American College of Radiology™

May 6, 2024

The Honorable Ami Bera, M.D.
Cofounder, Health Care Innovation Caucus
Member, Bipartisan House Task Force on AI
172 Cannon House Office Building
Washington, DC 20515

Re: State of AI in Health Care RFI; Comments of the American College of Radiology

Dear Dr. Bera:

The American College of Radiology (ACR)—a professional association representing over 40,000 diagnostic radiologists, interventional radiologists, nuclear medicine physicians, radiation oncologists, and medical physicists—appreciates the opportunity to comment on the March 20, 2024, request for information addressing the "State of Artificial Intelligence (AI) in Health Care." The ACR applauds the ongoing efforts of Congressman Bera and the Bipartisan House Task Force on AI to explore AI policy opportunities in health care.

Radiology has been the vanguard for AI innovation and regulation within the health care sector. Established in 2017 to address radiology AI topics, the ACR Data Science Institute (DSI)[1] promotes safe, effective, and trustworthy AI use through initiatives, such as:

- Communicating AI policy needs and recommendations to Congress and agencies.
- Defining radiology AI use cases and making these openly accessible online.
- Maintaining a centralized informational database of Food and Drug Administration (FDA)-authorized radiology AI products (AI Central.org[2]) to substantially improve AI transparency, empowering radiologists to find solutions befitting their practices, patient populations, and needs.
- Educating end-users on real-world practice matters of AI implementation, governance, and use considerations.
- Creating mechanisms to support clinical acceptance testing and real-world performance monitoring of radiology AI throughout its lifecycle.
- Developing standards, guidance, and accreditation around health care AI uses.
- Working with the international medical community to develop AI ethics and implementation guidelines.

Additionally, the ACR Center for Research and Innovation (CRI)[3] serves as a resource for ACR members, academia, and industry on clinical trial design and management. The CRI uses new knowledge derived from AI and other data-driven methods to inform research, practice, and reimbursement. While the ACR does not develop AI products for commercial use, ACR data scientists within the DSI and CRI have developed AI models for member education, demonstration projects, and to inform research, practice, and reimbursement.

[1] https://www.acrdsi.org/
[2] https://aicentral.acrdsi.org/
[3] https://www.acr.org/Research/Clinical-Research

Implementation

Q1. *How extensively is AI currently being implemented in health care institutions and other settings across the country?*

A1. Health care AI implementation is increasing across all medical specialties. There are approximately 250 FDA-authorized radiology-specific AI products intended for use in a variety of applications that use diagnostic imaging data, including examination triage, abnormality detection, diagnosis, and advanced quantification of disease. There has been a steady increase in deployment of the AI tools for radiology; however, adoption rates across the disparate clinical environments that provide radiology patient care may vary depending on available resources. While some larger institutions have extensive technological resources and experience with implementing commercial or self-developed AI tools, smaller hospitals may not have dedicated resources. Licensing costs, support staffing, and the relative lack of reimbursement are barriers to acquisition, appropriate governance, and high-quality medical uses of AI.

Q2. *What areas of health care are benefiting the most from AI integration, and what are the primary challenges hindering further adoption?*

A2. Ultimately all medical specialties benefit from integrating safe, effective, and transparent AI into health care systems. Currently, the primary implementation base has benefited from AI-promotion of operational efficiencies. Appropriate medical use of AI could potentially offer returns on investment through future reimbursement mechanisms and/or streamlined care.

There are several challenges that may hinder adoption of AI in health care:

- Lack of reimbursement or other funding source such as endowments or research grants.
- Technical challenges for end-users to confirm, on local data, that the AI will work as expected in their practice.
- Technical challenges in many health care facilities to ensure the AI software performance remains stable over time after initial acceptance testing and deployment.

AI tools for imaging may fail to work as expected if they are used with data and in clinical environments that substantially differ from those used for model training. These differences could relate to disparate equipment, imaging protocols, and patient populations. Additionally, AI model *performance* in a specific local environment can degrade over time as the input data into the AI model changes when the imaging practice changes devices, software, and protocols (technical parameters for how to scan and reconstruct actual images). To be considered safe and effective, AI tools and clinical data outputs must be trusted by the physician end-users and the patients they serve.

Q3. *What are the various applications of AI in clinical or operational contexts?*

A3. There are multiple interpretative and non-interpretive applications for AI in radiological subspecialties, including, but not limited to:

Interpretive

- Identification of individual imaging findings which can represent time sensitive, critical conditions and resulting triage of these examinations for immediate interpretation.

