American College of Radiology™

April 7, 2025

Dockets Management Staff (HFA-305)
Food and Drug Administration
5630 Fishers Lane, Rm. 1061
Rockville, MD 20852

Re: (Docket No. FDA-2024-D-4488) Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations; Draft Guidance for Industry and Food and Drug Administration Staff; Comments of the American College of Radiology

The American College of Radiology (ACR)—a professional association representing more than 40,000 physicians practicing diagnostic radiology, interventional radiology, radiation oncology, and nuclear medicine, as well as medical physicists—appreciates the opportunity to comment on the draft guidance, "Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations" (Docket No. FDA-2024-D-4488), published by the U.S. Food and Drug Administration (FDA) on January 7, 2025. The ACR supports the advancement of safe, effective, and clinically useful artificial intelligence-enabled device software functions (AI-DSFs) for use by radiologists to improve patient care and appreciates the FDA's extensive efforts to clarify the management of risk throughout the total product life cycle (TPLC).

ACR Data Science Institute

The ACR established its Data Science Institute (DSI) in 2017 to advance safe, effective, and clinically useful radiology AI innovations. The ACR DSI collaborates with radiology professionals, industry, government, patients, and other stakeholders in developing programs and tools in support of the implementation of AI applications that will help improve patient care. Initiatives include:

- Defining clinically relevant use cases intended to guide the development of useful imaging AI (Define-AI).
- Establishing the first national recognition program for safe and effective implementation of AI in imaging practices (ARCH-AI), and making it broadly available in 2024.
- Creating opportunities for monitoring the effectiveness of AI models, including via the first large-scale quality registry for AI performance monitoring in real-world clinical practice (Assess-AI).

- Sharing information about radiology AI models with radiologists to help them choose what works for their practices and their patients (AI Central).
- Addressing the regulatory, legal, and ethical issues associated with radiology AI.

General Recommendations

Require more specificity about "intended users"
Intended user information in public-facing regulatory documentation, such as 510(k) summaries, is often ambiguous. For example, intended users of AI-DSFs with outputs that appear to necessitate specialist-level physician interpretation may be broadly described as the "healthcare system" or "healthcare professional," or professions with varying levels of clinical training and experience. As a result, AI-DSFs may be marketed by vendors or implemented in real-world practices in ways unanticipated by device reviewers.

Address site-level validation
AI-DSF performance may vary across clinical settings due to variations in workflow, systems/devices, and patient demographics. Typically, clinical sites must validate AI-DSF accuracy and effectiveness prior to device deployment and medical use. FDA should ideally review manufacturers' plans for enabling site-level validation by provider facilities, including any related capabilities and recommendations.

Enhance post-market monitoring standards/requirements
The FDA's draft guidance encourages, but does not require, post-market monitoring as a TPLC risk control for AI-DSFs. Specific post-market monitoring mechanisms, such as periodic reporting or third-party audits, are not clearly outlined. Potential AI-DSF performance issues or other unintended consequences may go undetected without risk- and setting-appropriate post-market monitoring protocols. Manufacturers' plans, instructions, and/or recommendations to clinical sites for TPLC monitoring should be made available to FDA staff for consideration during pre-market review.

Clarify TPLC considerations for Predetermined Change Control Plans
The TPLC draft guidance references the Predetermined Change Control Plan (PCCP) guidance, which describes an approach for manufacturers to prospectively specify and seek premarket authorization for intended modifications to an AI-DSF (e.g., to improve device performance) without needing to submit additional marketing submissions or obtain further FDA authorization before implementing such modification consistent with the PCCP. However, FDA offers minimal insight into several TPLC considerations for AI-DSFs that have been modified pursuant to PCCPs in either of the guidance documents. PCCP-specific considerations should be added throughout all relevant sections of the TPLC draft guidance.

Discuss plans for AI-DSFs with unique safety/effectiveness considerations
The draft guidance narrowly applies to AI-DSFs with improvements that are locked by the manufacturer, including modifications that are rolled out to providers pursuant to the

manufacturer-defined PCCP. However, FDA should consider addressing novel oversight considerations for the following subtopics:

- TPLC risk management of autonomous AI and other AI-DSFs used by end-users without sufficient qualifications to appropriately serve as risk mitigations.
- Adaptive AI systems that evolve based on new data inputs, including specific provisions for real-time model updates and continuous validation.
- Generative AI (GenAI)-enabled DSFs, including "intended use" identification considerations and differentiation of GenAI-DSFs from non-DSFs in clinical sites.
- AI-enabled imaging reconstruction techniques, which may impact image quality and diagnostic interpretation.
- How sponsors should establish "fit-for-purpose" evidence when actual clinical data is scarce, and how to evaluate synthetic data in settings where obtaining large real datasets is not feasible.

If such content is not ready to be added to the initial version of the guidance, FDA should consider a brief section highlighting the agency's plans to address these and other subtopics in future revisions.

