CRA and CCC Response to Request for Information on Accelerating the Adoption and Use of Artificial Intelligence as Part of Clinical Care

February 17, 2026

Authors: Kevin Butler (University of Florida), Rayid Ghani (Carnegie Mellon University), Rachel Greenstadt (New York University), Ming C. Lin (University of Maryland), Manish Parashar (University of Utah), Blaine Reeder (University of Missouri), Tammy Toscos (Parkview Health), Cathy Wu (University of Delaware), and Haley Griffin (CRA/CCC).

This response is prepared by the Computing Research Association (CRA), assisted by the Computing Community Consortium (CCC) and CRA-Industry (CRA-I). CRA is an association of over 270 North American computing research organizations, both academic and industrial, and partners from six professional computing societies.

The mission of CCC, a CRA subcommittee, is to enable the pursuit of innovative, high-impact computing research that aligns with pressing national and global challenges.

The mission of CRA-I, a CRA subcommittee, is to convene industry partners on computing research topics of mutual interest and connect them with CRA's academic and government constituents for mutual benefit and improved societal outcomes.

Please note any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the authors' affiliations, or of the National Science Foundation, which funds the CCC.

1. What are the biggest barriers to private sector innovation in AI for health care and its adoption and use in clinical care?

There is a notable public reluctance in the U.S. to adopt AI for health-related outcomes (Patients' Trust in Health Systems to Use Artificial Intelligence), stemming from the high stakes and inherent privacy associated with clinical wellness. It is crucial that R&D resources are devoted to ensuring that the American people feel like their data, diagnosis, and care are trustworthy and secure in the age of AI.

Below is a list of some of the barriers to private sector innovation in this area:

- Inadequate access to large-scale, relevant, and useful data.
- Insufficient input and buy-in from clinical practitioners and organizational leadership.

- Disproportionate focus on technical implementation (AI models, metrics, etc.) and efficiency over collaborative, clinical outcome-driven design and evaluation.
- Underfunded research infrastructure and a lack of access to compliant computer/storage environments and systemic support.

3. For non-medical devices, we understand that use of AI in clinical care may raise novel legal and implementation issues that challenge existing governance and accountability structures (e.g., relating to liability, indemnification, privacy, and security). What novel legal and implementation issues exist and what role, if any, should HHS play to help address them?

There is significant ambiguity in the existing definitions of medical devices (e.g., FDA) in the age of AI, and what organizations are covered under HIPAA. This is especially concerning due to the private patient medical data that is oftentimes being shared. While the results that AI models produce are likely to improve if they are trained on more data, there are not currently strong assurances that patient data can remain private and secure if it is provided to a model, either in training, fine-tuning, or with a prompt at inference time. Here are a few implementation applications and related research questions:

- Recording devices collecting patient-clinical interactions, and AI generating post-visit summaries (how is the data being stored? What is the data being used for in the future? How can informed consent be ensured?)
- Software non-medical devices (what organizations are covered under HIPAA? When does software become a medical device?)
- Providing AI models with non-clinical data (how can we prevent models from sharing private data across different patients?)
- Reliability and integrity of AI-generated recommendations (How can providers evaluate and monitor systems that may hallucinate, produce over-confident outputs, or degrade over time? What validation and auditing standards are needed for systems used in clinical workflows?)
- The use of physical non-medical devices (How do robots fit into the clinical care environment? How can privacy and security be ensured?)

4. For non-medical devices, what are the most promising AI evaluation methods (pre- and post-deployment), metrics, robustness testing, and other workflow and human-centered evaluation methods for clinical care? Should HHS further support these processes? If so, which mechanisms would be most impactful (e.g., contracts, grants, cooperative agreements, and/or prize competitions)?

HHS should further support AI evaluation methods for non-medical devices. In particular, there is a need for usability and acceptability studies at all stages from conceptualization, design, iterative testing, lab and field tests, translation to real-world settings, and implementation. There is also a need for participatory design from the very beginning of a project that includes robust

discussions with direct users and stakeholders (see Generative AI-Derived Information About Opioid Use Disorder Treatment During Pregnancy: An Exploratory Evaluation of GPT-4's Steerability for Provision of Trustworthy Person-Centered Information for an example of a study that is examining the trustworthiness of generative AI at a very early research stage). Through these conversations, researchers can determine the on-the-ground needs for proposed applications, and decide on goals based on workflows and implementations that will function in a clinical setting. Standard and operationalized safety evaluation frameworks should also be implemented to provide guardrails that build trust in generative AI applications (see Health Consumers' Use and Perceptions of Health Information from Generative Artificial Intelligence Chatbots: A Scoping Review for more information about distrust in generative AI applications in healthcare).

5. How can HHS best support private sector activities (e.g., accreditation, certification, industry-driven testing, and credentialing) to promote innovative and effective AI use in clinical care?

We encourage HSS to support industry-academia partnerships across AI and clinical care, as many university researchers have expertise in training and credentialing that the private sector could greatly benefit from.

We also recommend regulatory reforms that facilitate broader access to anonymized health data from EHR providers, while simultaneously incentivizing health systems to collaborate with researchers for better data contextualization (Enabling the AI Revolution in Healthcare).

