NATIONAL HEALTH COUNCIL

April 7, 2025

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

RE: Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products [FDA-2024-D-4689]

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National Health Council (NHC) appreciates the opportunity to submit comments on the Food and Drug Administration (FDA) draft guidance, "Considerations for the Use of Artificial Intelligence to Support Regulatory Decision-Making for Drug and Biological Products." The integration of artificial intelligence (AI) and machine learning (ML) technologies into regulatory decision-making holds tremendous promise for advancing drug and biologic development, streamlining regulatory processes, and ultimately improving patient care and public health outcomes.[1,2] However, as these innovative technologies are increasingly deployed within the regulatory framework, it is essential to address the inherent complexities and challenges they present, particularly regarding safety, effectiveness, transparency, and ensuring that regulatory decisions are grounded in the best available science to support meaningful outcomes for all patient populations.[3,4]

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.

---

[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] Fahimeh Mirakhori and Sarfaraz K. Niazi, "Harnessing the AI/ML in Drug and Biological Products Discovery and Development: The Regulatory Perspective," *Pharmaceuticals* 18, no. 1 (2025): 47, https://doi.org/10.3390/ph18010047.

[3] 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.

[4] 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.

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.

With this broad, multi-stakeholder perspective, we stress that AI regulations must prioritize patient safety, consistent clinical performance across intended-use populations, and public trust. The NHC has published Principles on Health AI that emphasize patient benefit, scientific integrity, transparency, privacy, and accountability. 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 proactively developing a regulatory framework to guide the integration of AI applications in regulatory processes and appreciates the opportunity to provide input on this significant initiative. 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 and considerations regarding critical issues surrounding the use of AI in regulatory decision-making for drug and biologic products.[5] Specifically, our recommendations address:

- Rigorous validation and verification of AI models
- Transparency and representativeness in data collection
- Robust human oversight and interpretability
- Comprehensive post-market surveillance
- Stringent cybersecurity and data integrity protections
- International regulatory harmonization
- Ethical and patient-centered considerations

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.

[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/.

AI in Drug Discovery and Preclinical Development

The integration of AI and ML into drug and biologic development presents transformative opportunities while introducing critical regulatory and ethical challenges. These technologies have the potential to revolutionize drug discovery, accelerate clinical trials, enhance pharmacovigilance, and improve manufacturing processes—ultimately leading to more efficient and effective treatments for patients.[6] However, the use of AI in these early phases raises questions about how to ensure that emerging tools contribute to a reliable, transparent, and patient-centered regulatory system.[7,8]

The NHC acknowledges that many applications of AI in early-stage discovery fall outside the FDA's direct regulatory scope. At the same time, as insights generated by AI increasingly shape the selection of drug targets, biomarkers, or dosing strategies that inform regulatory submissions, it is essential that any AI-derived outputs used to support regulatory decision-making meet appropriate standards of empirical validation and scientific transparency.[9]

AI can significantly enhance drug discovery by analyzing vast datasets to detect patterns often missed by traditional approaches, thus accelerating early-stage development.[10] By leveraging ML algorithms, AI can facilitate hypothesis generation, structure-activity relationship modeling, and biomarker discovery, contributing to more efficient and targeted treatments.[11] Nonetheless, the reliability and generalizability of AI-driven insights depend heavily on the quality of training data and the rigor of validation methodologies.[12]

To this end, the NHC encourages the FDA to establish clear expectations for how sponsors validate and document AI-driven findings that are submitted as part of the

---

[6] Sarfaraz K. Niazi, "The Coming of Age of AI/ML in Drug Discovery, Development, Clinical Testing, and Manufacturing: The FDA Perspectives," *Drug Design, Development and Therapy* 17 (September 2023): 2691–2725, https://doi.org/10.2147/DDDT.S424991.

[7] E. Zaidan and I. A. Ibrahim, "AI Governance in a Complex and Rapidly Changing Regulatory Landscape: A Global Perspective," *Humanities and Social Sciences Communications* 11 (2024): 1121, https://doi.org/10.1057/s41599-024-03560-x.

[8] M. Khair ElZarrad, Aaron Y. Lee, Rose Purcell, and Scott J. Steele, "Advancing an Agile Regulatory Ecosystem to Respond to the Rapid Development of Innovative Technologies," *Clinical and Translational Science* (March 23, 2022), https://doi.org/10.1111/cts.13267.

[9] Mirakhori and Niazi, "Harnessing the AI/ML in Drug and Biological Products Discovery and Development," 47.

