RADIOLOGICAL SOCIETY OF NORTH AMERICA
820 JORIE BLVD., SUITE 200, OAK BROOK, IL 60523-2251
TEL 1-630-571-2670 FAX 1-630-571-7837
RSNA.ORG

RSNA® Radiological Society of North America

February 23, 2026

The Radiological Society of North America (RSNA) is a leading global organization representing over 52,000 radiologists and medical imaging professionals across 150 countries. Radiology and medical imaging are among the most data-intensive fields in medicine, and AI-driven technologies have already begun transforming clinical practice. Radiology has experienced the highest rate of medical AI tool development and deployment, with more than 75% of the over 1,000 Food and Drug Administration (FDA)-cleared AI algorithms designed for radiological applications.

RSNA's guiding principle when considering how to best accelerate the adoption and implementation of AI in clinical care is unwavering commitment to patient safety. Federal oversight of AI should be grounded in a risk-based, proportional approach—placing the strongest regulatory focus on applications that directly impact clinical decision-making or patient outcomes. Streamlined regulatory pathways and clear guidance are essential but must be paired with robust safety guardrails to protect patients. When thoughtfully designed, regulation does not hinder innovation; instead, it establishes a trusted foundation for progress. This ensures that AI technologies strengthen, rather than compromise, the integrity of medical practice and patient care.

Our comments, organized below by specific RFI questions, encompass issues relevant to all agencies at the Department of Health and Human Services (HHS), including but not limited to OSTP/ONC. This approach ensures a holistic perspective on regulatory, policy, and programmatic considerations.

Specific Questions

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

As highlighted in President Trump's AI Action Plan, the primary obstacles to realizing AI's full potential in healthcare are not simply the availability of models, tools or technical capability. Rather, accelerating the adoption and implementation of AI in clinical care requires the careful building of trust among clinicians, patients, and health systems. Because clinical decisions have direct, and often irreversible consequences for patients, accelerating the adoption of AI in healthcare requires a thoughtful and balanced regulatory approach that ensures patient safety without stifling innovation. Vendors and healthcare organizations face implementation challenges that are distinct from those in other industries, including variability in patient populations, workflow complexity, and the need for continuous clinical oversight.

The successful adoption of AI in clinical care depends not only on smart regulatory policy, but on confidence across the healthcare ecosystem. Regulators, clinicians, developers, and health system leaders must share common expectations around transparency, explainability, and ongoing performance monitoring. Without consistent frameworks to support these expectations, uncertainty among care providers and patients regarding the risks and trustworthiness of AI is a key barrier to adoption, even when AI tools demonstrate technical promise.

Organizational readiness is also a barrier to more widespread implementation of AI. Many healthcare institutions lack the internal data infrastructure, governance processes, and technical expertise needed to evaluate, deploy, and monitor AI tools effectively. Investments in AI literacy, standardized data practices, and operational infrastructure, across both public agencies and private healthcare organizations, are necessary to support safe, and effective implementation. Such investments help ensure that AI tools are used as intended, interpreted appropriately by clinicians, and monitored over time to maintain clinical performance and patient safety.

Additional barriers include the fact that AI-enabled tools are constantly evolving, and that performance can change across care settings. Many radiology practices report that AI algorithms do not perform as expected when applied to their specific patient populations, necessitating rigorous local testing, and validation prior to clinical implementation. Radiology practices have also learned that AI models must be monitored closely to prevent "diagnostic drift," in which even minor changes to algorithms, input data, or clinical context can result in clinically meaningful differences in output. These challenges are amplified by the current lack of standardization across AI tools, as algorithms designed for similar purposes may generate outputs using different formats, thresholds, or assumptions. As a result, modifications to AI tools must be transparent and readily available for revalidation by end users to ensure continued alignment with clinical expectations and patient safety.

2) **What regulatory, payment policy, or programmatic design changes should HHS prioritize to incentivize the effective use of AI in clinical care and why? What HHS regulations, policies, or programs could be revisited to augment your ability to develop or use AI in clinical care? Please provide specific changes and applicable Code of Federal Regulations citations.**

AI has the potential to improve the efficiency, quality, and consistency of clinical care, but its successful implementation depends on a regulatory approach that enables innovation while maintaining high standards for safety and AI performance. Clear, predictable, and adaptive regulatory pathways can reduce uncertainty for developers and healthcare organizations while reinforcing trust in AI-enabled clinical tools.

