NATIONAL HEALTH COUNCIL

April 7, 2025

Dockets Management Staff
U.S. Food and Drug Administration
5630 Fishers Lane, Room 1061
Rockville, MD 20852

RE: Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations [FDA-2024-D-4488]

Submitted via regulations.gov

To Whom It May Concern,

The National Health Council (NHC) appreciates the opportunity to submit comments on the Food and Drug Administration (FDA) draft guidance, "Artificial Intelligence-Enabled Device Software Functions: Lifecycle Management and Marketing Submission Recommendations." AI-driven medical technologies have the potential to improve diagnostic accuracy, enhance clinical decision-making, and increase efficiency in patient care.[1] However, their integration into health care also raises critical challenges, including transparency, consistency in patient outcomes, and the need for strong regulatory oversight.[2,3]

The NHC is uniquely positioned to provide input on this issue. Created by and for patient organizations over 100 years ago, the NHC brings diverse organizations together to forge consensus and drive patient-centered health policy. We promote increased access to affordable, high-value, equitable, and sustainable health care. Made up of more than 170 national health-related organizations and businesses, the NHC's core membership includes the nation's leading patient organizations. Other members include health-related associations and nonprofit organizations including the provider, research, and family caregiver communities; and businesses and organizations representing biopharmaceuticals, devices, diagnostics, generics, and payers.

---

[1] C.S. Ajmal, S. Yerram, V. Abishek, et al., "Innovative Approaches in Regulatory Affairs: Leveraging Artificial Intelligence and Machine Learning for Efficient Compliance and Decision-Making," *AAPS Journal* 27 (2025): 22, https://doi.org/10.1208/s12248-024-01006-5.

[2] Ciro Mennella, Umberto Maniscalco, Giuseppe De Pietro, and Massimo Esposito, "Ethical and Regulatory Challenges of AI Technologies in Healthcare: A Narrative Review," *Heliyon* 10, no. 4 (February 29, 2024): e26297, https://doi.org/10.1016/j.heliyon.2024.e26297.

[3] Pouya Kashefi, Yasaman Kashefi, and AmirHossein Ghafouri Mirsaraei, "Shaping the Future of AI: Balancing Innovation and Ethics in Global Regulation," *Uniform Law Review* 29, no. 3 (August 2024): 524–548, https://doi.org/10.1093/ulr/unae040.

With this broad, multi-stakeholder perspective, we stress that AI regulations must prioritize patient safety, consistent performance, and public trust. The NHC has published Principles on Health AI that emphasize patient benefit, transparency, privacy, accountability, and clinical relevance across patient populations. These principles guide our comments and align closely with the FDA's goals for responsible AI integration. We believe a robust regulatory framework can both encourage AI-driven innovation and ensure these technologies uphold the highest standards of ethics, scientific rigor, and patient-centeredness. In the spirit of collaboration, we offer the following recommendations to strengthen the FDA's guidance.

Summary of Recommendations

The NHC commends the FDA for its proactive efforts to establish a rigorous regulatory framework for AI-enabled medical devices. This draft guidance is a critical step toward ensuring that AI-driven technologies meet high standards for effectiveness, accountability, and continuous oversight. By addressing key issues such as risk management, consistency in AI-driven decision-making, and total product lifecycle (TPLC) oversight, the guidance provides an essential foundation for responsible AI deployment in health care. We welcome the opportunity for ongoing engagement with the FDA to advance a robust regulatory environment that maximizes the benefits of AI for patients and public health.

Based on the NHC's Principles on Health AI, we provide detailed recommendations to further refine the guidance, drawing on research global regulatory insights, and stakeholder perspectives to support a framework that fosters both innovation and patient protection.[4,5] Specifically, our recommendations address:

- Ethical and patient-centered considerations
- Comprehensive risk management, lifecycle oversight, post-market surveillance, and algorithmovigilance
- Rigorous performance consistency and data representativeness
- Robust human oversight and interpretability
- Stringent cybersecurity and data integrity protections

These recommendations aim to support the FDA's efforts to refine its regulatory framework, encouraging innovation while ensuring that AI technologies maintain the highest standards of safety, accuracy, and reliability.

[4] Rachel Yi Ling Kuo et al., "Stakeholder Perspectives Towards Diagnostic Artificial Intelligence: A Co-produced Qualitative Evidence Synthesis," *eClinicalMedicine* 71 (2024): 102555, https://doi.org/10.1016/j.eclinm.2024.102555.

[5] National Health Council, *NHC Statement on Artificial Intelligence and Health Care: Promise and Pitfalls*, statement for the record submitted to the Senate Finance Committee, February 8, 2024, https://nationalhealthcouncil.org/letters-comments/nhc-statement-on-artificial-intelligence-and-health-care-promise-and-pitfalls/.

Ethical and Patient-Centered Considerations

As AI/ML-enabled medical devices become increasingly embedded in clinical care, it is essential that their regulatory framework addresses not only technical performance but also the broader ethical and patient-centered dimensions that accompany their deployment. These technologies have the potential to shape critical health decisions, and the FDA's guidance must reflect the responsibility to ensure that patient safety, autonomy, transparency, and accountability remain at the forefront of innovation.[6,7]

One of the most critical ethical considerations is informed decision-making.[8] Patients and clinicians must understand when and how AI tools influence diagnostic, monitoring, or treatment recommendations. The FDA should require that sponsors clearly disclose the role of AI within a device, using language that is accessible to both medical professionals and lay users.[9] Labeling materials should describe how the technology works, its intended use, known limitations, and the circumstances under which human oversight is recommended. These measures will preserve patient autonomy, support shared decision-making, and enhance trust in AI-enabled tools.