[10] Vinay Kumar and Kunal Roy, "Embracing the Changes and Challenges with Modern Early Drug Discovery," *Expert Opinion on Drug Discovery* (March 19, 2025), https://doi.org/10.1080/17460441.2025.2481259.

[11] Shruti Singh, Rajesh Kumar, Shuvasree Payra, and Sunil K. Singh, "Artificial Intelligence and Machine Learning in Pharmacological Research: Bridging the Gap Between Data and Drug Discovery," *Cureus* 15, no. 8 (August 30, 2023): e44359, https://doi.org/10.7759/cureus.44359.

[12] K.K. Mak, Y.H. Wong, and M.R. Pichika, "Artificial Intelligence in Drug Discovery and Development," in *Drug Discovery and Evaluation: Safety and Pharmacokinetic Assays*, ed. F.J. Hock and M.K. Pugsley (Cham: Springer, 2024), https://doi.org/10.1007/978-3-031-35529-5_92.

regulatory dossier. These AI models should serve as complementary tools to traditional preclinical methods, not replacements for established scientific validation.[13] Sponsors should be prepared to demonstrate the biological and chemical plausibility of AI-generated findings and to submit documentation detailing data provenance, feature selection, validation methods, and performance metrics to allow greater regulatory scrutiny and help ensure clinically meaningful outcomes.[14]

In addition, sponsors should ensure that the datasets used to develop AI models reflect the intended patient population and disease variations.[15,16] Failure to do so risks overlooking meaningful variability and could reinforce historical limitations in drug development. While the NHC is not advocating for direct oversight of discovery tools, we do emphasize the importance of ensuring that any AI-informed decisions that reach the regulatory interface are supported by reproducible, well-validated evidence. An approach by FDA that preserves flexibility and innovation in early discovery while promoting transparency and scientific integrity in regulatory submissions ultimately serves the best interests of patients and reinforces confidence in AI-enabled innovation.

Validation and Verification of AI Models

The NHC recommends that the FDA adopt explicit, rigorous validation and verification requirements for AI models intended for regulatory decision-making in drug and biologic development. Given the critical role these technologies play in influencing decisions that directly affect patient safety, therapeutic effectiveness, and public trust, it is imperative that AI-driven outputs undergo thorough assessments of accuracy, robustness, and reproducibility. To achieve this, the NHC encourages the FDA to set clear, context-specific performance standards and validation methodologies, along with ongoing verification procedures, to ensure continuous reliability and effectiveness.

To establish a sound basis for trust in AI models, the NHC encourages the FDA to require sponsors to define, justify, and rigorously apply performance metrics tailored to each model's intended use and regulatory context. Depending on the application, critical metrics might include sensitivity and specificity, especially when predicting adverse events or determining patient stratification in clinical trials. Sponsors should also measure positive and negative predictive values to confirm that AI-generated predictions remain reliable in practical clinical scenarios.[17] Additionally, metrics like the

[13] Xinyue Hao, Emrah Demir, and Daniel Eyers, "Exploring Collaborative Decision-Making: A Quasi-Experimental Study of Human and Generative AI Interaction," *Technology in Society* 78 (September 2024): 102662, https://doi.org/10.1016/j.techsoc.2024.102662.

[14] Filippo Pesapane et al., "The Translation of In-House Imaging AI Research into a Medical Device Ensuring Ethical and Regulatory Integrity," *European Journal of Radiology* 182 (January 2025): 111852, https://doi.org/10.1016/j.ejrad.2024.111852.

[15] Jiarui Xie, Lijun Sun, and Yaoyao Fiona Zhao, "On the Data Quality and Imbalance in Machine Learning-Based Design and Manufacturing—A Systematic Review," *Engineering* 45 (February 2025): 105–131, https://doi.org/10.1016/j.eng.2024.04.024.

[16] Mirakhori and Niazi, "Harnessing the AI/ML in Drug and Biological Products Discovery and Development," 47.

area under the receiver operating characteristic curve (ROC-AUC) can help evaluate how effectively a model distinguishes between outcomes over a range of decision thresholds, while precision and recall illustrate the balance between correct identifications and potential misclassifications.[18] Whichever metrics are chosen, sponsors must thoroughly justify their selection, clearly demonstrating relevance and suitability in the given context.