Examples include intracranial hemorrhage, pulmonary embolus, fractures, pneumothorax, and others.
- Detection and/or characterization of imaging findings, examples include pulmonary nodule detection and breast cancer detection.
- Quantification applications, including evaluation of liver fat and iron, pulmonary emphysema, brain cortical thickness, quantification of the volume of diseased brain (plaques) in multiple sclerosis.
- Automated measurements and volumetric quantification of lesions in cancer management.
- "Opportunistic screening" which is the extraction of imaging biomarkers from imaging studies, performed for an unrelated reason, for the purpose of risk stratification of patients. Examples include automated detection and quantification of atherosclerotic disease, bone density, muscle mass, body mass index, and quantification of abdominal and subcutaneous adipose tissue.

Non-Interpretive

- Imaging protocol management, including radiation dose reduction techniques.
- Reference tools, including guidelines-based support and clinical trials matching.
- Workflow optimization, including tools for communication, scheduling, care coordination, and improving patient satisfaction/assistance.
- AI-enabled reporting solutions.

Q4. *How does AI distinguish itself from other health care technologies? How does AI support existing health care technologies?*

A4. Software applications in health care are fundamental to effective patient management, and AI can enable or enhance many of these applications. A subset of AI applications can even perform tasks physician-experts cannot, such as detailed quantitative functions, the potential to detect or characterize certain diseases earlier than humans, and the ability to better predict patient outcomes with certain therapy options. To realize its full potential, AI must integrate with the vast array of existing health care software and databases (including, but not limited to, electronic health records) to bring pertinent data to the physician-expert in a manner that augments their expertise and personalizes patient care.

Q5. *What measures can be employed to guarantee proper reimbursement and coverage for AI technologies in health care?*

A5. AI reimbursement is a novel and complex policy challenge. Establishment of Current Procedural Terminology (CPT) or other Healthcare Common Procedure Coding System (HCPCS) codes for every available AI tool or use case is problematic, and likely infeasible, due to the number of codes that would be needed through such a mechanism. Coding systems do not have pathways for creating procedure codes for services that are already performed by physicians—consequently, AI that perform tasks physicians are often unable to perform during routine workflows, such as advanced quantitative functions, have fared better with the establishment of Category III CPT codes and a few Category I codes. Additionally, the Centers for Medicare and Medicaid Services (CMS) has provided some reimbursement for Category I codes and reimbursement to hospitals through the Hospital Outpatient Prospective Payment System (HOPPS). However, it remains unclear that all reimbursed uses of AI are currently adding value to patients or the health system.

Appropriate and predictable reimbursement for AI is necessary to prevent a two-tier system of care where some patients will have access to AI as part of their care and others will not. To avoid wasteful spending, reimbursement must be based on whether the AI tool adds value to patients and the health system; and whether it is safe, effective, and trustworthy for widespread clinical use. The Centers for Medicare and Medicaid Services (CMS) currently has physician-led mechanisms for determining whether a new service or medical device provides patient care value, and physician leadership must continue in any future AI reimbursement paradigm.

Parameters CMS should consider to ensure safe and effective adoption of AI tools that are valuable to patient care might include:

- Demonstration of the AI's value to individual patients or the health system, as per physician-led evaluations.
- FDA clearance or approval of the AI tool. The sponsor of the FDA-authorized AI tool must also provide transparent information about its creation.
- The provider has adequate policies and procedures in place to ensure safe, effective, and ethical use of AI. This should include governance ensuring the tool is used only by end-users with appropriate subspecialty qualifications, and that end-users are empowered to recognize and mitigate risk by competently overruling the AI result when appropriate.
- The facility is proactively protecting sensitive patient data (privacy/cybersecurity) that serves as the input or output of the AI.
- The provider performs AI acceptance testing before the tool is deployed in patient care.
- The provider continuously monitors real-world AI performance longitudinally, identifies and mitigates performance drift, and provides feedback to the manufacturer and/or appropriate oversight body if the lessons learned are applicable to other stakeholders.

The above criteria for the safety and effectiveness of healthcare AI could be enabled by third parties contracted by the provider facility, such as via practice accreditation of high-quality AI use and registry-enabled testing/monitoring. The ACR has decades of experience in partnering with government agencies to accredit radiology practices—including via a CMS program for advanced diagnostic imaging[4] and a FDA program for mammography[5]—and can leverage this expertise to facilitate AI monitoring, implementation, governance, and use.