Strong privacy protections are also essential.[10] AI/ML-enabled devices often rely on large datasets to support real-time learning, device adaptation, or retrospective performance monitoring.[11] The FDA should require that sponsors implement robust data governance practices—such as de-identification protocols, secure data storage, and access controls—that comply with applicable privacy laws and prevent misuse or unauthorized disclosures. Sponsors must also communicate how patient data are collected, processed, and used throughout the AI lifecycle, ensuring transparency and maintaining public confidence.

Also important is the need for clearly defined and delineated accountability.[12] When AI-generated outputs inform clinical decisions, there must be clear lines of responsibility

---

[6] Mitul Harishbhai Tilala et al., "Ethical Considerations in the Use of Artificial Intelligence and Machine Learning in Health Care: A Comprehensive Review," *Cureus* 16, no. 6 (June 15, 2024): e62443, https://doi.org/10.7759/cureus.62443.

[7] A. Shoghli, M. Darvish, and Y. Sadeghian, "Balancing Innovation and Privacy: Ethical Challenges in AI-Driven Healthcare," *Journal of Reviews in Medical Sciences* 4, no. 1 (2024): 1–11, https://doi.org/10.22034/jrms.2024.494112.1034.

[8] Moritz Reis, Florian Reis, and Wilfried Kunde, "Influence of Believed AI Involvement on the Perception of Digital Medical Advice," *Nature Medicine* 30, no. 11 (November 2024): 3098–3100, https://doi.org/10.1038/s41591-024-03180-7.

[9] Tilala et al., "Ethical Considerations in AI and Machine Learning," e62443.

[10] World Health Organization, *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models*.

[11] Nazish Khalid et al., "Privacy-Preserving Artificial Intelligence in Healthcare: Techniques and Applications," *Computers in Biology and Medicine* 158 (May 2023): 106848, https://doi.org/10.1016/j.compbiomed.2023.106848.

among developers, clinicians, and manufacturers. The FDA should require that sponsors document mechanisms for post-market surveillance, root-cause investigation of errors or malfunctions, and audit trails that track data inputs, model updates, and user interactions. In the event of a safety concern, these systems will enable regulators and developers to respond quickly, determine responsibility, and implement corrective actions. Such transparency is critical for ensuring that patients are protected and that the use of AI remains subject to rigorous oversight.[13]

The FDA should also consider the potential risks of premature AI deployment. While rapid access to innovative technologies can benefit patients, insufficient validation or poorly defined human-AI interaction protocols can introduce new safety risks.[14] Any regulatory flexibility granted to accelerate AI implementation—such as phased approvals or adaptive modification pathways—must be accompanied by strong post-market monitoring requirements to ensure ongoing protection.

By embedding these ethical and patient-centered safeguards into the regulatory framework for AI/ML-enabled medical devices, the FDA can support responsible innovation that enhances patient care while maintaining the core principles of safety, accountability, and transparency.

Comprehensive Risk Management and Lifecycle Oversight

AI-enabled medical devices diverge from traditional medical technologies in that they evolve continuously through iterative learning, software updates, and real-time refinements. While these characteristics can improve diagnostic accuracy and enhance clinical decision-making, they also present novel regulatory challenges that must be addressed.[15] The FDA's Total Product Lifecycle (TPLC) approach, as envisioned in its draft guidance on AI/ML-enabled devices, is a critical starting point for managing software-driven modifications, real-world performance monitoring, and continuous risk assessments.[16]

Within this framework, Predetermined Change Control Plans (PCCPs) can enable the controlled evolution of AI models post-market. However, additional regulatory mechanisms are necessary to balance innovation with patient protection.[17] We

---

[12] Dane Bottomley and Donrich Thaldar, "Liability for Harm Caused by AI in Healthcare: An Overview of the Core Legal Concepts," *Frontiers in Pharmacology* 14 (December 14, 2023): 1297353, https://doi.org/10.3389/fphar.2023.1297353.

[13] M.A.K. Akhtar, M. Kumar, and A. Nayyar, "Transparency and Accountability in Explainable AI: Best Practices," in *Towards Ethical and Socially Responsible Explainable AI*, vol. 551, *Studies in Systems, Decision and Control*, (Cham: Springer, 2024), https://doi.org/10.1007/978-3-031-66489-2_5.

[14] Mennella et al., "Ethical and Regulatory Challenges of AI," e26297.

[15] Mennella et al., "Ethical and Regulatory Challenges of AI," e26297.

[16] Kavitha Palaniappan, Elaine Yan Ting Lin, and Silke Vogel, "Global Regulatory Frameworks for the Use of Artificial Intelligence (AI) in the Healthcare Services Sector," *Healthcare* 12, no. 5 (February 28, 2024): 562, https://doi.org/10.3390/healthcare12050562.

encourage the FDA to further explore and expand pre-certification strategies—building on efforts like the Digital Health Pre-Certification Program—so that organizations with demonstrated quality and transparency can benefit from more efficient pathways for lower-risk AI modifications, while high-risk or novel AI technologies remain subject to rigorous premarket review.[18,19] Establishing an "AI regulatory sandbox," where manufacturers and the FDA can collaborate to test and refine models before full deployment, would also promote real-world validation while maintaining safety oversight.[20,21] These strategies not only enhance the adaptability of regulatory oversight but also position the United States to lead on the global stage. While international regulators such as the European Union and the United Kingdom's Medicines and Healthcare products Regulatory Agency (MHRA) have advanced post-market controls and emphasized the importance of human oversight, the FDA has an opportunity to define and elevate the global benchmark for AI oversight.[22,23,24,25,26,27] By actively

---

[17] Yu-Hao Li, Yu-Lin Li, Mu-Yang Wei, and Guang-Yu Li, "Innovation and Challenges of Artificial Intelligence Technology in Personalized Healthcare," *Scientific Reports* 14 (2024): article 18994, https://doi.org/10.1038/s41598-024-70073-7.