The FDA should also emphasize validation approaches that assess the robustness and generalizability of AI-driven predictions. Sponsors must show that their AI models consistently deliver accurate, reproducible results when validated against independent datasets different from those used in initial development. These datasets should encompass multiple clinical trials or observational studies, reflecting diverse patient demographics, disease presentations, and care settings. Validation should also occur across various geographic locations to confirm real-world reliability.[19] In addition, temporal validation—assessing model performance over time or across evolving clinical conditions—should be required to ensure AI models maintain predictive accuracy as medical practices and patient populations change.[20]

Given AI models can perpetuate patterns present in their training data, potentially skewing outcomes, it is essential that the FDA explicitly require comprehensive evaluations of how models perform across clinically relevant subgroups.[21] Sponsors must identify any differential performance across patient subgroups defined by demographics or other relevant characteristics, and the FDA should mandate transparent subgroup reporting that details any disparities in model accuracy, sensitivity, or specificity. When substantial performance gaps are identified, sponsors should clearly document their mitigation strategies, explaining how they plan to correct or manage these disparities to protect patient outcomes. Furthermore, FDA guidance should include requirements for clear labeling and disclosure of any remaining performance limitations or biases, ensuring clinicians and patients are aware of potential constraints in real-world contexts.

Finally, to support transparency, reproducibility, and traceability, the FDA should require thorough documentation of AI model validation and verification. Sponsors should submit

---

[17] 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.

[18] 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.

[19] 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.

[20] 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.

[21] 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.

detailed records on dataset provenance, data collection methods, annotation protocols, and quality control procedures. Likewise, all aspects of model architecture—feature selection, hyperparameter tuning, and assumptions made during development—should be clearly explained. Validation protocols must be meticulously described, outlining experimental designs, statistical testing approaches, performance criteria, and the results from validation exercises (including those by third-party evaluators, if applicable). This level of documentation will enable robust independent review, facilitate the reproducibility of results, and promote greater regulatory transparency.

Transparency and Representativeness in Data Collection

The NHC emphasizes that the successful integration of AI models into regulatory decision-making for drugs and biologics relies on high-quality, transparent, and representative data. To that end, the FDA should require sponsors to disclose comprehensive information about their datasets, including detailed descriptions of data sources, collection and processing methodologies, and characteristics of the patient populations represented. Such transparency is critical not only for independent review and reproducibility but also for identifying potential imbalances or limitations that may affect reliability and consistency of model outputs across different patient populations.

To ensure AI-driven regulatory decisions are both fair and generalizable, the FDA should require that training datasets reflect the full spectrum of patient populations for which a drug or biologic is intended. Sponsors must demonstrate adequate demographic and clinical variability—such as age cohorts, disease stages, and other subgroup characteristics—within their datasets.[22] If gaps in representation exist, sponsors should clearly explain these shortcomings and outline the implications for model performance. In such cases, the FDA is encouraged to prompt sponsors to fill significant data gaps or adopt strategies to address these limitations during validation.

Beyond model development, AI and ML are increasingly employed in clinical trial design, patient recruitment, and risk stratification. By analyzing electronic health records, genetic profiles, and real-world data, AI can help identify patient subgroups most likely to benefit from investigational treatments, thereby improving trial efficiency and success rates. However, the use of AI in patient selection raises concerns about inaccuracies, limitations in data quality, and uneven performance across different populations. Sponsors should therefore be required to disclose the datasets used to train any AI-based recruitment tools, including key demographic and clinical attributes, to avoid inadvertently excluding certain populations or perpetuating existing disparities.[23,24] The FDA should also establish clear guidelines for validating AI-driven

[22] 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.

[23] 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.

[24] 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.

selection criteria across varied patient populations and clinical settings to confirm consistent performance and broad applicability.

As AI takes on a larger role in adaptive trial designs, the FDA should develop guidance on incorporating AI-generated insights into protocol amendments, interim analyses, and dose-optimization strategies. Sponsors must demonstrate that these AI-driven modifications enhance study efficiency without compromising patient safety or data integrity, ensuring that real-time AI interventions strengthen—rather than undermine—the rigor of clinical research.

The FDA should further encourage the creation of publicly available, well-curated reference datasets to support robust AI model development and validation. Such standardized benchmarks would help sponsors calibrate and validate their AI models more consistently, improving comparability across regulatory submissions. To achieve this, the FDA could collaborate with other regulatory bodies, research institutions, industry consortia, and patient organizations to develop data repositories representing the full range of patient populations and clinical scenarios, thereby enhancing overall transparency in regulatory decision-making.

