Submitted via email to*: ostp-ai-rfi@nitrd.gov*

March 14, 2025

AI Action Plan
Attn: Faisal D'Souza
National Coordination Office
National Science Foundation
2415 Eisenhower Avenue
Alexandria, VA 22314

RE: Request for Information on the Development of an Artificial Intelligence (AI) Action Plan

Dear Mr. D'Souza:

The Healthcare Trust Institute appreciates the opportunity to submit comments on the Request for Information on the Development of an Artificial Intelligence (AI) Action Plan (Plan) issued by the Networking and Information Technology Research and Development (NITRD) National Coordination Office (NCO), National Science Foundation on behalf of the Office of Science and Technology Policy (OSTP), and published in the Federal Register on February 6, 2025'').[1]

The Healthcare Trust Institute (HTI) is an alliance of healthcare organizations committed to promoting and implementing effective privacy and security protections for health information that engender trust in the healthcare system and allow for the advancement of treatments, cures and improved healthcare quality for individuals and populations. HTI members, which include companies and organizations from across the U.S. healthcare economy, agree that a strong national privacy standard for health information is needed to protect sensitive data and spur medical innovation.

The RFI is in response to Executive Order 14179, which calls for the development of the Plan to define the priority policy actions needed to sustain and enhance America's AI dominance, and to ensure that unnecessarily burdensome requirements do not hamper

[1] See 90 Fed. Reg. at 9088 (February 6, 2025).

private sector AI innovation. OSTP seeks input on the highest priority policy actions that should be in the Plan, including cybersecurity, data privacy and security, innovation, and competition.

While AI has been in use in limited areas of health care and the life sciences for decades, we are now entering a period of rapidly evolving AI developments that are transforming the healthcare landscape. AI tools are today being used in health care across a much broader spectrum of functional areas and is becoming critically important for improving diagnoses and care, efficient health administration, and reducing unnecessary costs by streamlining tasks and mundane processes. We believe the potential for the further beneficial deployment of AI in the healthcare industry is enormous within a policy framework that nurtures its growth while always keeping the safety and wellbeing of the patient front and center.

To achieve this, it is critical to establish a national standard or framework to ensure regulatory harmonization and avoid undue burden, inefficiencies, and potential safety issues, of having to comply with a myriad of overlapping, and potentially inconsistent, state laws. Within the last year alone, hundreds of state bills have been introduced seeking to regulate almost every aspect of the use of AI, and the pace of new state bills on AI is only increasing, with more and more being enacted each year. Compliance with this proliferation of state laws will be extremely challenging and onerous for the vast majority of organizations that do not operate exclusively within one state. For these reasons we support explicit federal preemption of state laws regulating AI, especially for key provisions where conflicting requirements would become impracticable or overly burdensome. However, to the extent that states do provide guidance on issues not addressed in federal guidance, they should be encouraged to take a coordinated approach using common models or frameworks to avoid inconsistencies and variability.

A federal AI framework should be established in coordination with industry, rather than being solely government-driven.[2] It should be based on widely-used and well-respected frameworks that rely on a risk-based approach, such as the AI Risk Management Framework (RMF) issued by the National Institute for Standards and Technology (NIST). This framework was developed with cross-industry and business perspectives in mind, ensuring that its approach responsibly addresses the risks of AI without being overly prescriptive so as to stifle innovation or competition. Quantitative measures, data specifications, and reporting requirements should be based on nationally adopted voluntary consensus-based standards.

