

AI Regulator: Biopharma Misinterprets FDA Guidance on Machine Learning

Written by The Life Science Feed Editorial Team · Published 15 July 2026 · Edited by Mara Voss

The biopharmaceutical industry faces a complex regulatory landscape when integrating artificial intelligence and machine learning (AI/ML) into drug discovery, development, and clinical trial operations. Companies often struggle to understand the Food and Drug Administration's (FDA) expectations for these rapidly evolving technologies. A former FDA AI regulator, now in industry, argues that many biopharma firms are fundamentally misinterpreting the agency's existing guidance, leading to unnecessary caution and missed opportunities for innovation.

KEY TAKEAWAYS

- **The Pivot:** Biopharma's conservative interpretation of FDA AI/ML guidance is slowing innovation, not protecting against regulatory risk.
- **The Data:** The FDA's 2023 guidance on AI/ML in drug development focuses on fit-for-purpose validation, not rigid prescriptive methods.
- **The Action:** Companies should engage with the FDA early and often, demonstrating robust internal validation processes rather than waiting for explicit, exhaustive regulatory frameworks.

Biopharmaceutical companies are increasingly exploring artificial intelligence and machine learning (AI/ML) to accelerate drug discovery, optimize clinical trial design, and enhance patient monitoring. The promise of these technologies is substantial, offering potential efficiencies in identifying novel drug targets, predicting patient responses, and streamlining data analysis. But the path from promise to widespread adoption is fraught with regulatory uncertainty, or at least, perceived uncertainty. Many firms approach AI/ML integration with extreme caution, fearing a misstep that could delay or derail a drug's approval. This hesitancy often stems from a conservative reading of the FDA's evolving guidance on AI/ML in medical products.¹

The core of the issue, according to a former FDA AI regulator, lies in a misunderstanding of the agency's philosophy. The FDA, by its nature, is a science-based regulatory body. Its guidance documents, particularly in emerging fields like AI/ML, aim to establish principles for ensuring safety and efficacy, not to provide a rigid, prescriptive checklist for every conceivable application. Companies, however, frequently interpret these principles as strict rules, leading them to over-engineer solutions or avoid AI/ML altogether in areas where a more pragmatic, risk-based approach would suffice. This often results in a 'wait and see' mentality, where firms delay implementing AI/ML tools until explicit, exhaustive regulatory frameworks are published, which may never fully materialize given the rapid pace of technological advancement.¹

The agency's actual stance

The FDA's approach to AI/ML in drug development is rooted in the concept of 'fit-for-purpose' validation. This means the level of validation required for an AI/ML algorithm should be commensurate with the risk associated with its intended use. For instance, an AI tool used to identify potential drug candidates in preclinical research, far removed from direct patient impact, requires a different level of scrutiny than an algorithm used to determine patient eligibility for a clinical trial or to monitor adverse events in real-time. The agency expects developers to understand the limitations of their models, characterize their performance, and manage potential biases, but it does not dictate a single, universal method for achieving these goals.²

A common misinterpretation centers on the FDA's emphasis on 'locked' versus 'continuously learning' algorithms. While the agency has expressed concerns about algorithms that change in real-time without adequate oversight, particularly in high-risk medical devices, this does not translate to a blanket prohibition on adaptive models in drug development. For many applications, such as optimizing patient recruitment or predicting trial dropouts, an adaptive model, properly validated and controlled, could offer significant advantages. The key is demonstrating robust governance, clear performance metrics, and a well-defined change management process, not avoiding adaptability altogether. The agency is looking for assurance that the model remains safe and effective throughout its lifecycle, regardless of its learning paradigm.³

Biopharma companies also frequently overcomplicate the documentation requirements. While thorough documentation is essential, the FDA is primarily interested in the scientific rigor and clinical relevance of the AI/ML application. This includes clear descriptions of the data used for training and validation, the model architecture, performance metrics (e.g., accuracy, precision, recall, F1-score), and a robust assessment of potential biases and their mitigation strategies. Companies often get bogged down in generating excessive technical documentation that does not directly address these core regulatory concerns, diverting resources from critical validation efforts. A concise, clinically focused summary of the AI/ML tool's capabilities and limitations, supported by robust internal data, is often more effective than an exhaustive technical manual.⁴

Consider the application of AI in clinical trial patient selection. An AI algorithm designed to identify patients who meet complex inclusion/exclusion criteria could significantly reduce screening failures and accelerate enrollment. The FDA would expect to see data demonstrating the algorithm's accuracy in identifying eligible patients (e.g., sensitivity of 92% and specificity of 88% compared to human review), its generalizability across different patient demographics, and a clear process for human oversight and intervention. The agency would not necessarily require a full re-validation for minor updates to the algorithm, provided the developer has a pre-specified plan for managing such changes and demonstrating continued performance. This iterative approach, where validation is proportional to risk and change, is often overlooked by companies seeking a one-time, definitive approval.⁵

Another area of confusion involves the use of real-world data (RWD) and real-world evidence (RWE) in conjunction with AI/ML. The FDA has explicitly encouraged the use of RWD/RWE to support regulatory decision-making, particularly for post-market surveillance and label expansion. AI/ML tools can be instrumental in extracting meaningful insights from vast RWD datasets. But companies often hesitate, citing concerns about data quality, representativeness, and the potential for confounding. While these are valid concerns, the agency's guidance provides frameworks for assessing RWD quality and relevance. The expectation is not for perfect data, but for transparent characterization of data limitations

and appropriate analytical methods to address them. An AI model trained on RWD to predict treatment response, for example, would need to demonstrate its predictive accuracy (e.g., an AUC of 0.85) and provide a clear explanation of the RWD sources and any data imputation strategies.⁶