Specific Recommendations

How To Use This Guidance (Pages 5-7)
FDA explains how the draft guidance uses the term "validation" in contrast to common usage (lines 276-284); however, the document confuses terms in other sections. FDA should add an expanded glossary aligning each AI lifecycle phase with the recommended submission materials.

Device Description (Pages 8-9)
FDA proposes the following among other information included in Device Descriptions: "*A description of the intended users, their characteristics, and the level and type of training they are expected to have and/or receive. Users include those who will interpret the output. When relevant, list the qualifications or clinical role of the users intended to interpret the output. Users also include all people who interact with the device including during installation, use, and maintenance. For example, users may include technicians, health care providers, patients, and caregivers, as well as administrators and others involved in decisions about how to deploy medical devices, and how the device fits into clinical care"* (lines 381-388).

The ACR supports improved granularity and specificity of intended users within device descriptions. To best ensure device safety and effectiveness, the "end-user intended to interpret the output" (as described separately from other intended users in this draft guidance) should be qualified to provide the same medical procedure or service without

the AI-DSF output. This end-user serves as the immediate risk mitigation between the AI-DSF output and the patient.

For example, radiologists are trained to review and interpret medical imaging input data independently of any AI-DSF outputs they may use in their workflow. By contrast, non-imaging specialists providing imaging care enabled by AI-DSFs would be reliant on that AI-DSF output and therefore unable to identify subtle inaccuracies and performance issues due to drift and other variables. In this example, the qualified radiologist end-user would intrinsically serve as a device risk mitigation; the unqualified end-user could not serve in that same capacity.

User Interface and Labeling (Pages 10-15)
FDA should add information requesting any expected characteristics (e.g., functional capabilities, experience and knowledge levels, and level of training) of intended users to the labeling subsection on "Model Output" (page 13, lines 577-578). This information is currently only discussed under the "Model Input" subsection (page 13, lines 571-575). Specificity is particularly critical for understanding the minimum recommended training and experience of any user(s) primarily responsible for interpreting the AI-DSF output.

Performance Validation (Pages 27-32)
FDA should add content for manufacturers to promote and enable site-specific performance validation to ensure the AI-DSF performs as expected with each clinical site's systems/devices and patient populations.

The FDA should add more specificity and examples with respect to the "assessing the performance of the human-device team" (lines 1143-1173). The draft guidance recognizes the need to assess users' interaction with AI outputs, but only in general terms of human factors or reader studies.

Device Performance Monitoring (Pages 32-34)
FDA should add content to the draft guidance describing standards and rigorous mechanisms for post-market device performance monitoring, such as the use of third-party registries. Content should address establishing predefined triggers for additional analyses and concrete strategies for communicating performance drift to users. FDA should also describe any information needed by agency reviewers to evaluate AI-DSF performance monitoring tools.

Proactive monitoring is presented in the draft text as recommended but typically optional. However, AI-DSFs are sensitive to variability in data inputs, including changes in patient populations and—as is often the case in radiology—changes in other devices/systems that present image data to AI-DSFs. FDA should consider ways to decrease pre-market burden, for example via elimination of reader improvement studies, as a trade-off for risk-appropriate post-market monitoring.

Cybersecurity (Pages 34-36)
AI-specific cybersecurity risks are discussed in a limited, generalized manner without practical examples. FDA should map each recommended control to known AI-specific attack scenarios.

Public Submission Summary (Pages 37-39, bulleted items on lines 1525-1541)
510(k) summaries and other public submission summaries are currently the primary public means of finding regulatory information about devices using FDA's website. FDA should require that 510(k) summaries provide more detailed information on intended users, particularly regarding manufacturer-recommended minimum qualifications of end-users responsible for clinically interpreting the AI-DSF outputs. Additionally, FDA should require descriptions of manufacturer-recommended site-level performance validation and monitoring mechanisms or processes.

Appendices E and F - Model Card (Pages 50-64)
The ACR supports algorithm nutrition labels, or "Model Cards," that provide highlighted safety and effectiveness information to intended users and others. When situationally appropriate, FDA should require that the Model Card display relevant pediatric use considerations, such as a brief safety statement discussing whether the respective AI-DSF was trained and tested with pediatric data and/or relevant instructions for pediatric uses.

Additionally, the Model Card should prominently display specific and unambiguous manufacturer-recommended qualifications for intended users of the AI-DSF, focusing exclusively on the clinical end-users who are medically responsible for interpreting the AI-DSF output data. This could include interpreting end-users' medical specialties. To avoid confusion, the Model Card should not include unhelpfully broad descriptions or secondary professions/roles that may otherwise be included in the device's intended user population (some examples include "healthcare system," "healthcare professional," "hospital workers," etc.).

Thank you for your consideration of these comments. The ACR invites continued collaboration with FDA to improve the draft guidance and advance safe, effective, and clinically useful AI-DSFs. Please contact Michael Peters, Senior Director, Government Affairs, at mpeters@acr.org, with questions.

Sincerely,

Dana H. Smetherman, MD, MPH, MBA, FACR
Chief Executive Officer
American College of Radiology