Finally, transparency must include how sponsors address potential imbalances in their training data. Sponsors should document explicit assessments of representativeness, identifying whether specific populations or clinical circumstances are over- or insufficiently represented. The FDA should require sponsors to detail how they plan to mitigate any identified deficiencies—whether by incorporating supplemental data, refining training methodologies, or applying post-processing techniques to strengthen the reliability and applicability of model outputs. By ensuring that any inherent limitations or biases are recognized, addressed, and clearly communicated to clinicians and patients, the FDA can help sustain public trust in AI-driven regulatory processes.[25]

Human Oversight and Model Interpretability

The NHC underscores that human oversight and interpretability are foundational to ensuring that AI used in drug and biologic development aligns with clinical realities and patient needs. AI models should not operate as isolated “black boxes” but instead support human expertise, with outputs that are transparent, explainable, and actionable by clinicians and regulatory reviewers. To achieve this, the FDA should encourage sponsors to adopt oversight approaches proportional to the model’s risk level and clinical impact, ensuring flexibility while upholding accountability.

For AI systems intended to inform regulatory decisions, sponsors should document the internal logic of their models—including model architecture, input-output relationships, and the rationale behind key design choices. Where greater complexity or opacity exists (as in “black box” models), sponsors should employ supplementary techniques such as feature attribution or visualizations to support interpretability. Rather than prescribing one-size-fits-all requirements, the FDA could allow sponsors to propose oversight

[25] 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.

mechanisms suited to their model's context, provided they clearly demonstrate that users can meaningfully review and understand AI outputs.

In scenarios involving high-consequence decisions—such as trial eligibility, safety signal detection, or risk stratification—the NHC recommends that FDA guidance promote human-in-the-loop processes. These do not need to be prescriptive but should establish a baseline expectation that clinicians or regulators retain the ability to verify, contextualize, or override AI-generated recommendations when appropriate. By encouraging thoughtful oversight practices that reflect the model's role and potential impact, the FDA can safeguard public trust without unduly burdening innovation.[26,27]

Finally, to further support real-world usability, the FDA should encourage sponsors to conduct human factors studies to ensure that AI outputs are clearly presented and actionable. These variables, tailored to the intended users, can help identify where improvements are needed to support confident decision-making. Finally, labeling and communication materials should offer plain-language summaries of what the AI does, its intended use, and any limitations—reinforcing the idea that AI is a tool to enhance, not replace, human judgment.

Post-Market Surveillance and Algorithmic Drift

The NHC urges the FDA to adopt rigorous, AI-specific post-market surveillance requirements for drugs and biologics. Unlike static decision-making tools, AI/ML models can evolve over time due to shifting data distributions, changes in clinical practices, or inherent model drift. Without robust monitoring and intervention strategies, these shifts may erode performance, reduce reliability, and jeopardize patient safety. Accordingly, FDA guidance should clearly mandate continuous, comprehensive post-market monitoring to promptly detect, evaluate, and correct any degradation or drift in model performance.[28]

AI can also strengthen pharmacovigilance by analyzing real-world data (including social media and electronic health records) to detect emerging safety concerns. Automated systems can identify adverse drug reactions more rapidly than traditional reporting methods.[29] However, the accuracy of AI-driven pharmacovigilance depends on transparent data sources, robust detection algorithms, and expert review. To ensure

---

[26] 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 (April 2021): 890–894, https://doi.org/10.1093/jamia/ocaa268.

[27] 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.

[28] 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.

[29] 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.

reliability, the FDA should require sponsors to validate the sensitivity, specificity, and clinical relevance of AI-generated safety alerts. Human oversight must also be incorporated to review flagged signals and contextualize their significance before any regulatory action is taken.[30]

To further guide industry, the FDA should establish explicit expectations for employing AI in spontaneous adverse event reporting, emphasizing that AI augments rather than replaces established pharmacovigilance processes. Sponsors must document how AI-generated safety signals are prioritized, validated, and communicated to regulators and health care providers. Clear, standardized requirements will help the FDA maintain a balance between the benefits of AI-enhanced pharmacovigilance and the need for accuracy, reliability, and patient safety.

In managing algorithmic drift, the FDA should require sponsors to implement protocols for continuous real-world performance assessment post-deployment. Such protocols must define explicit performance benchmarks—based on the model's original validation—that must be maintained throughout its lifecycle. Sponsors should detail how often they will reassess the model's outputs, what data sources will be used, and what criteria will trigger corrective actions if performance deviations occur. Ongoing monitoring is critical for safeguarding patient safety and ensuring any significant performance declines are swiftly identified and addressed.