Given the different uses, types, and level of risks of AI in different sectors, as well as the very different regulatory contexts, the issuance of more specific requirements, if any, should be delegated to existing sector-specific federal agencies with subject matter expertise. For example, the regulation of high-risk AI in the health sector should be

[2] See PUBLIC LAW 104-113 NATIONAL TECHNOLOGY TRANSFER AND ADVANCEMENT ACT OF 1995 Including Amendment by Public LAW 107-107, section 1115 on Dec 28, 2001, UTILIZATION OF CONSENSUS TECHNICAL STANDARDS BY FEDERAL AGENCIES, and 2016-01606 (81 FR 4673) Revision of OMB Circular No. A-119, “Federal Participation in the Development and Use of Voluntary Consensus Standards and in Conformity Assessment Activities”

delegated to the Department of Health and Human Services (HHS), which should work in partnership with health care organizations to develop a consensus on what constitutes high-risk AI in health care. This is particularly important in the health sector, where various HHS agencies, such as the Food and Drug Administration (FDA), the Centers for Medicare & Medicaid Services (CMS) and the Assistant Secretary for Technology Policy and the Office of the National Coordinator (ASTP/ONC) have already begun to issue regulations and guidance on the use of AI within their areas of jurisdiction. A one-size-fits-all approach to AI regulation, as is being proposed by some states, is destined to fail as it becomes clear that overlaying a totally new framework over existing, often very different regulatory regimes is impractical and, in many cases, infeasible. For example, the FDA has issued multiple guidance documents on AI and machine learning in software as a medical device. Similarly, HHS has issued regulations under Section 1557 of the Affordable Care Act addressing the use of patient care decision support tools in clinical care, and ASTP/ONC has issued transparency and risk management requirements for Predictive Decision Support Interventions that are part of certified health information technology (HIT), such as electronic health records. Any AI regulation should harmonize with these existing requirements as well as with other regulatory requirements applicable to these entities. It is also critical that the agencies within HHS and across other federal departments that may have jurisdiction over activities of the health care sector coordinate their approach to AI regulation so that there is consistency and harmonization of requirements across the sector.

Given the centrality of data to AI training and solutions, a national privacy standard should be a foundational requirement for a national policy on AI. Without the assurance that their personal data will remain protected and used and disclosed appropriately, patients and consumers will lack the trust essential for their acceptance of AI solutions. While more and more states are passing comprehensive privacy laws[3], the patchwork of state laws set different standards and provide different rights, which is confusing and difficult to navigate for patients and consumers. It also inhibits innovation by requiring businesses to comply with a patchwork of different, often inconsistent laws. This is not only costly and burdensome but creates inefficiencies with no counterbalancing privacy benefits. It also risks the imposition of operational and compliance barriers as data may be used for training in one state but not another, or subject to different requirements, restrictions, or conditions from state to state.

We have attached a copy of our privacy principles, which we believe should form the basis for any national privacy law. Some of the key principles include appropriate use limitations consistent with consumers' reasonable expectations, consumer rights, transparency in the form of plain English privacy notices, data minimization, and a risk-based approach to cybersecurity that relies on existing best practices and frameworks, such as NIST's cybersecurity framework. It is also critical that a national privacy law

[3] To date, 20 states have adopted generally applicable privacy laws, and this number is expected to grow as long as Congress does not pass a national privacy law that applies to personal data, including personal health data not already subject to HIPAA. See https://iapp.org/resources/article/us-state-privacy-legislation-tracker/.

have a meaningful enforcement mechanism in the form of tiered penalties based on the degree of wrongdoing, and that it harmonize with existing federal privacy regulatory frameworks, such as HIPAA.

A federal AI framework will also bring greater regulatory certainty, which will create a more fertile environment for the development of AI solutions, as major investment in AI, particularly in different aspects of healthcare, depends on a clear understanding of the "rules of the road." Thus, regulation based on adoption of national standards can promote American competitiveness and protect innovation through investments in AI. Both AI developers and those deploying AI in healthcare settings need the assurance that the regulatory environment will not be hostile to the application of AI, whether through overly prescriptive requirements or undue burdens, such as through ongoing evaluations and assessments.

