and potential AI laws and regulations should strike a balance between the right to opt-out and the benefits of AI in improving healthcare delivery. Overly restrictive opt-out requirements will involve redesigning software architectures and physical spaces, maintaining dual workflows when technically and operationally feasible (AI-assisted and non-AI), and training clinicians to manage both systems concurrently. These challenges are exacerbated by the difficulties of interoperability in healthcare technology systems. This complexity could deter healthcare organizations from adopting AI solutions. We suggest an alternative approach to opt-out provisions by ensuring transparency, such as clearly informing patients about the use of AI in their care. This can be achieved by incorporating general consent for AI use as part of the transparency process. Additionally, it is crucial to ensure robust oversight by maintaining high standards for safety and efficacy through continuous monitoring of model outputs and outcomes, including human-in-the-loop oversight.

This balance will ensure that patients have a clear understanding of their data rights, how their data is being used and protected, while reducing burdens on AI deployment.

Support Personalized and Accessible AI Education

We recognize the vital importance of education in the AI landscape and strongly support its growth and practical application. **We advocate for a multifaceted approach to AI education, moving away from a one-size-fits-all model.** Establishing educational best practices for AI in healthcare would help identify knowledge gaps and create customized learning programs while maintaining privacy standards. We recommend adopting impact-driven best practices for general workforce training, and more stringent oversight for AI-focused medical education.

Data Privacy and Security Throughout the AI Lifecycle

AI developers should integrate privacy-preserving techniques (such as differential privacy and data minimization) throughout the entire AI lifecycle. It is essential to establish clear guidelines for providers regarding AI data governance, including the sourcing, storage, and sharing of data. These guidelines should be based on industry-driven best practices that align with existing regulations such as HIPAA. We caution against excessive restrictions on AI training data, which could limit model development and innovation in areas like predictive analytics in healthcare, in addition to broad limitations on AI data sharing that prevent collaboration.

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. While healthcare providers bear responsibility for the care they provide, the **developers of commercial AI products must be accountable for the safety and reliability of their products and required to be truthful in marketing their products**, especially since safety, bias, privacy and security, or other harms may be caused by a flaw in the tool itself. Commercial AI and machine learning (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, along with collaboration from developers in properly monitoring AI technology performance over time.

The FAH appreciates the opportunity to address these critical issues and looks forward to collaborating with OSTP as the agency continues to consider how best to ensure the use of AI to realize benefits for patients, providers, and society at large, while mitigating risks, including through existing risk management frameworks. If you have any questions or wish to discuss these issues further, please do not hesitate to contact me or a member of my staff at 202-624-1500 .

Sincerely,