Grantees of the "America's Seed Fund", particularly Small Business Innovation Research (SBIR) and Small Business Technology Transfer (STTR), should receive guidance on accreditation, certification, industry-driven testing, and credentialing.

How to sustain models in health systems is a very critical area for support and innovation. Peter Emby at Vanderbilt coined an AI monitoring term *"algorithmovigilance"* during covid times [see 2021 paper] and now his lab has developed a platform they call **VAMOS** that *"aims to support systematic and proactive monitoring of AI tools in healthcare organizations. Our findings and recommendations can support the design of AI monitoring systems to support health systems, improve quality of care, and ensure patient safety."* [see 2025 paper].

6. Where have AI tools deployed in clinical care met or exceeded performance and cost expectations and where have they fallen short? What kinds of novel AI tools would have the greatest potential to improve health care outcomes, give new insights on quality, and help reduce costs?

Positive Deployment Use Cases

AI tools have a lot of potential to enhance the doctor-patient interaction when deployed in a clinical setting. Below are several examples from a recent study (To Do No Harm — and the Most Good — with AI in Health Care):

- capture of the recorded patient visit,
- prioritizing and analysis of test and imaging results,
- differential diagnosis,
- plan for therapy and discussion of alternatives,
- instructions for the patient and caregivers,
- appointment scheduling and other administrative functions,
- responses to patient questions at optimal levels of literacy.

Another relevant study (Ambient Artificial Intelligence Scribes and Physician Financial Productivity) indicates that, "adoption of ambient artificial intelligence (AI) scribes, which generate clinical documentation from audio recordings and are associated with reduced documentation time and burnout, is increasing," and concludes that there is a promising return on investment when physicians adopt the tools.

Negative Deployment Use Cases

Implementing and using unvalidated models in clinical care is incredibly dangerous for patients. For example, the widely adopted Epic Sepsis Model that is supposed to predict cases of sepsis has been shown to not be very accurate (External Validation of a Widely Implemented Proprietary Sepsis Prediction Model in Hospitalized Patients). Considering that sepsis occurs in 7% of hospitalizations, relying on this model is a potentially deadly misstep.

AI tools are only worthwhile if they are used by the individuals responsible for patient care. AI-generated trigger alerts in the electronic health record, such as sepsis, are often ignored. However, if systems are autonomous, and there is no human in the loop, AI tools that have not been validated with local data and maintained over time could result in significant harm to patients. A recent survey showed that only 44% of hospitals reported using local data for model validation (Current Use And Evaluation Of Artificial Intelligence And Predictive Models In US Hospitals). This study also showed that models built in house, vs. simply adopted by a vendor, were more likely to use local data for validation. However, such in house built systems typically happen at health systems that are better resourced. What is needed is more incentives and resources allocated for successful hyperlocal implementation and maintenance of predictive models.

7. Which role(s), decision maker(s), or governing bodies within health care organizations have the most influence on the adoption of AI for clinical care? What are the primary administrative hurdles to the adoption of AI in clinical care?

Hospital administrators have significant influence on the adoption of AI for clinical care, and their main concern is the hospital budget. They are incentivized to support AI if there is a return on investment in both profit and improved patient outcomes.

A significant challenge to the implementation of AI technology is that the core users (practicing physicians, nurses, etc.) are oftentimes not sufficiently consulted on what would best support their work with patients.

Julia Adler-Milstein PhD, Chief of Clinical Informatics & Digital Transformation at the University of California San Francisco (UCSF), has worked on ways to support executives and providers that are making decisions about AI use in healthcare. This is important work due to the pressure on health systems to adopt AI for cost savings. Julia's team has worked with UCSF Health to build the "*Impact Monitoring Platform for AI in Clinical Care* (IMPACC)," which seeks to "bridge the gap between the rapid evolution of AI technologies used by clinicians and the essential need for robust, ongoing assessment of their efficacy, safety, and equity" (First Continuous AI-monitoring Platform in Clinical Care). This is an important model for health systems to consider adopting for safe and effective use of AI in healthcare.

8. Where would enhanced interoperability widen market opportunities, fuel research, and accelerate the development of AI for clinical care? Please consider specific data types, data standards, and benchmarking tools.

There is a significant need for better/compliant infrastructure (example: NAIRR-Secure) and accessibility to data sources at the intersection of AI and clinical care. We recommend incentivizing the interoperability of electronic health records (EHR) to allow for data exchange, building on the widely used Fast Healthcare Interoperability Resources (FHIR) system. This would result in more open-source data, which is needed for benchmarking tools. Many existing regulations are in place to encourage EHR vendors to share their data through these systems, but there is a need for additional strategies to enforce compliance.

As new devices are developed, a consistent data format needs to be implemented and enforced so that medical and non-medical devices that are applied to the clinical care setting can communicate with each other. However, existing devices are not going to be discontinued overnight, so new tools need to be developed to help convert standards and promote interoperability. A mapping exercise of devices and tools would be very helpful.