Barriers to innovation - AI innovation in medical imaging is often constrained by regulatory frameworks designed for static technologies rather than adaptive systems. The FDA's authority to regulate medical devices is rooted in the 1976 Medical Device Amendments to the 1938 Federal Food, Drug, and Cosmetic Act, which were written for human-operated, and largely unchanging, devices. As a result, AI tools are regulated under Software as a Medical Device (SaMD) and Software in a Medical Device (SiMD) frameworks that do not fully accommodate continuous learning, recalibration, or iterative improvement.

Under current approaches, many AI models are required to remain "locked" following clearance or approval, limiting the ability to implement updates that could preserve accuracy, reliability, and safety over time. This rigidity can be misaligned with the realities of clinical practice, where patient populations, imaging protocols, and care environments evolve continuously. In such settings, the inability to make controlled updates may increase, rather than reduce clinical risk.

RSNA supports regulatory approaches that enable safe, monitored adaptation of AI tools while maintaining appropriate oversight. Smart regulation should distinguish between modifications that can be managed within predefined boundaries through ongoing monitoring and those that warrant more comprehensive re-evaluation.

Enabling this distinction would allow AI tools to improve responsibly over time, support long-term clinical performance, and better align regulatory oversight with the dynamic nature of AI technologies and real-world clinical use.

Outdated or misaligned frameworks - RSNA supports efforts to modernize regulatory approaches to better reflect the lifecycle realities of AI. Many existing device regulations were developed to govern episodic software updates with fixed performance characteristics. Rigid application of these frameworks to AI tools risks slowing innovation without improving patient safety. Unlike traditional software, AI system behavior may vary across clinical contexts, patient populations, imaging protocols, and patterns of use, even in the absence of explicit software updates.

Effective oversight must therefore address not only how a system is designed to operate, but how it performs in real-world clinical settings over time. Clinicians and regulators should have appropriate visibility into the parameters, assumptions, and trade-offs influencing AI outputs, particularly when those outputs inform clinical decision-making. Such transparency is essential to support accountability, enable meaningful monitoring, and allow users to recognize when an AI tool may be operating outside its intended context.

RSNA supports risk-based and proportional regulatory approaches that focus oversight on AI applications with direct influence on clinical decisions and patient outcomes. Adaptive regulatory models that account for context-dependent performance and real-world use can better align oversight with how AI functions in practice, supporting innovation while maintaining high standards of safety and clinical reliability.

Underused administrative tools - Existing mechanisms such as FDA's Predetermined Change Control Plans (PCCP) and related pilot programs offer important opportunities for safe, adaptive AI deployment. However, these tools remain voluntary and inconsistently implemented. Broader adoption and clearer alignment of these mechanisms across federal agencies could reduce regulatory redundancy, improve predictability for developers and health systems, and support timely access to safe and effective AI innovations.

While RSNA recognizes that current PCCP frameworks emphasize vendor self-monitoring, we have concerns that self-monitoring alone may be insufficient in complex clinical environments such as radiology. Radiology practices often deploy multiple AI tools from different vendors simultaneously, each with distinct functions, outputs, and update cycles. These tools may interact with shared workflows, data sources, and clinical decision processes in ways that individual vendors may not be positioned to fully observe or anticipate. In such settings, reliance solely on vendor-reported deviations risks leaving clinically meaningful performance issues undetected.

To strengthen oversight while preserving flexibility, RSNA recommends a dual-reporting model that allows both vendors and end users, like radiologists and healthcare organizations, to report observed deviations in AI tool performance directly to the FDA. In addition, RSNA encourages the FDA to establish a formal mechanism through which end users can request modifications or updates to PCCPs based on real-world experience. Radiologists may identify context-specific performance issues, population-related limitations, or workflow interactions that were not evident during initial development or testing, but that materially affect safety or effectiveness in practice. Creating structured pathways for this feedback would promote collaboration among developers, clinicians, and regulators, support continuous improvement, and help ensure that AI tools remain aligned with evolving clinical needs while prioritizing patient safety.

Structural incompatibility and need for clarity - Effective adoption of AI in clinical care depends on sustained cross-agency coordination among ASTP/ONC, FDA, and the Centers for Medicare & Medicaid Services (CMS). Greater alignment across these agencies on definitions, data standards, and post-market surveillance expectations would reduce duplication, close regulatory gaps, and promote trust among developers, healthcare organizations, clinicians, and patients.