[18] U.S. Food and Drug Administration, *Marketing Submission Recommendations for a Predetermined Change Control Plan for Artificial Intelligence/Machine Learning-Enabled Device Software Functions: Draft Guidance for Industry and Food and Drug Administration Staff*, last modified April 2023, https://www.fda.gov/media/161815/download.

[19] Palaniappan, Lin, and Vogel, "Global Regulatory Frameworks," 562.

[20] OECD, "Regulatory Sandboxes in Artificial Intelligence," *OECD Digital Economy Papers*, no. 356 (2023), OECD Publishing, Paris, https://doi.org/10.1787/8f80a0e6-en.

[21] Sara Gerke, Timo Minssen, and Glenn Cohen, "Ethical and Legal Challenges of Artificial Intelligence-Driven Healthcare," in *Artificial Intelligence in Healthcare*, ed. Adam Bohr and Kaveh Memarzadeh (June 26, 2020), 295–336, https://doi.org/10.1016/B978-0-12-818438-7.00012-5.

[22] European Commission, "Regulatory Framework Proposal on Artificial Intelligence," *Shaping Europe's Digital Future*, accessed March 11, 2025, https://digital-strategy.ec.europa.eu/en/policies/regulatory-framework-ai.

[23] Mateo Aboy, Timo Minssen, and Effy Vayena, "Navigating the EU AI Act: Implications for Regulated Digital Medical Products," *NPJ Digital Medicine* 7 (September 6, 2024): 237, https://doi.org/10.1038/s41746-024-01232-3.

[24] Health Canada, *Pre-Market Guidance for Machine Learning-Enabled Medical Devices*, last modified February 5, 2025, https://www.canada.ca/en/health-canada/services/drugs-health-products/medical-devices/application-information/guidance-documents/pre-market-guidance-machine-learning-enabled-medical-devices.html.

[25] Medicines and Healthcare products Regulatory Agency (MHRA), *MHRA's AI Regulatory Strategy Ensures Patient Safety and Industry Innovation into 2030*, last modified April 30, 2024, https://www.gov.uk/government/news/mhras-ai-regulatory-strategy-ensures-patient-safety-and-industry-innovation-into-2030.

[26] U.S. Food and Drug Administration, Health Canada, and Medicines and Healthcare products Regulatory Agency, *Good Machine Learning Practice for Medical Device Development: Guiding Principles*, October 2021, https://www.fda.gov/media/153486/download.

shaping international regulatory frameworks—through participation in venues such as the International Medical Device Regulators Forum (IMDRF) and by promoting standards for AI validation, interpretability, and lifecycle governance—the FDA can ensure that U.S. regulatory science serves as the foundation for global best practices.[28] Reinforcing this leadership through harmonized policies will drive innovation, enhance cross-border consistency, and maintain the United States' position at the forefront of AI/ML-enabled medical device regulation.[29]

Managing Performance Risks and Long-Term Safety

As AI-enabled devices iterate and learn over time, performance may drift, and previously unintended patterns can emerge in clinical decision-making.[30,31] Changes in the underlying patient population, data inputs, or disease presentation can adversely impact sensitivity and specificity. Addressing performance drift is especially critical to ensure that model outputs remain valid and unbiased across patient populations.[32] Manufacturers should establish clear, pre-approved performance benchmarks that every software update or learning cycle must meet. Independent validation studies and real-world performance monitoring should be mandatory to confirm that iterative learning does not degrade model integrity or introduce disparities.[33,34]

Strong cybersecurity measures further bolster safety by mitigating data breaches or integrity attacks. AI-driven devices can become vulnerable targets when malicious

---

[27] Department of Health and Social Care (UK), *A Guide to Good Practice for Digital and Data-Driven Health Technologies*, last updated January 19, 2021, https://www.gov.uk/government/publications/code-of-conduct-for-data-driven-health-and-care-technology/initial-code-of-conduct-for-data-driven-health-and-care-technology.

[28] D.C. Higgins and C. Johner, "Validation of Artificial Intelligence Containing Products Across the Regulated Healthcare Industries," *Therapeutic Innovation & Regulatory Science* 57, no. 4 (July 2023): 797–809, https://doi.org/10.1007/s43441-023-00530-4.

[29] A. Homeyer, C. Geißler, L. O. Schwen, et al., "Recommendations on Compiling Test Datasets for Evaluating Artificial Intelligence Solutions in Pathology," *Modern Pathology* 35, no. 12 (December 2022): 1759–1769, https://doi.org/10.1038/s41379-022-01147-y.

[30] Keyvan Rahmani et al., "Assessing the Effects of Data Drift on the Performance of Machine Learning Models Used in Clinical Sepsis Prediction," *International Journal of Medical Informatics* 173 (May 2023): 104930, https://doi.org/10.1016/j.ijmedinf.2022.104930.