It is important that the regulation of AI be risk-based, taking into account the potential impact and possible harm of the AI tool in question. Developers should be accountable for data quality, model performance, and design flaws, while deployers should be responsible for how they apply AI tools in practice. This approach allows AI developers and users in a healthcare setting to allocate resources appropriately and proportionate to the potential harm of the specific AI use case. Users should be able to tailor risk and impact assessments to the type of AI tool, its intended use and context, potential harms, and changes in the internal and external environment. This approach will result in healthcare applications with the highest risk having the highest guardrails, such as requiring more frequent review or human intervention, consistent with the RMF. On the other hand, AI applications in scientific research and development (R&D) pose a low risk and should be incentivized, as they have the ability to boost innovation and the discovery of new drugs, treatments, materials, equipment, processes, etc.

Finally, an important component of a federal AI action plan should be patient and consumer education on the value and benefits of AI to consumers. Too often, particularly in the area of health care, patients and consumers are warned about the risks and potential harm associated with AI solutions without also being told how AI may help them lead healthier and more productive lives, bring down health care costs, and lead to more accurate diagnoses and improved care. While the risks of AI should not be minimized or overlooked ,and patient safety should always be the first priority, a communication strategy that discusses the benefits of AI while promoting transparent standards to enhance confidence by end users and those impacted by AI decision-making, will help consumers understand what society has to gain from the beneficial application of AI. This will in turn provide broad-based public support for government actions to promote AI development and innovation. We have attached our "AI Best Practices and Uses Cases in Healthcare," which elaborates on the ways in which AI is currently being implemented in health care settings.

Thank you for your consideration of our comments. Please do not hesitate to contact me at tina@hctrustinst.com or 202-750-1989 if you have any questions.

Sincerely,

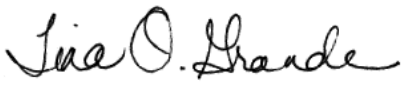

Tina O. Grande
President, Healthcare Trust Institute

Enclosures: HTI Privacy Principles, AI Best Practices and Uses Cases in Healthcare

PRINCIPLES ON HEALTH INFORMATION PRIVACY
BOTH INSIDE AND OUTSIDE OF HIPAA

1. Robust privacy and security protections for personal health information is essential for trust in the healthcare system, which is the foundation for the delivery of quality care and patient safety.

2. All personal health information, whether falling within or outside HIPAA, should be subject to regulation to ensure that it is used in a manner consistent with an individual's reasonable expectations. Uses for other purposes should require an individual's authorization and, where feasible, privacy-enhancing technologies should be implemented.

3. Entities collecting and holding personal health information should be required to have risk-based physical, administrative and technical safeguards in place to protect that information from misuse and threats, including cyberattacks. These safeguards should evolve as technology evolves and be consistent with nationally recognized frameworks, such as the National Institute for Science and Technology (NIST) Cybersecurity Framework.

4. Protections for personal health information should be established at the national level to ensure consistency, clarity and compliance as individuals and data increasingly travel across state lines. It is also essential to avoid data masking to the detriment of patient care and safety, and to ensure that the vision of national interoperability for health data exchange can be realized, leading to better care coordination and improved health outcomes.

5. The principles of minimum necessary and data minimization should be central to collection and processing of personal health information, including through use of de-identified data or privacy-enhancing technologies where feasible. The use of de-identified data is critical to allow for important and beneficial public purposes, such as medical research and public health. To engender consumer and patient trust and public support, recipients of deidentified data should be prohibited from attempting to re-identify the data.

6. Individuals should be provided clear and simple privacy notices that explain how an entity collects and processes personal health information, as well as the individual's rights and choices with respect to their health data. These rights should generally include the right to request access and the right to request corrections.

7. The Health Insurance Portability and Accountability Act (HIPAA) framework, which has been the cornerstone for the protection of patient health information in the health care sector for almost a quarter of a century and is well-understood and trusted by patients and health care organizations alike, should remain the framework for the regulation of patient health information in the health care industry. HIPAA is tailored to health care delivery and payment, and permits the sharing of medical information for treatment, payment and healthcare operations consistent with the reasonable expectations of patients.