Clearer coordination is particularly important at the intersection of regulation, interoperability, and reimbursement. Consistent expectations for data capture, performance monitoring, and reporting would enable healthcare organizations to implement AI tools more efficiently while supporting oversight across the full lifecycles of use. Harmonized approaches would also help ensure that safety monitoring and accountability mechanisms scale across institutions and care settings.

Complementary refinement of data privacy frameworks, including those established under the Health Insurance Portability and Accountability Act (HIPAA), could further support responsible AI development and deployment. Clear, consistent guidance on de-identification, data minimization, and secondary use would enable the appropriate use of large, privacy-protective datasets for AI training, validation, and ongoing performance monitoring. Establishing secure and compliant pathways for multi-institutional data aggregation would improve model generalizability, reduce performance variability across settings, and support more reliable clinical use.

Finally, clearer interpretive guidance, developed in collaboration with clinical experts and professional societies, would help ensure that safety and compliance expectations are transparent, consistent, and practical to implement. Such clarity is essential to fostering innovation while maintaining strong protections for patients and ensuring public trust.

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?**

The increasing use of AI tools that are not regulated as medical devices, including general-purpose decision support, workflow automation, and documentation tools, introduces novel legal and implementation considerations when these technologies are deployed in clinical care. These include questions related to clinical accountability, liability allocation, data provenance, privacy protections, and cybersecurity, particularly when AI outputs inform or influence clinical decision-making without clear regulatory oversight.

While RSNA does not recommend that such tools should be regulated as medical devices, clearer federal guidance is needed to delineate expectations for transparency, documentation, and accountability when non-medical AI tools are used in clinical environments. HHS can play a valuable role by coordinating cross-agency guidance that clarifies how existing frameworks, including privacy, security, and information governance requirements, apply to these tools; support best practices for responsible deployment; and reinforce that ultimate clinical responsibility rests with licensed healthcare professionals. Such guidance would help reduce uncertainty for health systems and clinicians while maintaining appropriate safeguards for patients.

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)?**

The broader adoption of AI in radiology and healthcare depends on trust among clinicians, patients, healthcare organizations, and regulators. As described above, to cultivate this trust, post-deployment monitoring of AI tools must be standardized, with clear mechanisms for evaluating real-world performance and safety over time. RSNA urges HHS to support ongoing validation studies of AI models in different clinical settings; establish transparent reporting mechanisms for AI efficacy and safety; and encourage collaboration between regulatory agencies, professional societies, and healthcare institutions to ensure continuous oversight while avoiding unnecessary administrative burden.

RSNA has gained experience in validating AI-enabled diagnostic imaging tools that have relevance for evaluation of non-medical devices. RSNA has curated large, expert-annotated datasets for use in a series of AI challenge competitions and made them freely available through its Medical Imaging Resource for AI (MIRA) platform (https://mira.rsna.org). These datasets span multiple imaging modalities and clinical specialties and include training, validation, and sequestered test partitions. These resources are used to support public benchmarking and independent evaluation of AI-enabled medical devices. Model performance is assessed using standard metrics selected to determine accuracy and generalizability. Together, these datasets and evaluation approaches provide a foundation for transparent, reproducible assessment aligned with FDA and National Institutes of Health (NIH) priorities for lifecycle evaluation of AI-enabled medical devices.

Meaningful evaluation of AI-enabled medical devices in radiology requires multidimensional performance assessment. Core metrics include diagnostic accuracy measures (such as sensitivity, specificity, and area under the curve); performance consistency across patient populations and clinical contexts; workflow impact; technical efficiency of algorithm execution; and usability of AI outputs for clinical end users.

Equally important is assessment of performance stability over time. Indicators of longitudinal reliability include calibration drift, changes in error rates, diagnostic concordance with clinical reference standards, and structured feedback from users regarding confidence and appropriate use. Reliability should also be evaluated across imaging devices, manufacturers, modalities, protocols, and acquisition conditions to ensure consistent performance in real-world clinical environments.

Future benchmarking and assessment frameworks should continue to evolve toward comprehensive, multidimensional evaluation. In addition to diagnostic performance, these frameworks should incorporate analysis of execution boundaries and failure modes; clarity and interpretability of outputs; quality of explanatory information provided to users; and consistency of communication within clinical workflows. Such approaches will support responsible AI deployment, continuous improvement, and sustain trust in AI-enabled clinical tools.