[31] Nineta Polemi, Isabel Praça, Kitty Kioskli, and Adrien Bécue, "Challenges and Efforts in Managing AI Trustworthiness Risks: A State of Knowledge," *Frontiers in Big Data* 7 (May 9, 2024): 1381163, https://doi.org/10.3389/fdata.2024.1381163.

[32] James L. Cross, Michael A. Choma, and John A. Onofrey, "Bias in Medical AI: Implications for Clinical Decision-Making," *PLOS Digital Health* 3, no. 11 (November 7, 2024): e0000651, https://doi.org/10.1371/journal.pdig.0000651.

[33] Joohon Sung and John L. Hopper, "Co-evolution of Epidemiology and Artificial Intelligence: Challenges and Opportunities," *International Journal of Epidemiology* 52, no. 4 (August 2023): 969–973, https://doi.org/10.1093/ije/dyad089.

[34] Joanna Reniewicz et al., "Artificial Intelligence/Machine-Learning Tool for Post-Market Surveillance of *In Vitro* Diagnostic Assays," *New Biotechnology* 78 (2023): 37–45, https://doi.org/10.1016/j.nbt.2023.11.005.

actors seek to manipulate outputs or corrupt data feeds.[35] Such threats pose serious risks to clinical decisions and patient care. We therefore recommend that the FDA require robust threat modeling and security protocols. Clear guidelines on encryption, access controls, and periodic security testing will help maintain device reliability. By emphasizing both external validation and cybersecurity planning, the FDA can ensure that adaptive AI remains effective, accountable, and aligned with the highest standards of patient care.

Building AI Fail-Safes and Contingency Plans

Given the inherent uncertainties in AI model behavior, even well-validated devices can fail in unexpected ways. Manufacturers should therefore incorporate predefined fail-safe mechanisms and contingency plans into their regulatory submissions, detailing how the device will respond when encountering unfamiliar data, operating conditions, or other indicators of potential malfunction. These measures may include automatically reverting to manual operation, triggering alerts when anomalies are detected, or requiring clinician confirmation for any high-risk outputs.[36] For example, an AI-assisted radiology tool should give radiologists the ability to override automated interpretations if the system encounters inputs outside its training parameters. Similarly, AI-driven devices for dosing or treatment management must include built-in protections to prevent erroneous recommendations if the model's confidence or validity cannot be assured.

Post-Market Surveillance and Algorithmovigilance

Although the FDA's draft guidance on AI/ML-enabled devices addresses post-market management strategies, clearer expectations for continuous performance tracking are needed to ensure ongoing accuracy and reliability in real-world settings. AI models can drift as a result of changing data distributions, new clinical practices, or inherent algorithmic evolution. These shifts can reduce performance, create unanticipated biases, and jeopardize patient safety. To address these risks, the FDA should require structured surveillance mechanisms that include standardized reporting of malfunctions, anomalies, and clinically significant performance changes.[37] Periodic revalidation of device sensitivity and specificity—supported by real-world outcomes data—will help confirm that updates and software modifications do not adversely affect patient care.[38] Regular performance reporting, including timely notification of deviations, can strengthen Agency oversight and allow for proactive intervention when emerging safety signals appear.

---

[35] International Hospital Federation, "Artificial Intelligence and Cybersecurity in Healthcare," *News & Insights*, October 3, 2023, https://ihf-fih.org/news-insights/artificial-intelligence-and-cybersecurity-in-healthcare/.

[36] Institute for Healthcare Improvement, *Patient Safety and Artificial Intelligence: Opportunities and Challenges for Care Delivery*, May 2024, https://www.ihi.org/sites/default/files/2024-05/PATIEN~1_1.PDF.

[37] Snigdha Santra et al., "Navigating Regulatory and Policy Challenges for AI-Enabled Combination Devices," *Frontiers in Medical Technology* 6 (November 27, 2024), https://doi.org/10.3389/fmedt.2024.1473350.

[38] Karen Zhou and Ginny Gattinger, "The Evolving Regulatory Paradigm of AI in MedTech: A Review of Perspectives and Where We Are Today," *Therapeutic Innovation & Regulatory Science* 58, no. 3 (March 25, 2024): 456–464, https://doi.org/10.1007/s43441-024-00628-3.

A structured system of "algorithmovigilance," modeled after pharmacovigilance, would further enhance the monitoring and accountability of AI-enabled medical devices.[39,40] This framework should incorporate registries or post-market studies for devices operating in high-risk clinical environments, enabling regulators to track real-world performance data and ensure that AI systems remain within acceptable parameters over time. Automated monitoring tools capable of detecting deviations in performance metrics can supplement these efforts by alerting both the Agency and manufacturers when a model strays beyond validated thresholds. Such real-time monitoring would ensure that sponsors address emergent issues promptly and that any substantial performance degradations are promptly corrected.

The FDA should encourage the use of post-market performance reports summarizing both positive and negative real-world findings, including trends in predictive accuracy, reliability, or patient outcomes. Enhanced transparency in these reports would foster trust by allowing regulators to verify sustained compliance with premarket claims and enabling broader awareness of how the model behaves over time. Collaboration could be advanced further through the creation of standardized surveillance systems or registries to track AI performance across sponsors and health care settings. Such a collective approach would allow stakeholders to compare notes on emerging issues, adopt best practices, and accelerate improvements.