Updated: May 2024

8. Regulation of personal health information outside the HIPAA regulations should harmonize with the HIPAA framework, using similar concepts and definitions where appropriate, such as treating data deidentified in accordance with HIPAA as deidentified data for all purposes.

9. Privacy protections must be enforced through meaningful penalties and a mechanism for individuals to be able to report violations without fear of retaliation.

AI Best Practices and Uses Cases in Healthcare

AI is advancing healthcare in multiple ways, from diagnosing diseases and helping to find new cures to streamlining administrative processes and improving workflows. Along with its enormous promise to transform healthcare, AI brings new potential risks. It is essential that these risks be managed through appropriate guardrails while allowing the full potential of AI in healthcare to be realized through ongoing innovation. Below we outline some AI best practices designed to achieve this goal, as well as describe some current and potential use cases for AI in healthcare.

A. AI Best Practices

Build Trust. Trust is the underpinning of the healthcare system, and any deployment of AI should be designed and implemented in a manner that builds trust by adhering to certain key principles such as transparency, accountability, fairness and respect for individual rights and privacy. Individuals should be assured that AI is deployed in an ethical manner with the ultimate goal of benefitting society.

Human Involvement. Humans should be involved throughout the AI life cycle, from the initial design to deployment to ongoing monitoring of results. This is essential not only to maintain trust, but also to ensure that AI tools perform as intended. AI should facilitate, not replace, critical decision-making, and AI solutions should be developed with input from and in collaboration with clinicians and others in the healthcare industry that will use them. It is through combining the respective strengths of AI and humans that the AI can achieve its highest and best use in healthcare.

Robust privacy and security controls. Privacy and security of data should be of the highest priority. The concept of privacy by design and default should be the guiding principle in AI development and deployment. Robust privacy and security protocols should be implemented, including use limitations, physical, administrative, and technical safeguards, data minimization, cybersecurity controls, and ongoing privacy and security training of personnel. Organizations should follow national standards for the collection, processing, and transfer of personal information not already subject to HIPAA, including standards for the use and protection of personal data in AI solutions, such as the NIST Cybersecurity Framework.

High Quality Data. The performance of AI is directly dependent on the quality of the data inputs that are being used to train the AI models. Training data should be as reliable, representative, and complete as possible with respect to the population it is intended to serve across a range of factors such as age, gender, ethnicity, and health status. From the inception of data collection to the use of the data to produce insights, the goal should be to build algorithms that are reliable, accurate, unbiased, and that protect the patient from harm.

Ongoing Monitoring and Evaluation. Performance benchmarks should be established to ensure that the AI model is performing as intended and that decision-making through it is appropriate. Decisions should be reviewed for potential bias and model outputs should be continuously monitored to avoid drift, overtraining, AI hallucination and other AI risks.

Transparency. Disclosure and openness are essential to promote trust and acceptance of AI in health care. AI developers should explain how their products work, including information on both their capabilities and limitations. This is important not only to set realistic expectations, but so that healthcare organizations can make informed decisions as to how to best use these tools. Healthcare organizations, in turn, should disclose to patients and consumers when they are interacting with AI tools.

Oversight and governance. A governance body should be put in place to oversee the development and deployment of AI in the organization. This body should establish policies and standards for the safe and responsible use of AI in the organization, including controls to identify and mitigate potential risks, ensure data integrity, protect privacy, and evaluate performance. It should also be responsible for continuous monitoring and improvement of AI tools as well as compliance with regulatory requirements.

B. AI Use Cases in Healthcare[1]

AI tools are increasingly being used in health care to improve patient outcomes, enhance operational and administrative efficiency, and advance medical research. Below we provide examples of AI use cases in these three key areas.