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?**

RSNA has a long-standing role in advancing technological innovation in medical imaging and in convening stakeholders across academia, industry, and government to translate emerging technologies into clinical practice. We stand ready to collaborate with HHS and other federal partners to support the responsible development, validation, and deployment of AI-enabled clinical tools.

RSNA's contributions reflect sustained investment in the foundational infrastructure required for scalable and trustworthy AI in healthcare. These efforts include leadership in the adoption of the DICOM imaging standard, which accelerated the digitization and interoperability of radiology; long-standing support for the Integrating the Healthcare Enterprise (IHE) initiative to improve interoperability between diagnostic systems, including AI-enabled tools in diagnostic imaging, and electronic health records; and the development of widely adopted semantic standards, including RadLex, the LOINC-RSNA Radiology Playbook, and RadElement, that enhance the machine readability and consistency of medical imaging data.

In addition, RSNA has played a significant role in advancing AI research and evaluation through the organization of eleven international AI challenge competitions since 2017, which engage global research communities in developing and validating algorithms for clinically relevant tasks. RSNA also collects, curates, and annotates large-scale medical imaging datasets through its Medical Imaging Resource Center (MIRA) platform to support training, benchmarking, and independent evaluation. Further, RSNA has helped collect, curate, and submit data to the Medical Imaging Data Resource Center (MIDRC), which aggregates and annotates multi-institutional imaging data to support AI research, validation, and generalizability.

Together, these efforts position RSNA as an effective capable of supporting public–private collaborations that advance AI innovation while emphasizing interoperability, validation, transparency, and patient safety. Continued collaboration between federal agencies and professional societies such as RSNA can help ensure that AI development is aligned with real-world clinical needs and supported by durable technical and governance infrastructure.

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?**

Where AI tools have met expectations *- AI tools have most consistently met or exceeded performance and cost expectations when they are applied to narrowly defined, high-volume clinical tasks with well-characterized workflows and measurable operational impact.* In radiology and other data-rich specialties, examples include:

- Workflow prioritization and triage, particularly for time-sensitive conditions, where AI can support earlier review and reduced turnaround times without displacing clinical judgment.
- Quantitative image analysis and measurement support, such as segmentation, volumetry, and standardized measurements, which improve consistency and reduce manual effort.

- Operational and technical quality support, including detection of incomplete studies, protocol deviations, and image quality issues that reduce repeat imaging and downstream costs.
- Automation of low-risk administrative tasks, such as routing, protocol suggestions, and documentation support, where benefits are realized through efficiency gains rather than diagnostic substitution.

In these settings, AI provides value when it is tightly integrated into existing clinical systems, minimizes additional cognitive or documentation burden, and produces outputs that are easy for clinicians to interpret and verify.

Where AI tools have fallen short *- AI tools have most often fallen short when deployed into complex clinical environments without sufficient attention to validation, integration, and ongoing governance.* Common challenges include:

- Performance variability across real-world settings, driven by differences in patient populations, imaging equipment, acquisition protocols, and local practice patterns.
- Limited generalizability and poorly defined boundary conditions, where tools perform well in development or trial settings but degrade in routine clinical use, particularly for uncommon presentations or edge cases.
- Workflow friction and usability challenges, including alert fatigue, unclear output presentation, and added steps that reduce clinician adoption even when algorithmic performance is adequate.
- Insufficient post-deployment monitoring, including limited transparency into version changes, delayed detection of performance drift, and lack of standardized mechanisms for real-world evaluation.
- Economic misalignment, where the entities responsible for purchasing AI tools do not directly realize downstream financial or operational benefits, complicating adoption decisions.
- Interoperability and standardization gaps, including inconsistent output formats across vendors and limited integration into structured clinical documentation and downstream care pathways.
- Unclear boundary conditions and off-label use, particularly when AI tools cleared for triage or workflow prioritization are used in practice as diagnostic or detection aids. In high-volume clinical environments, outputs intended to flag studies for expedited review may be interpreted as indicative of disease presence or absence, influencing diagnostic reasoning without having been validated for that purpose.
- Scalability and cumulative implementation burden, particularly as multiple AI tools are applied concurrently to the same clinical data. As AI adoption grows, clinicians and health systems increasingly contend with overlapping alerts, notifications, and outputs generated by multiple algorithms operating on the same imaging studies or clinical workflows. Without careful coordination, this can contribute to alert fatigue, reduced clinician attention to high-value signals, and difficulty attributing downstream effects to individual tools. In addition, monitoring the combined impact of multiple algorithms, including interactions between tools from different vendors, presents operational and governance challenges that are not well addressed by current implementation or oversight models.