While these mechanisms focus primarily on device performance and safety, AI can also bolster pharmacovigilance by analyzing comprehensive data sources—ranging from electronic health records to social media—to detect emerging safety issues more quickly than conventional methods.[41] To ensure such systems deliver reliable alerts, the FDA should require manufacturers to validate the sensitivity and specificity of the AI algorithms used for adverse-event detection. Human oversight remains indispensable: flagged signals should be reviewed and contextualized by clinical experts before any regulatory action is taken, preserving the balance between early detection and accurate interpretation.[42] By formally embedding algorithmovigilance requirements into AI/ML guidance, the FDA would both strengthen patient protections and reinforce public confidence in the evolving capabilities of AI-enabled medical devices.

---

[39] Peter J. Embi, "Algorithmovigilance—Advancing Methods to Analyze and Monitor Artificial Intelligence–Driven Health Care for Effectiveness and Equity," *JAMA Network Open* 4, no. 4 (April 15, 2021): e214622, https://doi.org/10.1001/jamanetworkopen.2021.4622.

[40] A. Balendran et al., "Algorithmovigilance: Lessons from Pharmacovigilance," *NPJ Digital Medicine* 7 (October 2, 2024): 270, https://doi.org/10.1038/s41746-024-01237-y.

[41] Ania Syrowatka et al., "Key Use Cases for Artificial Intelligence to Reduce the Frequency of Adverse Drug Events: A Scoping Review," *The Lancet Digital Health* 4, no. 2 (February 2022): e137–e148, https://doi.org/10.1016/S2589-7500(21)00229-6.

[42] Thomas P. Quinn et al., "Trust and Medical AI: The Challenges We Face and the Expertise Needed to Overcome Them," *Journal of the American Medical Informatics Association* 28, no. 4 (December 19, 2020): 890–894, https://doi.org/10.1093/jamia/ocaa268.

Performance Consistency and Data Representativeness

Ensuring that AI-driven devices perform consistently in different clinical contexts and across a broad array of patient populations is vital to maintaining patient safety and reliable health outcomes; it also sets the foundation for transparent regulatory submissions. Although AI holds promise in accelerating clinical decision-making, inconsistencies in real-world performance have been observed. Early successes in medical imaging, for example, did not always transfer well across facilities that used different devices, protocols, or patient characteristics, highlighting the need for continuous validation.[43] When datasets do not fully capture the range of clinical realities a device might encounter, performance gaps can emerge that may compromise real-world applicability. As such, sponsors should disclose essential information about data sources, collection methodologies, processing protocols, and any known limitations. If certain populations, clinical scenarios, or geographic areas are insufficiently represented, developers should explain how these gaps may affect model reliability and outline the steps they plan to take—such as additional data gathering, synthetic data augmentation, or post-processing techniques—to achieve balanced validation.[44,45,46,47]

Addressing Coverage Gaps, Subgroup Variation, and Corrective Measures

To establish clear benchmarks, the FDA should set explicit requirements for dataset representativeness, validation transparency, and the assessment of consistent performance across clinically relevant subgroups. In alignment with Good Machine Learning Practice (GMLP) Principle 3, manufacturers should break down key patient characteristics—such as age, clinical presentation, and disease variations—and illustrate how those characteristics reflect the scope of the device's intended use.[48] If significant coverage gaps arise, developers should conduct a risk assessment and propose mitigation strategies (for example, by introducing expanded datasets or refining algorithms) to ensure reliable outcomes for relevant populations. If significant coverage gaps arise, developers should conduct a risk assessment and propose mitigation

---

[43] World Health Organization. *Ethics and Governance of Artificial Intelligence for Health: Guidance on Large Multi-Modal Models.* Geneva: World Health Organization, 2024. https://iris.who.int/bitstream/handle/10665/375579/9789240084759-eng.pdf.

[44] Mohamed Khalifa and Mona Albadawy, "AI in Diagnostic Imaging: Revolutionising Accuracy and Efficiency," *Computer Methods and Programs in Biomedicine Update* 5 (2024): 100146, https://doi.org/10.1016/j.cmpbup.2024.100146.

[45] Anmol Arora et al., "The Value of Standards for Health Datasets in Artificial Intelligence-Based Applications," *Nature Medicine* 29, no. 11 (October 26, 2023): 2929–2938, https://doi.org/10.1038/s41591-023-02608-w.

[46] Michael H. Chin et al., "Guiding Principles to Address the Impact of Algorithm Bias on Racial and Ethnic Disparities in Health and Health Care," *JAMA Network Open* 6, no. 12 (2023): e2345050, https://doi.org/10.1001/jamanetworkopen.2023.45050.

[47] Richard J. Chen et al., "Algorithm Fairness in Artificial Intelligence for Medicine and Healthcare," *Nature Biomedical Engineering* 7, no. 6 (June 2023): 719–742, https://doi.org/10.1038/s41551-023-01056-8.