Additionally, FDA guidance must emphasize proactive strategies for handling algorithmic drift, including structured feedback mechanisms. Sponsors should set up processes to systematically review AI outputs in collaboration with clinical experts and regulatory personnel. These reviews should include validation checks against external data, direct evaluations of clinical outcomes, and the incorporation of user feedback, helping detect early signs of drift or decreased reliability. When a model deviates from established benchmarks, sponsors must investigate the cause, document findings, and enact corrections, such as retraining or recalibrating the model. Any major adjustments should be reported to the FDA and shared with stakeholders, including clinicians and patients.

The FDA should also require periodic submission of post-market performance reports detailing monitoring results, deviations in performance, corrective measures taken, and lessons learned. Periodic reporting increases transparency, allowing regulators to maintain active oversight of AI products, confirm continued compliance, and uphold patient safety. These reports should include longitudinal analyses of model performance since approval, spotlighting changes in predictive accuracy, reliability, or clinical outcomes. Such transparent oversight allows the FDA to proactively identify emerging issues, foster responsible drift management, and strengthen public confidence in AI-driven decision-making.

Finally, the FDA should encourage the creation of standardized surveillance systems or registries to track AI model performance across sponsors, health care settings, and product categories. A structured "algorithmovigilance" framework—analogous to pharmacovigilance for drugs—would enable regulators and stakeholders to monitor

[30] Quinn et al., "Trust and Medical AI," 890–894.

real-world AI performance, detect trends, and address new risks.[31] By supporting shared databases that collect outcome data across various AI implementations, the FDA can foster transparency, promote collaboration, and expedite best practices for managing AI models over time. This coordinated approach will advance understanding of real-world AI safety and effectiveness, ultimately reinforcing regulatory oversight and protecting patient interests.

AI in Drug Manufacturing and Quality Control

AI technologies are playing an increasingly prominent role in pharmaceutical manufacturing and quality control, offering the potential to enhance process automation, improve efficiency, and ensure consistency in drug production. AI-driven automation can optimize predictive maintenance, enable real-time process monitoring, and provide automated quality assessments, thus reducing manufacturing variability and reinforcing compliance with Good Manufacturing Practice standards. However, integrating AI into production settings presents new challenges, especially when it comes to validating and verifying these systems across varied environments.

To maintain high product quality and regulatory compliance, the FDA should establish clear guidance for validating AI models used in pharmaceutical manufacturing. This guidance should define requirements for model training, testing, and continuous performance monitoring within AI-based process control systems. Sponsors must submit documentation that demonstrates how AI-enabled manufacturing processes consistently produce high-quality products across different sites and batches. Additionally, the guidance should address algorithmic drift, ensuring that AI-driven decision-making does not inadvertently alter product consistency or efficacy over time. The FDA should also require manufacturers to develop contingency plans for AI model failures. These plans should outline procedures for reverting to traditional process controls if AI-driven decisions introduce unacceptable variability or errors in product quality. Given the critical nature of pharmaceutical production, robust fallback mechanisms are essential to protect product integrity, patient safety, and supply continuity.

By setting clear validation standards, mandating comprehensive documentation, and requiring contingency planning for AI-powered automation, the FDA can foster innovation in pharmaceutical manufacturing while safeguarding the reliability, quality, and compliance of drug products.

Addressing Cybersecurity and Data Integrity

The NHC recommends that the FDA establish explicit guidelines to address cybersecurity risks and safeguard data integrity in AI systems supporting regulatory decision-making for drug and biological products. Because these applications often integrate large datasets from multiple sources—frequently through cloud-based or networked platforms—they are particularly vulnerable to breaches, unauthorized access, and the intentional manipulation of AI models. To counter these threats, FDA

[31] 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.

guidance must mandate stringent measures that protect sensitive data and uphold the integrity and reliability of regulatory processes.[32]

To ensure robust cybersecurity, the FDA should require sponsors to conduct comprehensive risk assessments, identifying vulnerabilities in data storage, transmission, and processing. Sponsors must also detail mitigation strategies—such as advanced encryption, multi-factor authentication, strict access controls, and anomaly detection capable of flagging unauthorized alterations to AI outputs. In addition, documentation should describe continuous monitoring procedures and contingency plans for responding swiftly to breaches or incidents. Thorough, well-documented risk management practices will significantly bolster public trust in AI-driven systems.

Further, the FDA should require rigorous audit trails that log each step of AI-driven decision-making and system interaction. These logs should capture data inputs, model updates, any human overrides, and alerts from unusual activities, facilitating oversight and accountability. By enabling swift forensic analysis when anomalies occur, transparent audit records help safeguard against tampering or misuse. Sponsors must demonstrate their ability to securely manage and preserve these logs, ensuring traceability throughout the AI system lifecycle.