Efficacy, Accuracy, and Transparency

Q6. *What clinical evidence exists regarding the efficacy and accuracy of AI-driven health care solutions?*

A6. The FDA's premarket process requires developers to submit standalone performance data and/or multi-reader study data to provide reasonable assurance that their AI devices perform as intended and meet the special controls for the intended FDA clearance category. However, data submitted to the FDA may not be extensive or diverse enough to comprehensively validate performance in all future real-world scenarios across a variety of imaging equipment manufacturers, imaging protocols, end-user populations, and patient populations. As such, AI models used in clinical environments that differ from those where the models were trained may

[4] https://www.cms.gov/medicare/health-safety-standards/quality-safety-oversight-general-information/accreditation-advanced-diagnostic-imaging-suppliers
[5] https://www.fda.gov/radiation-emitting-products/mammography-quality-standards-act-and-program

not perform as specified in the FDA documentation. This reality needs to be conveyed by the producer and understood by the users of the technology and acceptance testing in the local environment is necessary. Furthermore, the full extent of the parameters in the training data, including equipment manufacturer and protocols used and patient population data, may not be transparently accessible to provider. For example, a particular AI model that was not trained on pediatric imaging studies might not be suitable for pediatric use. Other regulatory challenges associated with AI use in pediatric populations were highlighted in the Congressional AI Caucus-hosted briefing on December 5, 2023, featuring ACR member-speakers Richard A. Barth, MD, FACR, FAAP (Chair, ACR Commission on Pediatric Radiology) and Marla Sammer, MD, MHA, FAAP (Chair, ACR Pediatric AI Committee of the Commission on Informatics).[6]

Additionally, it is important to understand the applicable expertise of any end-users who participated in an AI tool's premarket performance study. For example, a radiology AI algorithm intended for highlighting abnormalities in diagnostic images shown in premarket testing to be effective when used by board-certified radiologists and/or other imaging physician specialists, may result in unreliable or poor performance if used by non-experts unqualified to review diagnostic imaging data. Because AI models will have some false positive and false negative outputs, end-users must have adequate training to recognize when a model fails on their own, especially when the model fails to identify an important finding. There was a recent example of *examination triage* algorithms intended for use by imaging specialists to prioritize their worklists being inappropriately used "off-label" by non-imaging clinicians as a "*diagnostic* algorithm," resulting in the FDA issuing a Letter to Health Care Providers in 2022 to warn against such practices.[7]

The American College of Radiology has worked with developers to update the ACR's database of FDA cleared AI products (AI Central.org[8]) to provide further transparency from developers regarding the parameters, including equipment manufacturers, protocols, and patient demographics used in training their AI models. The goal of the ACR Transparent-AI[9] initiative is to supply providers and end-users with more information about commercially available AI tools so they can accurately assess whether a particular product is likely to work in their facility and with their patients prior to purchase.

Q7. *What best practices are recommended to ensure sufficient availability and use of health data for AI-driven health care solutions?*

A7. The ACR has been an active participant in several projects to develop centralized data repositories including the National Cancer Institute's (NCI) Surveillance, Epidemiology, and End Results (SEER) Program.[10] While centralized data repositories are considered to be valuable in providing data for training and testing AI models, health care facilities are reluctant to moving patient data offsite out of privacy concerns. Therefore, the ACR's National Clinical Imaging Research Registry (ANCIRR) is establishing a virtual repository via a virtual distributed imaging data registry where patients' imaging data remain at the health care facility but are available for use in federated training or distributed validation of AI models. We believe this effort will provide

[6] https://www.acr.org/Advocacy-and-Economics/Advocacy-News/Advocacy-News-Issues/In-the-Dec-9-2023-Issue/ACR-Hosts-Hill-Briefing-With-Congressional-AI-Caucus
[7] https://www.fda.gov/medical-devices/letters-health-care-providers/intended-use-imaging-software-intracranial-large-vessel-occlusion-letter-health-care-providers
[8] https://aicentral.acrdsi.org/
[9] https://www.acrdsi.org/DSI-Services/AI-Central/Transparent-AI
[10] https://seer.cancer.gov/

necessary real-world data for AI training and validation, and significantly augment the available data resources needed to accelerate the development of new AI tools for clinical practice.