[48] U.S. Food and Drug Administration, *Good Machine Learning Practice for Medical Device Development: Guiding Principles*, last modified October 27, 2021, https://www.fda.gov/medical-devices/software-medical-device-samd/good-machine-learning-practice-medical-device-development-guiding-principles.

strategies (for example, by introducing expanded datasets or refining algorithms) to ensure reliable outcomes for relevant populations. If sponsors observe substantial differences in how a model performs for any subgroup, they should take corrective measures before seeking regulatory approval; where adjustments remain infeasible, labeling should fully disclose these limitations to clinicians and patients. To support evaluations of subgroup consistency, the FDA may wish to draw upon established standards and research in on performance variability within medical AI, providing guidance that allows flexibility in how sponsors test their models yet demands consistency in performance measurement.[49,50] Such clarity would help confirm that AI-enabled devices provide similar levels of clinical benefit and reliability for all identified subgroups of interest.[51]

Independent Validation and Generalizability

Even when in-house validation data appear promising, real-world performance may degrade if a model is never tested outside its initial development environment. Variations in patient demographics, clinical protocols, and data-capture processes can cause well-tuned algorithms to underperform or behave unpredictably in new settings. Independent validation studies offer a critical means of detecting site-specific variations, overfitting, or other hidden limitations that are unlikely to emerge in controlled or narrow data samples. These studies may involve testing on patient populations from different regions, time periods, or clinical institutions to better approximate the full range of real-world conditions the device is likely to encounter. If truly independent datasets are infeasible or unavailable, sponsors must thoroughly justify any reliance on internal cross-validation and provide detailed documentation of how these internal methods simulate real-world variability. By demonstrating that performance holds up across multiple contexts, sponsors can strengthen the credibility of their model's effectiveness and ensure that safety and accuracy claims are valid for a broad segment of potential end users.

Synthetic Data

To enhance generalizability and strengthen model development under constrained data conditions, the FDA should actively encourage the responsible use of data augmentation techniques, including the generation and application of synthetic data. Synthetic data—when carefully designed and validated—can serve as a valuable supplement to real-world datasets, enabling AI models to train on scenarios that may be logistically difficult to capture through traditional clinical data collection. This is particularly important for rare conditions, underdiagnosed populations, or edge cases

[49] Siân Carey, Allan Pang, and Marc de Kamps, "Fairness in AI for Healthcare," Future Healthcare Journal 11, no. 3 (September 19, 2024): 100177, https://doi.org/10.1016/j.fhj.2024.100177.

[50] Matthew G. Hanna et al., "Ethical and Bias Considerations in Artificial Intelligence/Machine Learning," *Modern Pathology* 38, no. 3 (March 2025): 100686, https://doi.org/10.1016/j.modpat.2024.100686.

[51] Anmol Arora et al., "The Value of Standards for Health Datasets in Artificial Intelligence-Based Applications," Nature Medicine 29, no. 11 (October 26, 2023): 2929–2938, https://doi.org/10.1038/s41591-023-02608-w.

that are essential for robust model performance but may not appear with sufficient frequency in conventional datasets.[52,53]

However, the generation and use of synthetic data must be subject to rigorous oversight. Sponsors should be required to demonstrate that synthetic datasets are derived from statistically sound methods, reflect realistic clinical features, and do not introduce artifacts or distortions that could mislead model training. Synthetic data should not be treated as interchangeable with empirical data unless validated against actual clinical outcomes. All validation exercises that incorporate synthetic data should be held to the same evidentiary standards as those using real-world data, including thorough assessments of sensitivity, specificity, and predictive value. Sponsors must also disclose the provenance of any synthetic data, describe the algorithms used to generate it, and outline steps taken to ensure that it preserves key clinical correlations without amplifying noise or introducing spurious patterns.

Furthermore, when synthetic data are used in support of a marketing submission, the FDA should require a detailed impact assessment evaluating how synthetic inputs may influence the model's predictions in real-world deployment. Sponsors should be prepared to explain the rationale for including synthetic data, identify which components of the model were trained or augmented using such data, and describe any additional safeguards employed to monitor post-deployment performance. By establishing clear expectations for the responsible use and validation of synthetic data, the FDA can support innovation while ensuring that AI-enabled devices remain accurate, clinically meaningful, and safe for real-world use.

Validation and Verification of AI Models

The NHC recommends that the FDA adopt explicit, rigorous validation and verification mandates for AI models intended to inform regulatory or clinical decision-making. Because these technologies can directly affect patient outcomes and public trust, the FDA should require clearly defined performance standards, carefully chosen validation methodologies, and ongoing verification procedures.

First, sponsors should provide detailed records of how data were collected, including annotation protocols, quality-control processes, and any assumptions regarding data labeling or preprocessing. This foundational documentation helps reviewers determine whether the input data—and the procedures by which it was gathered—adequately capture real-world conditions and faithfully represent the patient populations for which the device is intended.

Second, the FDA should require sponsors to report and justify a range of performance metrics, such as sensitivity, specificity, positive predictive value, negative predictive value, precision, recall, and the area under the receiver operating characteristic (ROC)

---

[52] Panteha Alipour and Erika E. Gallegos, "Leveraging Generative AI Synthetic and Social Media Data for Content Generalizability to Overcome Data Constraints in Vision Deep Learning," *Artificial Intelligence Review* 58, no. 5 (February 2025), https://doi.org/10.1007/s10462-025-11137-6.

[53] Vasileios C. Pezoulas et al., "Synthetic Data Generation Methods in Healthcare: A Review on Open-Source Tools and Methods," *Computational and Structural Biotechnology Journal* 23 (December 2024): 2892–2910, https://doi.org/10.1016/j.csbj.2024.07.005.

curve. These different metrics each offer unique insights into how consistently and accurately an AI model detects or predicts certain outcomes. By segmenting results by critical subgroups and clearly explaining why certain metrics are most relevant for a device's intended use, sponsors can demonstrate that their models meet appropriate clinical thresholds in practice.[54,55] If significant differences appear across subgroups, sponsors should address them prior to regulatory approval or disclose any remaining limitations in the device labeling.