These challenges underscore that successful AI adoption depends as much on implementation, governance, and context of use as on model accuracy.

AI capabilities with the greatest potential for future impact - *The greatest opportunities for AI to improve health outcomes, generate meaningful quality insights, and reduce costs lie in applications that extend beyond*

isolated detection tasks and instead support longitudinal care, decision-making, and system-level efficiency. Priority areas include:

- Longitudinal imaging intelligence, tools that track findings over time, identify clinically meaningful changes, and support follow-up recommendations and adherence, particularly for chronic disease and incidental findings.
- Context-aware decision support, tools that combine imaging with structured clinical data such as laboratory results, prior diagnoses, and clinical history to improve appropriateness, reduce unnecessary imaging, and support timely care.
- AI-enabled quality measurement and safety monitoring, including low-burden generation of performance metrics related to diagnostic consistency, technical quality, and care pathway adherence, supporting learning health system feedback loops.
- Protocol optimization and resource stewardship, tools that assist with exam selection, protocol tailoring, dose and contrast optimization, and scheduling, helping reduce repeat imaging and improve throughput.
- Safety-focused AI, including tools that detect preventable errors, high-risk scenarios, and breakdowns in communication, supporting systematic risk reduction rather than individual fault attribution.
- Population- and specialty-specific AI, particularly tools designed and validated for pediatric populations and other specialized care settings, such as rare disease, where general-purpose models often underperform.
- Foundation models capable of supporting multiple diagnostic tasks, including the assessment of several conditions within a single study or encounter. Unlike narrowly trained algorithms, these models may offer efficiencies by reducing the need to deploy and manage multiple task-specific tools. However, realizing their potential will require careful validation of performance across conditions, clear definition of intended use, and robust monitoring to ensure reliability as clinical complexity increases.
- AI tools designed to identify normal findings or the absence of disease, which may offer substantial value by reducing unnecessary follow-up, supporting workload management, and improving prioritization of abnormal studies. Reliable identification of normal examinations could help clinicians focus attention on higher-risk cases, but such tools must be held to high standards of sensitivity, transparency, and post-deployment monitoring given the consequences of missed pathology.
- Agentic AI workflows that automate multi-step clinical and operational processes, such as study routing, protocol selection, follow-up tracking, and communication of results. These tools have the potential to reduce administrative burden and improve consistency across care pathways, but they also introduce new governance challenges related to accountability, error propagation, and monitoring of downstream effects across interconnected systems.
- Ambient and context-aware AI models that integrate clinical conversations and electronic health record data, with the goal of supporting appropriate imaging selection and care coordination. By synthesizing information from clinician–patient interactions and structured clinical data, these tools could improve appropriateness, reduce unnecessary imaging, and support more personalized care. Effective deployment will require clear safeguards for privacy, transparency about how recommendations are generated, and careful evaluation of how such tools influence clinical decision-making.

Across these domains, the most impactful innovations will be those that are interoperable, transparent, and supported by standardized post-deployment monitoring, allowing performance to be maintained as clinical environments evolve.

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?

Based on extensive experience supporting AI development, validation, and deployment in medical imaging, RSNA observes that the most significant barriers to AI adoption in clinical care are administrative and operational rather than technical. These challenges consistently limit the ability of healthcare organizations to implement AI safely, effectively, and at scale.

A primary barrier is fragmented governance and unclear accountability for AI oversight. Many organizations lack standardized frameworks to guide validation, monitoring, and lifecycle management once AI tools are deployed. RSNA recommends that HHS support the development and dissemination of clear governance models and best practices that define roles, responsibilities, and escalation pathways for AI oversight, particularly for post-deployment monitoring.

Implementation and maintenance burden also pose a major challenge. Healthcare organizations must dedicate substantial resources to local validation, workflow integration, cybersecurity review, staff training, and ongoing performance monitoring. RSNA encourages HHS to prioritize policies and programs that reduce duplicative requirements, promote reuse of standardized validation and monitoring approaches, and leverage automation to lower the operational cost of safe AI deployment.