Q8. *What guardrails or accountability mechanisms could be set to ensure end-to-end transparency?*

A8. In addition to requiring product development and performance transparency, any future reimbursement for AI should be tied to a series of criteria that describe safe and effective clinical AI use. Physicians will ultimately be responsible for AI used in the care of their patients, and it is necessary that their software tools are free of unintended bias and can produce results that are verifiable by physician-experts. Therefore, AI outputs should be explainable to physicians to the extent feasible.

Additionally, health care AI policies should recognize and proactively prevent end-user "automation bias"—i.e., over-reliance on an AI model instead of clinical judgement. Physician-expert end-users are a key risk control to help identify and mitigate unexpected AI performance problems. End-users must be sufficiently qualified in the relevant field of medicine to judge whether a clinical algorithm is performing accurately, and they must be empowered by the health care facility and AI developer to overrule AI inferences and recommendations. Other examples of potential guardrails to ensure safe, effective, and trustworthy AI use—and ACR's readiness to assist government agencies with implementation of these guardrails—are discussed in A5 above.

Q9. *How can we ensure guardrails are put in place to mitigate risks such as disparate impact from racial, ethnic, and other biases?*

A9. If feasible, developers should provide demographics information of the patients used in their training data to end-users (see the discussion of the ACR Transparent-AI initiative in A6 above). Health care facilities should perform acceptance testing of clinical AI tools on local data to ensure effectiveness on their population.

Q10. *What are accountability mechanisms that can be put in place to ensure that there is an accurate spread of information?*

A10. FDA requirements and CMS payment policies must reflect the need for improved AI transparency, as provider and end-user access to information about product training/testing is key to unbiased, effective performance with the tool. In contrast to hardware medical devices, AI training datasets essentially determine product effectiveness, including the risk of bias. Moreover, FDA review of premarket performance data may not guarantee real-world performance in disparate clinical scenarios and across disparate patient populations. Congress should expand or clarify FDA's authority to enhance the public accessibility to product training and testing information to empower AI end-users to avoid or mitigate unintended bias. Such policies could leverage existing public-private AI initiatives, such as the ACR Transparent-AI initiative discussed in A6 above).

Q11. *Are there specific examples of AI applications that have significantly improved patient outcomes or streamlined health care processes?*

A11. Some studies have shown certain CAD-triage type AI applications, used appropriately by radiologists, have led to decreased turnaround time for interpreting studies with critical results or to decreased emergency department and hospital lengths of stay. Some AI applications may

also reduce radiologist interpretive errors by highlighting subtle findings for the radiologist to investigate further, thereby reducing diagnostic errors by serving as a secondary quality assurance review.

Ethical and Regulatory Considerations

Q12. *With the increasing reliance on AI in health care decision-making, what ethical and regulatory considerations need to be addressed to ensure patient safety, privacy, and equity?*

A12. Ethical use of AI demands that patient data is protected. Practices that allow patient data to remain safely behind the institutions' firewalls can often protect data more readily than those that transmit data offsite to developers or central repositories. While offsite data collection can be done with appropriate guardrails, processes such as federated learning and distributed validation where patient data remains onsite, can be equally beneficial while enhancing privacy/security (see A7 for information about relevant ACR collaborations with NCI and others).

Additionally, there are clear disparities in AI use—particularly for patients in under-resourced areas—that could lead to a two-tier system of care based on AI availability. A payment policy mechanism that supports appropriate AI access across various settings is necessary to ensure equitable AI access.

Q13. *How can the use of AI in health care provide benefits while safeguarding patient privacy in clinical settings?*

A13. Please see A12 and A7.

Q14. *What regulations, policies, frameworks, and standards should entities utilizing AI adhere to, and what mechanisms are in place or should be in place to supervise and enforce them?*

A14. For developers:

- Use diverse training and validation datasets to ensure algorithms are broadly generalizable in widespread clinical use in all intended and anticipated use scenarios.
- More narrowly define the intended user population to enable regulators and providers to better understand how the tool is safely and effectively used.
- Implement processes to identify occurrences of, or risks for, unintended bias and unexpected performance.
- Implement ethical practices to prevent exposure of sensitive patient health information.
- Transparently provide information about AI tools in an accessible manner and ensure AI outputs are explainable to physician-expert end-users who are empowered to overrule the AI.