Third, sponsors must show that their AI models maintain robustness across multiple real-world conditions. Validating against third-party datasets, or datasets spanning comprehensive geographic and institutional settings, helps confirm whether performance claims extend beyond controlled development environments. Temporal validation—tracking how outputs evolve as clinical practices or population traits change—is equally important to detect performance drift. By requiring sponsors to periodically re-test their models or incorporate real-world data, the FDA can ensure that AI-enabled devices retain their claimed accuracy and applicability over time.[56,57,58]

Finally, comprehensive documentation is essential for reproducibility and public confidence. The FDA should mandate that sponsors submit experimental designs, statistical testing strategies, performance thresholds, and any results from third-party evaluators. This should also include descriptions of how sponsors capture and analyze disparities in measured metrics, along with any root-cause analyses and remediation steps. Such transparency fosters independent review, clarifies model limitations, and helps uphold trust in AI-driven health solutions. By reinforcing these practices—encompassing dataset disclosure, careful validation, real-world generalizability checks, assessments of performance consistency across patient subgroups, and thorough documentation—the FDA can strengthen performance consistency in AI-enabled devices and safeguard reliable care throughout each model's lifecycle.[59]

---

[54] Sebastian Vollmer et al., "Machine Learning and Artificial Intelligence Research for Patient Benefit: 20 Critical Questions on Transparency, Replicability, Ethics, and Effectiveness," *BMJ* 368 (March 20, 2020): l6927, https://doi.org/10.1136/bmj.l6927.

[55] D. Ukalovic et al., "Prediction of Ineffectiveness of Biological Drugs Using Machine Learning and Explainable AI Methods: Data from the Austrian Biological Registry BioReg," *Arthritis Research & Therapy* 26 (2024): 44, https://doi.org/10.1186/s13075-024-03277-x.

[56] Zhaoyi Chen et al., "Applications of Artificial Intelligence in Drug Development Using Real-World Data," *Drug Discovery Today* 26, no. 5 (2021): 1256–1264, https://doi.org/10.1016/j.drudis.2020.12.013.

[57] Daehwan Ahn, Abdullah Almaatouq, Monisha Gulabani, and Kartik Hosanagar, "Impact of Model Interpretability and Outcome Feedback on Trust in AI," in *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems (CHI '24)* (New York: Association for Computing Machinery, 2024), Article 27, 1–25, https://doi.org/10.1145/3613904.3642780.

[58] Eike Petersen et al., "The Path Toward Equal Performance in Medical Machine Learning," *Patterns* 4, no. 7 (July 14, 2023): 100790, https://doi.org/10.1016/j.patter.2023.100790.

[59] Anastasiya Kiseleva, Dimitris Kotzinos, and Paul De Hert, "Transparency of AI in Healthcare as a Multilayered System of Accountabilities: Between Legal Requirements and Technical Limitations," *Frontiers in Artificial Intelligence* 5 (May 30, 2022): 879603, https://doi.org/10.3389/frai.2022.879603.

Human Oversight and Model Interpretability

The NHC emphasizes that human oversight and model interpretability are foundational to the responsible integration of AI/ML-enabled medical devices. As these technologies increasingly inform clinical decisions, their outputs must be clear, explainable, and meaningful to clinicians, regulators, and patients.[60] Ensuring that AI systems are understandable supports confidence in their use, facilitates safe clinical integration, and aligns with broader goals of patient-centered care.

To that end, the FDA should encourage sponsors to incorporate design features that promote interpretability for the range of intended users. Sponsors should be prepared to describe the logic underpinning model predictions, including key input variables, model assumptions, and decision thresholds. For more complex models, such as those based on deep learning, additional tools—such as visualizations or feature importance analyses—may help clarify the model's rationale. The FDA should continue to provide flexibility in how interpretability is achieved, recognizing that requirements may vary depending on a device's intended use, complexity, and risk level.

Human oversight plays a complementary role by ensuring that AI-generated recommendations are used appropriately and remain subject to expert judgment. For higher-risk applications, sponsors should outline processes for clinical review of AI outputs, including the conditions under which additional evaluation may be warranted. Rather than imposing prescriptive requirements, the FDA should emphasize the importance of context-specific oversight approaches that enable safe use of AI systems across a range of care environments.

Interpretability and oversight are closely linked to usability. AI-enabled tools must integrate effectively into clinical workflows without introducing unnecessary complexity or cognitive burden. The FDA should encourage human factors studies to ensure that AI outputs—such as alerts, probability scores, or visual annotations—are presented clearly and support effective decision-making.[61,62] These studies should include representative users across diverse clinical settings and should confirm that the device can be used safely and effectively in practice.[63]

Clear, accessible communication to patients is also essential. Labeling materials should explain the function of the AI component, its intended use, and any known limitations or conditions under which clinical input is recommended.[64] When AI models are adaptive,

---

[60] Tilala et al., "Ethical Considerations in AI and Machine Learning," e62443.

[61] Shuroug A. Alowais et al., "Revolutionizing Healthcare: The Role of Artificial Intelligence in Clinical Practice," *BMC Medical Education* 23 (September 22, 2023): Article 689, https://doi.org/10.1186/s12909-023-04873-5.

[62] Schinkel, van der Poll, and Wiersinga, "Artificial Intelligence for Early Sepsis Detection," 854.