1. Improve Clinical Care

AI is being used to facilitate clinical decision-making and improve care in many ways, from helping to diagnose diseases to improving treatments and finding new cures. Some use cases include:

[1] These use cases are drawn from a variety of sources, including Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care," a presentation by Manatt and the American Medical Association, May 1, 2024; the Testimony of Peter Shen, Head of Digital & Automation, Siemens Medical Solutions USA, Inc., at United States Senate Committee on Finance Hearing on "Artificial Intelligence and Health Care: Promise and Pitfalls," February 8, 2024; and Kaiser Permanente's May 6, 2024, response to Rep. Ami Bera's RFI on "The State of AI in Healthcare."

- To perform clinical activities more accurately and quickly.
  Example:
    - Patients undergoing a CT scan for lung cancer screening can be better positioned in the CT scanner to help optimize the resulting generated images, while minimizing the time the patient spends in the scanner. This is done by AI that is built into the CT scanner technology that allows the machines to identify human anatomy.

- To identify objects, patterns, and/or characteristics within data (often images).
  Examples:
    - Radiologists reviewing the resulting lung cancer screening CT images can utilize AI-guided computer software as a companion to the clinician to identify abnormalities, including the ability to measure the density and characterize the size of suspicious nodules that were previously not possible to visualize without the assistance of AI.

    - A physician orders an X-ray for a patient who presents with pain, swelling, and limited leg mobility. An AI tool can review the X-ray and identify an incidental nodule for further analysis by a radiologist.

    - Using computer vision programs to help clinicians identify subtle features that may be associated with increased risk of diseases. For example, AI-enabled imaging of the retina can be used to predict the risk of cardiovascular disease and stroke.

- To predict or forecast future events based on historical data and patterns.
  Examples:
    - A patient is discharged after hospitalization for heart failure. Using historic heart failure readmission rates and the patient's clinical data, an AI tool can predict the risk of the patient's hospital readmission.

    - AI tools can measure changes in brain volume over time (a predictor of neurogenerative diseases, such as Alzheimer's) by automatically segmenting different structures of the brain on an MRI image, measuring their volumes, and comparing these to data in a brain database. The AI tools can feed these comparative results into a report where deviations in volume from the norm are highlighted, providing neurologists with actionable, patient-specific volumetric data to diagnose and treat the patient more accurately.

    - Using predictive analytics to identify and address inpatient complications before they occur.

- Using predictive analytics to identify patients that would benefit from a care coordination intervention after hospital discharge which resulted in reduced hospital readmission rates.

- To summarize data inputs into shorter and more accessible outputs.
  Example:
  - A patient is admitted to an emergency room after suffering an epileptic event. A team of admitting healthcare providers review the patient's medical file to understand the patient's medical history, current medications, previous allergic reactions, and potential triggering factors. An AI tool can review the patient's medical history in totality, near- instantly identifying and summarizing key information for current clinical needs, such as recent medication changes affecting seizure threshold and a list of contraindicated drugs based on allergy history.

- To provide recommendations, guidance, or advice.
  Example:
  - A patient sees a provider every few months for a routine check-in; provider team conducts retrospective analysis of blood glucose measures from past few months. An AI tool can continually monitor a patient's blood glucose levels and (1) sends an alert to patient and provider when deviations occur and (2) provides recommended course of action (e.g., insulin level recommendation)

- To deliver safer and more accurate treatment.
  Examples:
  - To minimize the risk that healthy tissue around the cancer is not unnecessarily radiated, radiation physicists create a radiation treatment plan, which includes the tedious task of manually drawing the unique contours of the cancerous tumor. This manual contouring potentially delays the time to treatment for the patient. AI-enabled auto-contouring software can automatically detect these contours of the cancerous area, significantly speeding up the patient's time to treatment and potentially eliminating extraneous treatments.

  - Traditionally, a urologist identifies suspected areas of prostate cancer by manually reviewing written reports and pictograms of the prostate provided by radiology and as needed, acquires tissue samples from the areas in question using ultrasound-guided biopsy. An AI-tool is being developed to automatically segment suspect areas of the prostate and characterize and measure suspicious lesions in the prostate from MRI images. This qualitative and quantitative analysis may support the urologist's decision on whether a tissue biopsy is additionally required for diagnosis or if such invasive procedure can be avoided, which is significant in managing a

prostate cancer patient's well-being and minimizing unnecessary costs within the health system.