Scalability challenges are becoming increasingly pronounced as multiple AI tools are applied concurrently to the same clinical data. RSNA has observed that overlapping alerts, notifications, and outputs from different algorithms can contribute to alert fatigue and reduce clinician confidence. In addition, monitoring the cumulative impact of multiple AI systems, including interactions between tools from different vendors, is not well supported by current oversight models. RSNA recommends that HHS support scalable monitoring frameworks and interoperable standards that enable health systems to assess both individual and combined AI tool performance.

Interoperability and integration limitations further impede adoption. Inconsistent data standards and non-standardized AI outputs increase implementation complexity and limit the ability to efficiently monitor real-world performance. RSNA recommends continued federal investment in interoperable data standards, structured outputs, and standardized performance metrics to support safe and scalable AI integration into clinical workflows.

Finally, economic and regulatory uncertainty remains a persistent barrier. Healthcare organizations often face unclear reimbursement pathways, misalignment between implementation costs and downstream benefits, and uncertainty regarding post-deployment compliance expectations. RSNA encourages HHS to provide clearer guidance and alignment across regulatory, reimbursement, and interoperability initiatives to reduce uncertainty and support sustainable AI adoption.

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.**

As front-line users of AI-enabled medical devices and physicians ultimately responsible for patient care and safety, radiologists rely on AI as decision-support that complements, but does not replace, clinical judgment. To fulfill this role safely and effectively, AI outputs must be transparent, interpretable, and presented in a manner that allows clinicians to contextualize results within the broader clinical picture.

Given current difficulties in obtaining valid data for training and testing AI models, broader interoperability, data consistency and accessibility would contribute powerfully to accelerating innovation in AI for diagnostic imaging.

As described above, RSNA has conducted a series of AI challenges over the past ten years and is thus very familiar with the difficulty of collecting and annotating the datasets needed for AI model development. We have assembled a consortium of data-contributing sites from around the world and engaged expert radiologists to label imaging datasets for use cases associated with each of these challenges. Our experience has shown the high expense and limited scalability of this approach and the restricted applicability for clinical use of models generated through these methods.

In addition, the closed nature of the network of health information systems in clinical use and the absence of well-defined points of integration for new tools, including AI tools for diagnostic imaging, presents significant challenges to accelerating the development of AI for clinical care. To make them capable of efficient integration in the diagnostic imaging environment, developers building AI-enabled tools must invest in extensive research to determine how to position them within a complex, heterogeneous and inconsistently linked array of systems and workflows. While incumbent vendors now offer a variety of platforms for integrating AI tools, they are based on proprietary integrations and thus force developers into restrictive commercial relationships and limit access to best-of-breed tools for potential users. The absence of clear standardized APIs currently makes it difficult to integrate AI-enabled tools efficiently into workflow and to evaluate and monitor their performance.

RSNA believes that ASTP-ONC can accelerate innovation in AI by facilitating access to clinical data, including medical imaging, for AI development, validation, and monitoring. While our understanding is that ASTP-ONC has generally avoided imposing requirements regarding medical devices, RSNA believes that diagnostic imaging must be considered a core component of the patient medical record and that substantial benefits for patient care and technologic innovation can be realized through a consistent set of regulations that foster interoperability and improved access across health information technology platforms.

The USCDI defines an evolving set of data elements to be included in the electronic health record, and the Health IT Certification Program defines standard Application Programming Interfaces for accessing and using these data. RSNA supports expanding the USCDI to include effective access to imaging studies (e.g. through the proposed Imaging Reference data element) and including view, download and transmit requirements for diagnostic imaging in the Health IT Certification Program. Providing API-based access to imaging alongside other relevant clinical data and allowing use of this data for research purposes (with proper security and privacy protections) would facilitate the creation of multimodal datasets for building AI tools and enable multicenter, prospective validation of AI models using real-world clinical and imaging data.

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?**

Our understanding of patients' goals for integrating AI into healthcare closely align with those of RSNA. Chief among these is enhancing the quality of care and increasing access to healthcare services. AI holds significant potential to broaden access to high-quality, personalized healthcare for all Americans, particularly those in underserved and rural communities. Additionally, AI tools that streamline administrative workflows can help ensure patients receive the right care at the right time by reducing barriers and delays.