[63] Ulugbek Vahobjon Ugli Ismatullaev and Sang-Ho Kim, "Review of the Factors Affecting Acceptance of AI-Infused Systems," *Human Factors Review* 66, no. 1 (January 2024): 126–144, https://doi.org/10.1177/00187208211064707.

[64] U.S. Food and Drug Administration, *Good Machine Learning Practice for Medical Device Development: Guiding Principles.*

sponsors should describe how updates occur, how users will be informed of changes, and what safeguards are in place to ensure continued reliability.[65,66] By promoting interpretability, thoughtful oversight, and strong usability practices, the FDA can ensure that AI/ML-enabled medical devices operate as valuable tools that support clinical expertise and advance patient care.[67] These measures will reinforce user trust and promote responsible adoption of AI in health care settings.[68,69,70,71,72]

Addressing Cybersecurity and Data Integrity

The NHC recommends that the FDA establish explicit requirements to address cybersecurity risks and protect data integrity in AI/ML-enabled medical devices. These technologies often rely on cloud-based infrastructure, networked environments, and integration with electronic health record systems—introducing potential vulnerabilities to data breaches, unauthorized access, and manipulation of algorithmic functions. Because AI-enabled devices may interact with live clinical systems and continuously learn or update over time, ensuring secure and trustworthy system operation is essential to protecting patient safety and preserving public confidence.

To that end, the FDA should require sponsors to conduct comprehensive cybersecurity risk assessments throughout the AI device lifecycle. These assessments should identify potential vulnerabilities in data storage, transmission, and processing—including points of entry for external attacks and risks associated with interconnected systems. Sponsors must also describe mitigation strategies, including the implementation of advanced encryption, access controls, multi-factor authentication, secure software development practices, and anomaly detection systems designed to flag unusual behaviors or unauthorized modifications to model outputs.

---

[65] Mark Steyvers and Aakriti Kumar, “Three Challenges for AI-Assisted Decision-Making,” *Perspectives on Psychological Science* 19, no. 5 (July 13, 2023): 722–734, https://doi.org/10.1177/17456916231181102.

[66] Jacqueline G. You et al., “Clinical Trials Informed Framework for Real World Clinical Implementation and Deployment of Artificial Intelligence Applications,” *NPJ Digital Medicine* 8 (February 17, 2025): Article 107, https://doi.org/10.1038/s41746-025-01272-1.

[67] Vijaytha Muralidharan, Boluwatife Adeleye Adewale, Caroline J. Huang, et al. “A Scoping Review of Reporting Gaps in FDA-Approved AI Medical Devices.” *npj Digital Medicine* 7 (2024): 273. https://doi.org/10.1038/s41746-024-01270-x.

[68] Ismatullaev and Kim, “Review of the Factors Affecting Acceptance of AI-Infused Systems,” 126–144.

[69] Quinn et al., “Trust and Medical AI,” 890–894.

[70] Quinn et al., “Trust and Medical AI,” 890–894.

[71] Steven M. Williamson and Victor Prybutok, “Balancing Privacy and Progress: A Review of Privacy Challenges, Systemic Oversight, and Patient Perceptions in AI-Driven Healthcare,” *Applied Sciences* 14, no. 2 (2024): 675, https://doi.org/10.3390/app14020675.

[72] Sabale Mrunal M. et al., “Maintaining Data Safety and Accuracy Through Data Integrity (DI): A Comprehensive Review,” *Research Journal of Pharmacy and Technology* 17, no. 5 (2024): 2431–2440, https://doi.org/10.52711/0974-360X.2024.00381.

Robust documentation should accompany each device submission, outlining sponsors' protocols for continuous security monitoring, breach detection, and response plans in the event of a security incident. The FDA should also require sponsors to maintain detailed audit trails capturing key system events—such as data inputs, model versioning, user interactions, update deployments, and overrides of AI recommendations. These logs are essential for enabling post-incident investigations and for maintaining accountability throughout the product lifecycle. Sponsors must demonstrate their ability to manage and preserve audit logs securely, ensuring the traceability and reproducibility of system behavior across all deployment contexts.

In addition, the FDA should mandate independent verification and validation of cybersecurity and data integrity safeguards at multiple stages, including during premarket review and post-market surveillance. External validation of model performance and system resilience—particularly when tested on datasets and infrastructure not used during development—can help uncover potential vulnerabilities and reduce the risk of undetected model tampering or corruption. These safeguards are essential to ensure that AI-enabled devices remain reliable, accurate, and secure in the dynamic environments in which they operate.

By requiring strong cybersecurity practices and comprehensive data integrity protections, the FDA can help ensure that AI/ML-enabled medical devices maintain operational reliability, safeguard patient information, and uphold trust in AI-driven health technologies.

Conclusion

The NHC appreciates the opportunity to contribute to this critical regulatory initiative and looks forward to continued collaboration with the FDA to ensure that AI-driven drug and biologic development remains patient-centered, ethical, and scientifically rigorous. Thank you for your consideration and for your dedication to advancing safe and effective health innovation. For additional dialogue, please contact Kimberly Beer, Senior Vice President of Policy and External Affairs, at kbeer@nhcouncil.org or Shion Chang, Senior Director of Policy and Regulatory Affairs, at schang@nhcouncil.org.

Sincerely,

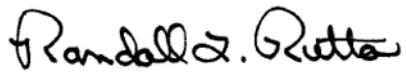

Randall L. Rutta
Chief Executive Officer