The former regulator emphasized that early engagement with the FDA is paramount. Companies often wait until late-stage development to discuss their AI/ML strategies, by which point significant resources have been invested in approaches that may not align with regulatory expectations. Proactive communication through pre-submission meetings, particularly for novel applications of AI/ML, allows developers to gain clarity on the agency's perspective and adjust their validation strategies accordingly. This collaborative approach can de-risk development and accelerate the path to market. The agency is generally receptive to innovative approaches, provided they are scientifically sound and adequately validated for their intended use.⁷

But the onus remains on the developer to demonstrate the AI/ML tool's reliability and clinical utility. This includes not only technical performance metrics but also a clear understanding of the clinical context in which the tool will be used. For an AI-powered diagnostic aid, for instance, the FDA would want to see how the tool integrates into existing clinical workflows, its impact on physician decision-making, and ultimately, patient outcomes. Merely demonstrating high technical accuracy in a laboratory setting is insufficient; the clinical relevance and practical applicability must be thoroughly evaluated. This often requires prospective studies or robust retrospective analyses that simulate real-world clinical scenarios.⁸

The challenge for biopharma is not a lack of FDA guidance, but rather a perceived ambiguity that leads to overly conservative interpretations. The agency has published several foundational documents, including the 'Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD)' and the 'Clinical Decision Support Software' guidance. These documents, while not specifically for drug development, establish core principles that are broadly applicable to AI/ML in health care. The 2023 discussion paper on 'AI/ML in Drug Development' further clarifies the agency's thinking, emphasizing a risk-based approach and the importance of transparency and human oversight.^{1,2,3}

The open-ended nature of some FDA guidance is not a flaw; it is a necessity in a rapidly evolving technological domain. The agency cannot predict every future application of AI/ML, nor can it issue prescriptive rules that would quickly become obsolete. Instead, it provides a framework for developers to demonstrate that their AI/ML tools are safe, effective, and fit for their intended purpose. Companies that embrace this principle-based approach, engage proactively with regulators, and focus on robust internal validation will be better positioned to leverage AI/ML for drug development. Those that wait for explicit, exhaustive instructions risk falling behind.⁹

CLINICAL IMPLICATIONS

The biopharma industry's reluctance to fully embrace AI/ML, driven by an overly cautious interpretation of FDA guidance, is a self-imposed handicap. Clinicians stand to benefit significantly from AI-accelerated drug development, potentially gaining access to novel therapies faster and with more precise indications. But if companies are too timid to innovate, these benefits will remain theoretical.

The FDA has consistently signaled a desire for innovation, provided it is underpinned by sound science and robust validation. The agency is not demanding perfection, but rather transparency about model limitations and

a clear demonstration of clinical utility. Companies that engage early and present a well-reasoned, risk-based validation strategy will likely find a receptive audience.

For patients, this regulatory inertia translates into slower progress. AI/ML tools could refine patient selection for trials, identify optimal dosing, and even predict treatment non-response, leading to more effective and personalized care. Delaying the integration of these technologies means delaying these patient-centric advancements.

Ultimately, the industry needs to shift its mindset from seeking explicit permission for every AI/ML application to demonstrating responsible innovation within the existing regulatory framework. The guidance is there; the interpretation needs to catch up to the agency's intent.

EDITORIAL & AI STANDARDS

AI tools are used solely to structure and summarise evidence from peer-reviewed primary sources. No AI-generated conclusions appear in The Life Science Feed without editor verification against the primary source.

REFERENCES

1. US Food and Drug Administration. Artificial Intelligence and Machine Learning in Drug Development. Discussion Paper. May 2023.
2. US Food and Drug Administration. Proposed Regulatory Framework for Modifications to Artificial Intelligence/Machine Learning (AI/ML)-Based Software as a Medical Device (SaMD). Discussion Paper. April 2019.
3. US Food and Drug Administration. Clinical Decision Support Software. Guidance for Industry and Food and Drug Administration Staff. September 2022.
4. European Medicines Agency. Reflection paper on the use of artificial intelligence in the medicinal product lifecycle. Draft. July 2023.
5. US Food and Drug Administration. Real-World Evidence. Guidance for Industry. December 2021.
6. US Food and Drug Administration. Best Practices for Conducting and Reporting Real-World Evidence Studies. Guidance for Industry. December 2021.
7. US Food and Drug Administration. Formal Meetings Between FDA and Sponsors or Applicants of PDUFA Products. Guidance for Industry. March 2015.
8. US Food and Drug Administration. Developing and Submitting Proposed Protocols for Randomizing Patients to a Clinical Trial. Guidance for Industry. October 2023.
9. US Food and Drug Administration. Digital Health Policy Navigator. Accessed January 2024.

LICENCE & RIGHTS

© 2026 The Life Science Feed. All rights reserved. This content is produced for medical education purposes. It may not be reproduced, distributed, or transmitted without prior written permission from The Life Science Feed.

MEDICAL DISCLAIMER

This content is intended for healthcare professionals only and does not constitute medical advice. Clinical decisions must be made by qualified practitioners based on individual patient circumstances.

FOLLOW THE LIFE SCIENCE FEED

Instagram @lifesciencefeed **LinkedIn** linkedin.com/company/thelifesciencefeed

thelifesciencefeed.com