For health care facilities and AI end-users:

- Implement AI governance to ensure AI is being used to solve relevant clinical problems. This must ensure appropriate oversight of the selection, acceptance testing, deployment, and longitudinal monitoring of AI models throughout the facility.
- Ensure tools are used by appropriately qualified end-users.
- Mitigate automation bias.
- Conduct acceptance testing to ensure the AI will perform as expected prior to clinical deployment.
- Implement longitudinal monitoring of AI performance.

As predicated by other policy initiatives to improve health care quality, including accreditation requirements mandated by the Mammography Quality Standards Act (MQSA) and the Medicare Improvements for Patients and Providers Act (MIPPA), incentives through payment policy are proven to promote adherence to best practices. We believe a similar practice accreditation-based approach could be used to promote safe and effective use of clinical AI tools. As mentioned in A5 above, the ACR has extensive experience as a CMS-designated accreditation organization of advanced diagnostic imaging suppliers and can serve in a similar accreditor role to enable value-based AI payment to accredited facilities that safely/effectively use AI.

Other Considerations

Q15. *What emerging trends do you foresee in the intersection of AI and health care?*

A15. Generative AI/foundation models are increasingly becoming utilized in for a wide range of clinical and non-clinical tasks. While these innovations promise to significantly improve future patient care, they can pose novel oversight challenges for regulators and health care providers, particularly when used for functions such as diagnosis or treatment.

Q16. *Are there any promising innovations or potential disruptions on the horizon that warrant attention from policymakers?*

A16. FDA may need authority expansions or clarifications from Congress to make necessary enhancements to the health care AI regulatory process including post-market surveillance mechanisms and programs to ensure the ongoing safety and effectiveness of AI such as "unlocked" (i.e., locally self-improving) tools. FDA may also need authority expansions or clarifications to post-market surveillance programs to address various novel regulatory challenges posed by Generative AI—for example, local use of foundation models by health care institutions to create or enable medical device software functions.

Q17. *Are there legislative measures that Congress can take to ensure access to safe, reliable AI healthcare services?*

A17. The health care sector is comparatively well-positioned for AI policy solutions, as there are existing federal agencies and regulatory programs that oversee medical devices, protect patient data, accredit health care facilities, fund research, etc. Congress should empower existing health care regulatory agencies. Examples of novel ways to leverage existing regulatory agencies to tackle AI-specific challenges include:

- ***Expand FDA's authority beyond medical devices*** to comprehensively encompass all health care AI, including those software functions that do not meet the current "medical device" definition as amended by the 21st Century Cures Act Sec. 3060. FDA could implement a safety- and effectiveness-tiered AI oversight system with various premarket vetting and post-market surveillance requirements based on what regulators, providers, and patients need to trust that the AI performs as intended in real-world clinical environments. This could leverage the capabilities of public-private partnerships such as the ACR Transparent-AI initiative discussed in A6 above. This new paradigm could be created within the Center for Devices and Radiological Health, Digital Health Center of Excellence, or via establishment of a new FDA Center for Health Care AI. Voluntary health IT certification and/or optional standards do not adequately substitute for FDA oversight.

- ***Create payment policy incentives within CMS to promote widespread clinical use of high-quality AI in healthcare.*** Quality and value-based reimbursement policy should focus on providing end-users with incentives designed to increase the adoption of clinical AI devices that are safe and effective (as outlined in A14) for use by qualified end-users, in AI accredited health care facilities, with appropriate privacy/security safeguards and governance practices. As mentioned in A5, the ACR is currently a CMS-designated accreditation organization of advanced diagnostic imaging suppliers, and could similarly perform accreditation of safe and effective use of AI to enable value-based reimbursement.

Thank you for your consideration of these comments. The ACR is committed to serving as a trusted health care AI policy resource to Congressman Bera and the Bipartisan House Task Force on AI. Please contact Cindy Moran, Executive Vice President, Government Relations, Economics, and Health Policy, at cmoran@acr.org with questions or input opportunities.

Sincerely,

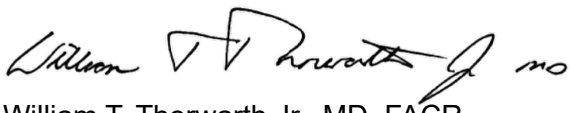

William T. Thorwarth Jr., MD, FACR
Chief Executive Officer
American College of Radiology