Interoperability typically focuses on clinical data, but social determinants of health — which account for 80 percent of health outcomes — must also be considered (Enabling the AI Revolution in Healthcare). Support for privacy-preserving cryptography—especially to enable secure and private machine learning and unlearning—is essential for advancing safe public health data analysis (Enabling the AI Revolution in Healthcare).

9. What challenges within health care do patients and caregivers wish to see addressed by the adoption and use of AI in clinical care? Equally, what concerns do patients and caregivers have related to the adoption and use of AI in clinical care?

At the end of the day, patients and their caregivers want to know what is going to happen to them, and what they can do about it. Researchers should work towards developing AI tools that can answer these two important questions.

Building on our response to question 1, the majority of Americans do not trust the use of AI in their healthcare. They are concerned about the security of their data, AI hallucinations, model poisoning, etc. It is incredibly important that AI and clinical care researchers partner with communications and public health professionals to share the steps being taken to ensure that the AI helping to inform health decisions is safe, accurate, and trustworthy. Without these steps in place, individuals could lose trust in AI and, by extension, in the healthcare infrastructure as a whole, leading to lower engagement and, consequently, worse health outcomes.

Hospitals and care facilities in rural communities often do not have as many resources to support patient care as those in urban areas, and in the AI era the gap between areas with access to AI infrastructure and without is becoming increasingly difficult to overcome. Additional resources should be devoted to supporting clinical care in states included in the Established Program to Stimulate Competitive Research (EPSCoR) and Institutional Development Award (IDeA) programs.

As AI is being developed for the clinical care setting, it should support interactions between clinicians and patients/caregivers, as well as between healthcare portals and patients/caregivers. Today, a lot of interactions are not person-person. Test results, diagnoses, and treatment recommendations are oftentimes communicated through a portal, which can be frustrating and cryptic for patients/caregivers. The portals should be personalized based on the needs, language and cultural background

10. Are there specific areas of AI research that HHS should prioritize to accelerate the adoption of AI as part of clinical care?

HHS should prioritize (1) AI-Driven Clinical Decision Support, (2) Human-AI Collaboration, and (3) Development and Implementation Trajectory Evaluation Frameworks.

<u>AI-Driven Clinical Decision Support</u>
New modeling methods need to be developed that are end to end, throughout the entire workflow of clinical care, and include reasoning and inference. Benchmarks and trials need to be developed.

Human-AI Collaboration
Clinicians need to be a key stakeholder in any AI application (especially embodied AI) to the clinical care sector, as they are the ones using it to impact patient health. Many other humans also need to be in the loop, including radiologists, patients, and caregivers.

Development and Implementation Trajectory Evaluation Frameworks
There is a need for in-context evaluation of the methods and systems impacted by AI use. The process should evaluate not only trustworthiness, but correctness. If the goal is for AI to generate new assistance/knowledge for diagnostic advice, the information produced *has* to be accurate, relevant, and understandable.

a. Are there published findings about the impact of adopted AI tools and their use in clinical care?

Yes, below is a list of relevant published findings we recommend:

- State of Clinical AI Report 2026 - ARISE Network (An annual synthesis of the most significant developments, evidence, and emerging challenges in clinical AI): Peter G. Brodeur, Ethan Goh, Emily Tat, Liam McCoy, David Wu, Priyank Jain, Rebecca Handler, Jason Hom, Laura Zwaan, Vishnu Ravi, Brian Han, Kevin Schulman, Kathleen Lacar, Kameron Black, Adrian Haimovich, Eric Horvitz, Adam Rodman, Jonathan H. Chen "State of Clinical AI 2026," ARISE Network, January 2026.
- AI-CARING NSF AI Institute has developed a number of AI tools for use in the clinical context particularly as they relate to mild cognitive impairment and older adults, and published a wide range of articles on the topic.
- A recent survey (2025) offers a summary of Current Use And Evaluation Of Artificial Intelligence And Predictive Models In US Hospitals

b. How does the literature approach the costs, benefits, and transfers of using AI as part of clinical care?

In a recent secondary analysis of the 2023 American Hospital Association Annual Survey Information Technology, Nong et al. identified how AI and predictive models are currently being implemented in healthcare organizations. The authors found that organizations "might not robustly evaluate AI because of cost, underestimation of potential negative impacts, barriers to governance including technical capacity, a lack of financial incentives, or the perception that regulation does not require it. If hospitals are not performing robust evaluation, further policy intervention will likely be necessary" (Current Use And Evaluation Of Artificial Intelligence And Predictive Models In US Hospitals).

In a discussion paper by Adler-Milstein et al, the authors discuss barriers and facilitators of clinical adoption of AI for medical diagnosis. They point out that "AI diagnostic decision support (AI-DDS) tools could reduce the cognitive burden on providers, mitigate burnout, and further enhance care quality." The authors also point out the criticality of understanding human adoption. Even the best model will not be used for good if implementation details foundational

for successful adoption are not considered - key factors outlined in the paper are 1) reason to use, 2) means to use, 3) method to use, and 4) desire to use (Meeting the Moment: Addressing Barriers and Facilitating Clinical Adoption of Artificial Intelligence in Medical Diagnosis).