- To reduce medical errors.
  Example:
  - AI technology has been trained to analyze pill size, shape, color, and markings, as well as see broken pills, foreign objects, or anything else that should not be in the vial. The system then uses this information to appropriately send prescriptions to a pharmacist for verification. This technology helps ensure that the correct medication is dispensed.

- Population health management.
  Example:
  - AI models are used to help identify health plan members most at risk of suffering a fall, having difficulty controlling blood pressure, struggling to adhere to medication regimens, more likely to require non-obstetric hospitalization in the next 12 months, and those who should likely be included in vaccine outreach campaigns. These populations can then be targeted with additional services or information.

2. Improve Health Care Administration

There are many promising AI use cases to streamline administrative processes, increase efficiency and improve productivity. Operational and workflow improvements reduce wait times, lower administrative costs, and improve the patient and healthcare provider experience.

Some uses cases include:

- Automate actions. Reduce staff time spent on administrative work by having the system, for example, answer patient questions, schedule follow-ups, provide the status of a claim or order or assist in finding an in-network provider. Using an AI-enabled tool to assist with these tasks frees up call center personnel to focus on more complex questions and issues, reducing wait times and lowering overhead costs.

- Simplify documentation. Reduce time spent at the keyboard by having the system, for example, use machine learning and/or ambient voice technology to facilitate scribe-like capabilities in real time to improve physician-patient interaction and improve administrative efficiency. This not only saves time but allows clinicians to focus on the patient during a visit, rather than having to constantly turn to a computer screen to input notes.

- Tailor communications. Communicate better by having the system, for

example, simplify clinical notes to patient-friendly language and instructions. AI tools are particularly adept at translating clinical jargon to plain English descriptions and generating an initial draft, which will then be reviewed and finalized by the clinician. For example, AI tools may use natural language processing algorithms to optimize patient-physician communications or draft patient discharge instructions for a provider to review, reducing provider burden and improving productivity.

- Summarize the chart. Reduce time spent searching the chart by having the system summarize recent notes before a visit or highlight key details in imaging studies or converting a radiologist's audio dictation into a structured summary and applies the BIRADS classification scheme automatically.[2]

- Finding patterns. Improving fraud and abuse detection by reviewing medical and pharmacy claims to detect unusual patterns or red flags that warrant closer scrutiny. This allows auditors to focus their resources where most needed, resulting in lower health care costs as fraudulent and abusive practices are more efficiently and effectively identified and addressed.

3. Research

AI is also now also a key component in developing new drugs and cures for some of the most challenging diseases, including through precision medicine and individually tailored drug therapies. Some use cases include:

- Improving the effectiveness and reducing the cost of clinical trials.
  Examples:
  - Using AI tools to identify potential participants and streamlining the monitoring and coaching of patients.

  - Using AI algorithms to examine in detail the available scientific literature and support the identification of genetic biomarkers associated with certain diseases, enabling more effective clinical trials and shorter periods to put treatments on the market.

- Improving cost-efficiency of drug development.
  Examples:
  - Creating virtual control groups to decrease or remove the need for "real" control groups in certain clinical trials. This results in selecting fewer patients for placebo or standard treatment, thereby increasing the cost-

[2] See "Future of Health: The Emerging Landscape of Augmented Intelligence in Health Care," presentation by Manatt and the American Medical Association, May 1, 2024 ("Manatt Presentation").

efficiency of drug development.

- Analyzing vast datasets like genomic data connected to a disease, detecting potential drug targets, and predicting a drug's efficacy and its potential side effects.

- Helping researchers to analyze and repurpose existing medicines to combat specific diseases, making the development of new drugs more cost-efficient and effective.