Lastly, the FDA should mandate independent verification and validation at multiple stages of the AI lifecycle. Through external reviews, sponsors must show that their models consistently generate reliable, reproducible results when applied to datasets not used during initial development. Such independent assessments mitigate risks associated with intentional or accidental data manipulation, reinforcing the scientific rigor of AI-driven regulatory decisions, safeguarding patients, and maintaining public confidence.

International Harmonization

The NHC urges the FDA to lead international harmonization efforts in AI regulation for drug and biologic development to ensure the United States remains at the forefront of regulatory science and innovation. By shaping global standards rather than merely aligning with them, the FDA can reinforce its leadership in AI oversight, ensuring that U.S. regulatory approaches set the benchmark for best practices worldwide. Harmonizing AI regulatory frameworks across major global authorities—such as the European Medicines Agency (EMA) and other international bodies—will enhance consistency, reduce regulatory uncertainty, and enable the United States to drive innovation while maintaining high standards for safety and efficacy.

The FDA should proactively engage in ongoing international collaborations to establish globally recognized best practices in AI oversight. Participation in organizations such as the International Council for Harmonisation (ICH) and the International Medical Device Regulators Forum (IMDRF) offers opportunities to shape emerging global AI regulatory

[32] 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.

standards.[33] Through these forums, the FDA can influence the development of shared definitions, validation methodologies, data governance frameworks, cybersecurity protocols, and interpretability requirements—ensuring that U.S. regulatory principles are not only reflected but positioned as the gold standard for AI-enabled drug and biologic development.

Additionally, the FDA should leverage international harmonization efforts to advance standardized datasets and reference models for AI validation. Establishing globally accepted reference datasets will enhance comparability across regulatory submissions, reduce variability, and improve validation accuracy.[34] By leading in this space, the FDA can streamline cross-border regulatory cooperation and reinforce the United States' role as the principal architect of AI-driven regulatory science. Furthermore, integrating insights from international counterparts will ensure U.S. regulations remain aligned with the evolving global landscape while preserving America's leadership in setting the direction for AI in health care.

Ethical and Patient-Centered Considerations

The NHC underscores that ethical considerations must remain central to any framework governing AI in drug and biologic development. Given the potential for AI to significantly influence patient outcomes and health care delivery, FDA guidance should explicitly address how patient autonomy, informed consent, privacy, and accessibility will be protected.[35] It is essential that patients and clinicians alike understand when and how AI models impact care decisions, thereby maintaining transparency and empowering patient autonomy in clinical settings.

One of the most critical ethical considerations is informed consent. Patients must be clearly informed whenever AI tools significantly impact regulatory or clinical decisions related to their care. To facilitate genuine understanding, the FDA should require sponsors to provide clear, patient-friendly labeling and disclosure materials describing the benefits, risks, and limitations of AI-driven technologies. By explaining when additional human judgment may be necessary and highlighting scenarios where AI recommendations should be interpreted with caution, these disclosures preserve patient autonomy and promote confident decision-making.

Equally crucial are the privacy implications of AI applications that often rely on large volumes of sensitive patient data.[36] The FDA should mandate strict adherence to privacy regulations—requiring robust anonymization, secure data management, and

---

[33] 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.

[34] 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.

[35] 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.

[36] Williamson and Prybutok, "Balancing Privacy and Progress," 675.

cybersecurity safeguards—to protect patients' confidential information. Sponsors must document their compliance with these protocols and inform patients about how their data is collected, used, and shared within AI-driven regulatory processes. By clearly communicating these practices, sponsors can maintain public trust and assure patients that their information is handled responsibly.

Finally, an effective ethical framework must include clearly defined accountability pathways for AI-driven decision-making. The FDA should outline the responsibilities of developers, sponsors, clinicians, and regulators if AI-generated recommendations contribute to adverse outcomes. Transparent documentation and audit trails will enable swift investigations, promote lessons learned from real-world incidents, and reinforce trust in AI-based technologies.[37] By comprehensively integrating these ethical considerations into AI governance, the FDA can ensure that innovation in drug and biologic development remains patient-centered, safe, and aligned with public values.

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,

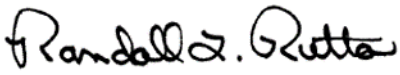

Randall L. Rutta
Chief Executive Officer

[37] 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.