Both patients and radiologists are optimistic about AI's capacity to improve disease detection and diagnosis, which can strengthen prevention strategies and lead to better outcomes through earlier identification and treatment. Importantly, there is broad consensus that AI should serve to support and enhance physicians' performance rather than replace their critical role in diagnostic decision-making. AI's strength lies in synthesizing complex medical histories, medication lists, and patient preferences, which enables more personalized treatment plans and smoother transitions of care. For caregivers managing chronic or complex conditions, improved care coordination made possible by AI is especially valuable.

In our view, many of the concerns that patients and caregivers have regarding the adoption of AI in clinical care center on issues of trust, transparency, and explainability. Patients seek clear explanations of how AI contributes to medical decisions, as well as assurance that clinicians understand and can appropriately question AI outputs. Central to building this trust is the clinician's ability to effectively validate and continuously monitor AI tools as they are implemented in real-world care settings.

10) **Are there specific areas of AI research that HHS should prioritize to accelerate the adoption of AI as part of clinical care?**
    a. **Are there published findings about the impact of adopted AI tools and their use clinical care?**
    b. **How does the literature approach the costs, benefits, and transfers of using AI as part of clinical care?**

The effectiveness of AI in healthcare depends on the strength of the research and data infrastructure that supports its development, evaluation, and deployment. Yet, there is a relative dearth of literature on the actual impact of AI-enabled tools in clinical care, and the limited evidence that exists is often conflicting.[1,2] RSNA recommends that HHS and its research funding agencies strategically invest in every stage of the AI development process to optimize clinical outcomes and ensure safe adoption in healthcare.

First, foundational research should focus on building robust, adaptable, and generalizable algorithms that can handle a wide variety of clinical scenarios and patient populations. This approach will help prevent performance

[1] Najjar, R. Redefining Radiology: A Review of Artificial Intelligence Integration in Medical Imaging; Diagnostics, 13(17):2760 (2023). doi: 10.3390/diagnostics13172760,

[2] Buijs, E., Maggioni, E., Mazziotta, F. et al. Clinical impact of AI in radiology department management: a systematic review. Radiol med 129, 1656–1666 (2024). doi.org/10.1007/s11547-024-01880-1.

variability and "diagnostic drift," ensuring AI solutions remain effective across different healthcare settings and over time. Next, translational and implementation research is crucial for integrating AI into real-world clinical workflows. These studies should address how AI tools can be efficiently embedded into existing systems, maintain high standards for quality and patient safety, and foster seamless collaboration between technology and clinicians. Additionally, research must examine human-AI interactions for both physicians and patients, exploring factors such as trust, explainability, and user experience to identify ways to build confidence and understanding in AI-assisted care. Finally, continued development and expansion of data and reporting standards is essential to support interoperability, enable benchmarking across institutions, and establish comprehensive, ongoing performance monitoring. Such standards will facilitate transparent evaluation and continuous improvement, helping healthcare organizations responsibly deploy AI and maintain excellence in patient safety and clinical care.

As AI tools become increasingly embedded in clinical practice, it is essential that radiologists, technologists, and other healthcare professionals are equipped with the knowledge needed to critically evaluate, interpret, and appropriately use these technologies. RSNA encourages HHS to support federally funded AI education and literacy initiatives tailored to healthcare professionals, with a focus on understanding intended use, limitations, workflow integration, and patient safety implications.

RSNA also recommends collaboration between HHS and professional societies to develop standardized education, training, and certification pathways that support AI competency in clinical care. Ongoing education initiatives should extend to both clinicians and patients, ensuring transparency around how AI is used in care delivery and reinforcing the role of clinicians as accountable decision-makers. Strengthening education and workforce preparedness will be essential to sustaining trust, safety, and effective adoption of AI across healthcare.

RSNA values the opportunity to provide input on this Request for Information and looks forward to continued collaboration with HHS Office of the Deputy Secretary and ASTP/ONC. Together, we aim to establish smart and balanced regulatory frameworks for AI in healthcare that uphold patient safety without stifling innovation. For additional information or questions, please contact RSNA's Director of Government Relations, Libby O'Hare (eohare@rsna.org).

Sincerely,

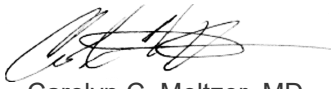

Carolyn C. Meltzer, MD
Chair of the Board
Radiological Society of North America